

Spin-Weighted Spherical Harmonics for Polarized Light Transport

SHINYOUNG YI, KAIST, South Korea
DONGGUN KIM, KAIST, South Korea
JIWOONG NA, KAIST, South Korea
XIN TONG, Microsoft Research Asia, China
MIN H. KIM, KAIST, South Korea

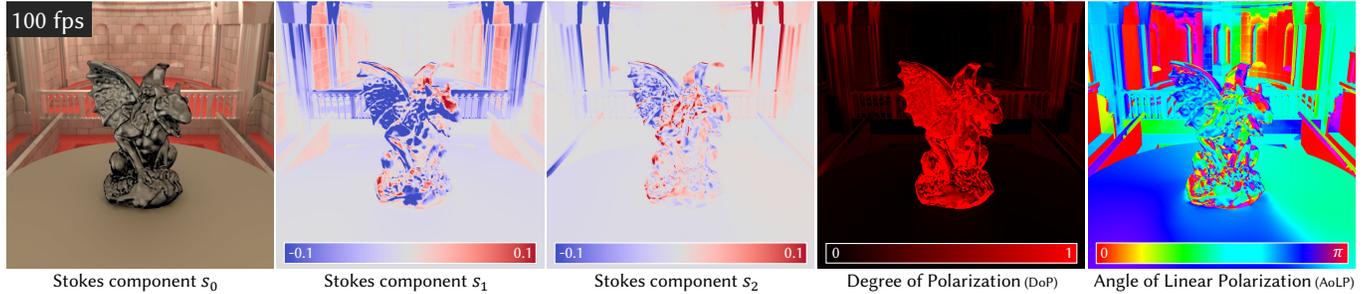


Fig. 1. Real-time rendering of polarized light has been unavailable due to its higher dimensional space of polarized light. We introduce a new frequency-domain analysis of polarized light transport and propose a new method, called *polarized spherical harmonics*, based on the spin-weighted spherical harmonics theory in physics. Our method provides a rotation-invariant representation and spherical convolution of Stokes vector fields, enabling efficient simulation and reproduction of polarized light interactions. We demonstrate the first real-time polarization rendering under polarized environmental illumination through polarized spherical harmonics. Refer to the supplemental video for real-time video results.

The objective of polarization rendering is to simulate the interaction of light with materials exhibiting polarization-dependent behavior. However, integrating polarization into rendering is challenging and increases computational costs significantly. The primary difficulty lies in efficiently modeling and computing the complex reflection phenomena associated with polarized light. Specifically, frequency-domain analysis, essential for efficient environment lighting and storage of complex light interactions, is lacking. To efficiently simulate and reproduce polarized light interactions using frequency-domain techniques, we address the challenge of maintaining continuity in polarized light transport represented by Stokes vectors within angular domains. The conventional spherical harmonics method cannot effectively handle continuity and rotation invariance for Stokes vectors. To overcome this, we develop a new method called polarized spherical harmonics (PSH) based on the spin-weighted spherical harmonics theory. Our method provides a rotation-invariant representation of Stokes vector fields. Furthermore, we introduce frequency domain formulations of polarized rendering equations and spherical convolution based on PSH. We first define spherical convolution on Stokes vector fields in the angular domain, and it also provides efficient computation of polarized light transport, nearly on an entry-wise product in the frequency domain. Our frequency domain formulation, including spherical convolution, led to the development of the

first real-time polarization rendering technique under polarized environmental illumination, named precomputed polarized radiance transfer, using our polarized spherical harmonics. Results demonstrate that our method can effectively and accurately simulate and reproduce polarized light interactions in complex reflection phenomena, including polarized environmental illumination and soft shadows.

CCS Concepts: • **Mathematics of computing** → *Functional analysis*; • **Computing methodologies** → *Rendering; Appearance and texture representations*.

Additional Key Words and Phrases: spherical harmonics, polarized rendering, polarimetric imaging, polarimetric appearance, theory of light transport, basis function

ACM Reference Format:

Shinyoung Yi, Donggun Kim, Jiwoong Na, Xin Tong, and Min H. Kim. 2024. Spin-Weighted Spherical Harmonics for Polarized Light Transport. *ACM Trans. Graph.* 43, 4, Article 127 (July 2024), 24 pages. <https://doi.org/10.1145/3658139>

1 INTRODUCTION

Polarization, imperceptible to the human eye, offers a wealth of auxiliary information about an object's shape and physically meaningful material characteristics. Consequently, polarization has been increasingly exploited in both the fields of computer graphics and vision for tasks of geometry modeling and appearance acquisition in recent years. It has shown extensive applications, including multispectral ellipsometry to obtain a polarimetric bidirectional reflectance distribution function (pBRDF) dataset [Baek et al. 2020], inverse rendering for acquiring polarimetric appearance and geometry [Baek et al. 2018; Hwang et al. 2022], and physically-based polarization rendering through the synthesis and analysis of polarized light transport [Jakob et al. 2022].

Authors' addresses: Shinyoung Yi, KAIST, South Korea, syyi@vclab.kaist.ac.kr; Donggun Kim, KAIST, South Korea, dgkim@vclab.kaist.ac.kr; Jiwoong Na, KAIST, South Korea, jwna@vclab.kaist.ac.kr; Xin Tong, Microsoft Research Asia, China, xtong@microsoft.com; Min H. Kim, KAIST, South Korea, minhkim@vclab.kaist.ac.kr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 0730-0301/2024/7-ART127

<https://doi.org/10.1145/3658139>

Polarization rendering, as opposed to traditional rendering that calculates light intensity, necessitates additional computation of further information. First, polarization rendering simulates a four-dimensional vector that consists of light intensity, two linear polarizations in horizontal/vertical and diagonal/antidiagonal directions, and circular polarization, commonly represented as a Stokes vector. This fully characterizes the polarization state of light as it travels along a ray. Second, different from the conventional vector, the Stokes vector physically quantifies the sinusoidal oscillation of light waves. Consequently, when the coordinate system rotates, the elements of the Stokes vector change at a rate twice as fast as the components of a conventional vector change. It is critical to account for this fact when transforming a Stokes vector for polarized light simulation. Third, the unique features of polarization rendering not only escalate the computational expenses significantly but also necessitate special transformation. The Stokes vector is determined within a local coordinate system of the progressing ray. Therefore, the polarimetric reflectance function must be configured in accordance with the coordinate systems of the incident and exitant Stokes vectors. This implies that additional coordinate conversions are required for polarization rendering. Lastly, the formulation of the polarimetric reflectance must consider both the incident and exitant Stokes vectors. This process is often presented as a Mueller matrix, a structure of sixteen components arranged in a four-by-four matrix format, of which components change depending on the incident/exitant light angles. The complexity of this matrix makes it sixteen times larger than a scalar value used in conventional reflectance functions, significantly increasing the computational cost of ray samples in polarization rendering [Wilkie and Weidlich 2012].

Due to the unique attributes of Stokes vectors in the representation of light polarization, conventional frequency-domain analysis of environment lighting [Ramamoorthi and Hanrahan 2001b], does not guarantee smooth transformations and invariance under rotations in polarization rendering. Moreover, the process of rendering specular reflections using spherical harmonics (SH), which necessitates spherical convolution [Sloan et al. 2002], functions efficiently for scalar fields that do not change under rotation. However, these conventional spherical harmonics cannot be applied for the transportation of polarized light. This is because, in polarization rendering, light is represented not as a scalar intensity but as a Stokes vector. In this work, we focus on addressing the following two main challenges with the goal of facilitating real-time polarization rendering via frequency-domain analysis.

Rotation invariance. Different from the conventional SH-based rendering that computes light intensity as a scalar value on a sphere, polarization rendering needs to simulate a Stokes vector field on a sphere. However, dealing with Stokes vector fields using conventional basis functions, including SH, which are designed for scalar fields, results in a *singularity problem*. This is commonly known as the Hairy Ball Theorem [Nash and Sen 1983]. Representing a Stokes vector field using conventional SH requires separating it into four scalar fields and one frame field, which assigns local frames for tangent planes for each point on a sphere. The Hairy Ball Theorem implies we cannot assign a smooth and consistent direction of unit tangent vectors at every point on the sphere without at least

one singular point, so that the resulting Stokes vector field combined with continuous scalar basis functions should have a singular point. This is a critical problem for rotation transformation in light transport. For instance, suppose we want to rotate Stokes vector fields by transforming basis functions. The transformation cannot guarantee the rotation invariance because this operator to the basis functions will create another singularity point in another location on the sphere.

To address the issue of rotation invariance in polarization rendering, we introduce a new method utilizing a spin-weighted spherical harmonics (SWSH) theory [Scanio 1977] in physics. These SWSH serve as orthonormal basis functions, and they can be classified based on how different spin fields behave on a sphere. The spin-0 spherical harmonics, which are equivalent to traditional spherical harmonics, represent scalar fields that remain unchanged under the rotation of local frames. On the other hand, spin-1 spherical harmonics represent vector fields on a sphere. These can transform under rotations in the same way as a typical vector, which indicates they possess a certain spin orientation. Finally, spin-2 spherical harmonics (S2SH) represent fields of such quantities as neither scalar nor ordinary vectors, which are characterized as multiple directions associated with each point, mirroring the properties of Stokes vector fields. To resolve the fundamental rotation invariance problem in polarization rendering, we employ S2SH, which handles Stokes vector fields by extending the domain of the basis function space from the sphere to the frame space while maintaining a spin-2 constraint, rather than improperly separating them into scalars and frame fields with singularity. This approach paves the way for real-time polarization rendering.

Spherical convolution. To accomplish efficient real-time rendering of specular reflection in conventional rendering, SH-based rendering has utilized the scalar spherical convolution of light intensity [Sloan et al. 2002]. For real-time polarization rendering to be feasible, it is also crucial to establish a spherical convolution of Stokes vector fields. In polarization rendering, the input and output for spherical convolution are represented as Stokes vectors. Therefore, the convolution kernel needs to be defined as a Mueller matrix. However, we observe that the Mueller matrix domain should have only one degree of freedom of the zenith angle of the kernel in relation to the zonal axis of spherical harmonics when generalizing conventional spherical convolution as rotation equivariant linear operators. The spherical convolution of a Mueller matrix to Stokes vector fields, which we formulate in this paper, has not yet been addressed in the field of computer graphics research.

We, therefore, introduce a new frequency-domain method for spherical convolution for Stokes vector fields. This method allows for the efficient yet precise convolution of approximated Stokes vectors, thereby enabling real-time rendering of polarized light. Our approach, which is based on polarized spherical harmonics (PSH), facilitates efficient computation, operating nearly on an element-by-element product basis. To this end, we jointly combine the spin-0 and spin-2 cases of spin-weighted spherical harmonics, incorporating a new theory concerning the frequency-domain analysis of pBRDF and spherical convolution of Stokes vector fields.

Further, we demonstrate a real-time technique for polarization rendering, the so-called precomputed polarized radiance transfer (PPRT), using our polarized spherical harmonics. See Figure 1 for an example. Our proposed method can efficiently and accurately simulate and replicate the approximated interactions of polarized light in complex reflection phenomena, including polarized environmental illumination and soft shadows.

2 RELATED WORK

2.1 Spherical Harmonics

A frequency-domain framework using spherical harmonics is introduced by Ramamoorthi and Hanrahan [2001a; 2001b] to computer graphics community. They represent environment maps [Ramamoorthi and Hanrahan 2001a] into SH coefficients and render environment map lighting by the product of coefficient vectors. Extending SH coefficients of diffuse albedo to store radiance self-transfer, including self-shadow and interreflection, their framework has been extended to precomputed radiance transfer (PRT) [Sloan et al. 2002]. The PRT method has various extensions which deal with dynamic shadow [Zhou et al. 2005], deformable objects [Sloan et al. 2005], and polygonal lights [Wang and Ramamoorthi 2018; Wu et al. 2020]. Benefits of some of these methods come from not only algorithmic enhancement but also analytic integrals related to SH, such as triple product [Zhou et al. 2005] and integrals on spherical polygons [Wang and Ramamoorthi 2018]. We refer to Kautz et al. [2005] and Ramamoorthi et al. [2009] for more history and overview of the field of precomputation-based rendering. Note that not only real-time rendering methods, application of SH to rendering also include physically based ray tracing [Belcour et al. 2018], which uses SH products as control variates and inverse rendering of reflectance [Ramamoorthi and Hanrahan 2001b] which projects BRDF and normal vectors into SH coefficients.

Other bases for spherical functions, including the Haar wavelet [Lessig and Fiume 2008; Ng et al. 2003, 2004], spherical Gaussians [Ritschel et al. 2012], and neural bases [Xu et al. 2022] have been discussed. Still, only spherical harmonics simultaneously hold orthonormality, rotation invariance, and a coefficient-wise product of spherical convolution. There is another recent approach to learning basis functions on the sphere rather than defining analytically by Xu et al. [2022], but their work produces no genuine basis that should satisfy linearity.

While SH provides a wide range of applications in computer graphics, as discussed above, there has been no extension of any of these methods to polarized light transport due to the difficulty of the continuity structure of Stokes vector fields.

2.2 Polarization

Polarization has played an important role in computer graphics. For example, polarized illumination enhances the reconstruction quality of 3D geometry and reflectance [Ba et al. 2020; Ghosh et al. 2011; Kadambi et al. 2015]. In addition, rendering [Jarabo and Arellano 2018; Mojzík et al. 2016] and reconstructing in both explicit geometry [Baek et al. 2018; Hwang et al. 2022] and radiance fields [Kim et al. 2023] polarized quantities themselves have also been investigated recently. These problems handle polarized appearance, which captures

what traditional scalar intensity-based appearance has not done and has been addressed as challenging problems due to more parameters and unconventional coordinate conversion problems. However, no frequency-domain methods have been developed. Jarabo and Arellano [2018]; Mojzík et al. [2016] introduce polarized ray tracing methods that consider the light source and material appearance as Stokes vectors and Mueller matrices, respectively, but there are no precomputed methods through basis functions that achieve real-time performance. Baek et al. [2020] captured image-based pBRDF datasets, but there is still a lack of methods of how to render their materials in runtime efficiently. In this context, we propose a new frequency-domain framework of polarized light transport, which implies polarized precomputed rendering, so that our novel rendering method achieves real-time performance and provides a novel way to render Baek et al. [2020]’s data-based pBRDF.

Certain studies have utilized polarized gradient illumination to capture the appearance of objects [Ghosh et al. 2009, 2011; Ma et al. 2007], which is related to spherical harmonics up to order 2. The utilization of polarized light in these studies is specific to scenarios where it is necessary to separate two scalar fields of diffuse and specular reflection. However, bases of Stokes vector fields have not been addressed in these studies.

The works mentioned above, including this one, use Mueller calculus formulations to deal with polarized light. However, physical light transport methods, such as those presented in recent works [Steinberg et al. 2022; Steinberg and Yan 2021a,b], have introduced a generalized Stokes parameters formulation based on optical coherence theory. This formulation combines the strengths of both Mueller and Jones calculus. However, it does not address the challenges in the angular domain, and its contributions are not relevant to our current scope.

For more concepts, history, and applications in computer graphics of polarization, we refer to Collett [2005], Wilkie and Weidlich [2012], and Baek et al. [2023].

2.3 Spin-Weighted Spherical Harmonics

Spin-weighted spherical harmonics theory is originally introduced by Goldberg et al. [1967]; Newman and Penrose [1966] to handle the symmetry of gravitational radiation in physics. Zaldarriaga and Seljak [1997] point out that spin-2 SH can encode the all-sky information of polarized light to the frequency domain in the context of the cosmic microwave background. Rotation invariance and coefficient rotations of SWSH are shown by Boyle [2013].

Note that SWSH has also been referred to as *generalized spherical harmonics* in some literature [Garcia and Siewert 1986; Keegstra et al. 1997; Kuščer and Ribarič 1959; Phinney and BurrIDGE 1973], and the relation between these names is pointed out by Rossetto [2009].

While SWSH formulation of Stokes vector fields already exists, to the best of our knowledge, we first formulate linear operators on Stokes vectors, including pBRDF, into SWSH coefficients.

Zaldarriaga and Seljak [1997] and Ng and Liu [1999] establish the SWSH formulation of the *correlation* operation between two Stokes vector fields in the perspective to analyze statistics of given data. While the correlations have some similarities to *convolutions*, these

are inherently different operations in terms of types of inputs and outputs. We focus on the convolution operation from the perspective of image processing and computer graphics, especially PRT.

Spherical convolution of Stokes vector fields has been discussed in Garcia and Siewert [1986], Ng and Liu [1999], and Tapimo et al. [2018]. However, their formulations are subsets of our formulation of polarized spherical convolution. Specifically, their polarized convolution kernels have one degree of freedom (DoF) [Ng and Liu 1999] or six DoF [Garcia and Siewert 1986; Tapimo et al. 2018] for each frequency band, while ours has 16 DoF. Based on this generalization, we first discover that polarized spherical convolution is equivalent to rotation equivariance linear operators on Stokes vector fields, with a proper sense of such linearity.

We refer to Section 8.3 for a more technical description of our novelty against existing work on SWSH.

3 OVERVIEW

The following is a brief outline of our paper's organization. In Section 4, we provide the theoretical foundations of traditional spherical harmonics, spherical convolution, and polarization of light in Mueller calculus. This section is included for the sake of readability, but expert readers may skip it, while Section 4.2.1 gives a brief introduction to the mathematical notations used in this paper. It will help the readers to understand the mathematical concepts presented in the paper. In Section 5, we discuss the challenges of applying existing spherical harmonics to Stokes vector fields. Our main method is presented in Section 6, which consists of the polarized spherical harmonics theory (Section 6.2) and polarized spherical convolution (Section 6.4). In Section 7, we demonstrate the first real-time polarized rendering method, followed by a discussion in Section 8 and a conclusion in Section 9. Tables 1 and 2 provide notations, symbols, and operators used in this paper. We also make our code available on our project website (<https://vclab.kaist.ac.kr/siggraph2024/>), which includes a step-by-step tutorial to help understand various quantities and equations.

4 BACKGROUND

4.1 Spherical Harmonics

This subsection briefly reviews the definition and core properties of spherical harmonics. In Supplemental Sections 1.4 and 2, we additionally provide a general theory of function spaces and bottom-up mathematical description of SH, including how some properties of SH are inherited from the general theory.

Spherical harmonics are spherical functions $Y_{lm} \in \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$, where $\mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C}) := \{f: \hat{\mathbb{S}}^2 \rightarrow \mathbb{C}\}$, which can be evaluated in spherical coordinates (θ, ϕ) as follows:

$$Y_{lm}(\theta, \phi) = A_{lm} P_l^m(\cos \theta) e^{im\phi}, \quad (1)$$

where $A_{lm} = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}}$ and P_l^m denotes the associated Legendre function of order l and degree m (Supplemental Equation (31b)). $\{Y_{lm} \mid (l, m) \in \mathbb{Z}^2, |m| \leq l\}$ is an orthonormal basis of $\mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$. In other words, any spherical function $f \in \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$ is equal to an

Table 1. Lists of notations and symbols used in this paper.

Notation		
$\mathbf{x}, \mathbf{y}, \dots \in \mathbb{R}^N$		Numeric N -dimensional vectors, lowercase Latin letters with boldface (including Stokes component vector)
$\mathbf{A}, \mathbf{B}, \dots \in \mathbb{R}^{M \times N}$		Numeric $M \times N$ matrices, uppercase Latin letters with boldface (including Mueller matrices)
$\vec{\mathbf{x}}, \vec{\mathbf{y}}, \dots \in \vec{\mathbb{R}}^N$		Geometric N -dimensional vectors, lowercase Latin letters accented single side arrow
$\vec{\mathbf{A}}, \vec{\mathbf{B}}, \dots \in \vec{\mathbb{R}}^{M \times N}$		Geometric $M \times N$ matrices, uppercase Latin letters accented single side arrow
$\overleftrightarrow{\mathbf{x}}, \overleftrightarrow{\mathbf{y}}, \dots \in \mathcal{S}_{\hat{\omega}}$		Stokes vectors (geometric), lowercase Latin letters accented both side arrow
$\overleftrightarrow{\mathbf{A}}, \overleftrightarrow{\mathbf{B}}, \dots \in \mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o}$		Mueller transforms (geometric), uppercase Latin letters accented both side arrow
Symbol		
$\hat{\omega} \in \hat{\mathbb{S}}^2$		Directions (unit vector), where $\hat{\mathbb{S}}^2$ is unit sphere
$\vec{\mathbf{F}} \in \vec{\mathbb{F}}^3$		Orthonormal frames in 3D, uppercase Latin letter F with boldface accented single side arrow
$\mathbf{R} \in \text{SO}(3)$		Numeric 3D rotation matrices
$\vec{\mathbf{R}} \in \vec{\text{SO}}(3)$		Geometric 3D rotation transforms
$\mathcal{F}(X, Y)$		Function space from X into Y , for any sets X and Y
$\mathcal{S}_{\hat{\omega}}$		Stokes space: set of all Stokes vectors of a ray along direction $\hat{\omega}$
$\mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o}$		Mueller space from $\mathcal{S}_{\hat{\omega}_i}$ to $\mathcal{S}_{\hat{\omega}_o}$
Operator		
$[\mathbf{s}]_{\vec{\mathbf{F}}} = \overleftrightarrow{\mathbf{s}}$		Stokes component vector \mathbf{s} to Stokes vector $\overleftrightarrow{\mathbf{s}}$ w.r.t. frame $\vec{\mathbf{F}}$
$[\overleftrightarrow{\mathbf{s}}]_{\vec{\mathbf{F}}} = \mathbf{s}$		Stokes vector $\overleftrightarrow{\mathbf{s}}$ to Stokes component vector \mathbf{s} w.r.t. a frame $\vec{\mathbf{F}}$
$[\mathbf{M}]_{\vec{\mathbf{F}}_1 \rightarrow \vec{\mathbf{F}}_2} = \vec{\mathbf{M}}$		Mueller matrix \mathbf{M} to geometric Mueller transform $\vec{\mathbf{M}}$ w.r.t. frames $\vec{\mathbf{F}}_1, \vec{\mathbf{F}}_2$
$[\vec{\mathbf{M}}]_{\vec{\mathbf{F}}_1 \rightarrow \vec{\mathbf{F}}_2} = \mathbf{M}$		Mueller transform $\vec{\mathbf{M}}$ to numeric Mueller matrix \mathbf{M} w.r.t. frames $\vec{\mathbf{F}}_1, \vec{\mathbf{F}}_2$
$z^* = x - yi$		Complex conjugation of $z = x + yi \in \mathbb{C}$
$\Re z, \Im z = x, y$		Real and imaginary parts of $z = x + yi \in \mathbb{C}$
$\mathbb{R}^2(z), \mathbb{C}([x, y]^T)$		Conversion between complex number $z = x + yi \in \mathbb{C}$ and $[x, y]^T \in \mathbb{R}^2$ (Eq. (48))
$\mathbb{R}^{2 \times 2}(z)$		Eq. (37), Conversion from complex number to 2D real numeric matrix
$\mathbb{C}_{\text{iso}}(\mathbf{M}), \mathbb{C}_{\text{conj}}(\mathbf{M})$		Eq. (44), Conversion from 2×2 real matrix \mathbf{M} to two complex numbers respectively

Table 2. List of rotations and inner products in various quantities.

Symbol	Operand	Eq. num.
$\vec{\mathbf{R}}\vec{\mathbf{x}}, \vec{\mathbf{R}}\vec{\mathbf{F}}$	Geometric vectors $\vec{\mathbf{x}} \in \vec{\mathbb{R}}^3$, and frames $\vec{\mathbf{F}} \in \vec{\mathbb{F}}^3$	
$\vec{\mathbf{R}}\vec{\mathbf{s}}\vec{\mathbf{s}}$	Stokes vectors $\vec{\mathbf{s}} \in \mathcal{S}$	Eq. (19)
$\vec{\mathbf{R}}_{\mathcal{M}}[\vec{\mathbf{M}}]$	Mueller transforms $\vec{\mathbf{M}} \in \mathcal{M}$	Eq. (59)
$\vec{\mathbf{R}}_{\mathcal{F}}[f](\hat{\omega})$	Scalar fields $f: \hat{\mathbb{S}}^2 \rightarrow \mathbb{C}$	Eq. (8)
$\vec{\mathbf{R}}_{\mathcal{F}}[\vec{f}](\hat{\omega})$	Stokes vector fields $\vec{f}: \hat{\mathbb{S}}^2 \rightarrow \mathcal{S}_{\hat{\omega}}$	Eq. (25)
$\langle \overleftrightarrow{\mathbf{s}}, \overleftrightarrow{\mathbf{t}} \rangle_{\mathcal{S}}$	Stokes vectors $\overleftrightarrow{\mathbf{s}}, \overleftrightarrow{\mathbf{t}} \in \mathcal{S}_{\hat{\omega}}$ (identical direction)	Eq. (18)
$\langle f, g \rangle_{\mathcal{F}}$	Scalar fields $f, g: \hat{\mathbb{S}}^2 \rightarrow \mathbb{C}$	Eq. (3)
$\langle \overleftrightarrow{f}, \overleftrightarrow{g} \rangle_{\mathcal{F}}$	Stokes vector fields $\overleftrightarrow{f}, \overleftrightarrow{g}: \hat{\mathbb{S}}^2 \rightarrow \mathcal{S}_{\hat{\omega}}$	Eq. (24)

infinite number of the linear combination of SH as

$$f = \sum_{l=0}^{\infty} \sum_{m=-l}^l f_{lm} Y_{lm}, \quad (2)$$

and the coefficient f_{lm} is computed as

$$f_{lm} = \langle Y_{lm}, f \rangle_{\mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})} := \int_{\hat{\mathbb{S}}^2} Y_{lm}^*(\hat{\omega}) f(\hat{\omega}) d\hat{\omega}, \quad (3)$$

where the integration over the sphere $\hat{\mathbb{S}}^2$ is defined with the solid angle measure $d\hat{\omega} = \sin\theta d\theta d\phi$, and z^* indicates the complex conjugate of an arbitrary $z \in \mathbb{C}$. Note that when the domain of an inner product is clear in context, we just write the inner product as $\langle Y_{lm}, f \rangle_{\mathcal{F}}$ for the sake of simplicity.

From Equation (3), a numeric vector which consists of such f_{lm} called *coefficient vector*, which encodes frequency-domain information of the spherical function f . While an infinite dimensional coefficient vector $[f_{00}, f_{1,-1}, f_{10}, f_{11}, \dots]^T$ represents continuously defined f without loss of information, we can take the projection of f on SH up to order l_{\max} , and store it into a finite coefficient vector $[f_{00}, \dots, f_{l_{\max}, l_{\max}}]^T$ of $O(l_{\max}^2)$ entries.

4.1.1 Coefficient matrix and radiance transfer. In rendering pipelines or other frequency-domain analysis, many methods can be represented as functions of spherical functions (linear operator). SH also represents linear operators on spherical functions into discrete coefficients, called *coefficient matrix*. Suppose that $T: \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C}) \rightarrow \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$ be a linear operator on spherical functions. Similar to Equation (3), the linear operator T can be represented by discrete SH coefficients $T_{l_o m_o, l_i m_i}$ as

$$T_{l_o m_o, l_i m_i} = \langle Y_{l_o m_o}, T[Y_{l_i m_i}] \rangle_{\mathcal{F}}, \quad (4)$$

where the subscript i and o in l and m stands for input and output. The evaluation of T at a function $f \in \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$ can be considered as a matrix-vector multiplication in the SH coefficient space as

$$\langle Y_{l_o m_o}, T[f] \rangle_{\mathcal{F}} = \sum_{l_i, m_i} T_{l_o m_o, l_i m_i} f_{l_i m_i}, \quad (5)$$

where $\langle Y_{l_o m_o}, T[f] \rangle_{\mathcal{F}}$ is the coefficient of the output function $T[f]$, obtained by Equation (3).

In computer graphics, a BRDF¹ $\rho: \hat{\mathbb{S}}^2 \times \hat{\mathbb{S}}^2 \rightarrow \mathbb{R}$ can be characterized by a linear operator $\rho_{\mathcal{F}}: \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C}) \rightarrow \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$ which acts as the rendering equation:

$$\rho_{\mathcal{F}}[L^{\text{in}}](\hat{\omega}_o) = \int_{\hat{\mathbb{S}}^2} \rho(\hat{\omega}_i, \hat{\omega}_o) L^{\text{in}}(\hat{\omega}_i) d\hat{\omega}_i, \quad (6)$$

for any incident radiance function of a direction L^{in} . Taking the matrix product of the SH coefficient matrix of $\rho_{\mathcal{F}}$, also called the radiance transfer matrix, and the coefficient vector of L^{in} is the core operation in the efficient environment lighting [Ramamoorthi and Hanrahan 2001b] and PRT [Sloan et al. 2002] methods.

Moreover, the isotropy constraint of the BRDF (in general, an azimuthal symmetric operator) yields increasing the sparsity of SH

¹We consider a cosine-weighted BRDF which already contains the term $|\hat{n} \cdot \hat{\omega}_i|$.

coefficients, which can be written with fewer indices as [Ramamoorthi and Hanrahan 2001b, 2002]

$$\rho_{l_o m_o, l_i m_i} = \delta_{m_o m_i} \rho_{l_o l_i m_i}, \quad (7)$$

where $\delta_{m_o m_i}$ indicates the Kronecker delta. Note that while a general linear operator requires $O(l_{\max}^4)$ SH coefficients in Equation (4), azimuthal symmetry described in Equation (7) reduces the number of coefficients to $O(l_{\max}^3)$.

4.1.2 Rotation invariance. One of the most important properties of SH is rotation invariance, which allows us to efficiently convert SH coefficients with respect to another frame without loss of information.

A rotation can be considered as a linear operator. Given rotation transform $\vec{R} \in \overline{SO}(3)$, the rotation on spherical functions rather than vectors is denoted by $\vec{R}_{\mathcal{F}}: \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C}) \rightarrow \mathcal{F}(\hat{\mathbb{S}}^2, \mathbb{C})$ and acts as

$$\vec{R}_{\mathcal{F}}[f](\hat{\omega}) = f(\vec{R}^{-1}\hat{\omega}). \quad (8)$$

The coefficient matrix of the rotation $\vec{R}_{\mathcal{F}}$ is obtained from Equation (4). It can be written with the Kronecker delta and a special function $D_{mm'}^l$, which is called a Wigner D-function as

$$\langle Y_{lm}, \vec{R}_{\mathcal{F}}[Y_{l'm'}] \rangle_{\mathcal{F}} = \delta_{ll'} D_{mm'}^l(\vec{R}). \quad (9)$$

The rotation invariance of SH is stated as the block diagonal constraint of the coefficient matrices of rotations due to the term $\delta_{ll'}$ in Equation (9), which is also visualized in Figure 10(a). This property also implies that we can commute the SH projection of a function and a rotation without loss of information. We refer to Supplemental Figures 3(a), 4, and 5 in Supplemental Section 2.3 for further description and visualization.

4.1.3 Spherical convolution. Spherical convolution is defined for a kernel $k: [0, \pi] \rightarrow \mathbb{C}$, a spherical function with azimuthal symmetry $k(\theta, \phi) = k(\theta)$, and any spherical function f as follows.

$$k * f(\hat{\omega}) = \int_{\hat{\mathbb{S}}^2} k(\cos^{-1}(\hat{\omega} \cdot \hat{\omega}')) f(\hat{\omega}') d\hat{\omega}'. \quad (10)$$

The definition of spherical convolution in Equation (10) is determined from its important properties, linearity, and rotation equivariance for f . Conversely, it is known that a rotation equivariant linear operator on spherical functions is equivalent to a convolution with some kernel k .

SH provide an efficient computation of this convolution. The SH coefficients of the convolution result, $f'_{lm} := \langle Y_{lm}, k * f \rangle_{\mathcal{F}}$ is evaluated by

$$f'_{lm} = \sqrt{\frac{4\pi}{2l+1}} k_{l0} f_{lm}, \quad (11)$$

which is just an element-wise product of the kernel and the input function in SH coefficients. Note that it is analogous to the convolution theorem of the Fourier transform in Euclidean domains.

In a rendering context, a BRDF is encoded to a coefficient matrix with $O(l_{\max}^4)$ space complexity. However, assuming Phong-like

BRDFs with rotation equivariance whose reflected lobe just rotates as the incident ray rotates, a BRDF can be represented as a spherical convolution kernel [Sloan et al. 2002], which can lead to more efficient computation from its $O(l_{\max})$ sparsity.

4.1.4 Real and complex SH. While SH defined in Equation (1) are complex-valued functions, real-SH Y_{lm}^R are also defined as follows:

$$Y_{lm}^R = \begin{cases} \sqrt{2}\Re Y_{lm} = \frac{1}{\sqrt{2}}(Y_{lm} + (-1)^m Y_{l,-m}) & m > 0 \\ Y_{lm} & m = 0 \\ \sqrt{2}\Im Y_{l|m|} = \frac{i}{\sqrt{2}}((-1)^m Y_{lm} - Y_{l,-m}) & m < 0 \end{cases} \quad (12)$$

We will sometimes call Y_{lm} defined in Equation (1) *complex* SH to distinguish from real ones. Note that the real SH also satisfy orthonormality and rotation invariance, but they always convert real-valued functions into real-valued coefficients.

For the rotation transform of real SH coefficients, it can be written similarly to complex SH as

$$\left\langle Y_{lm}^R, \vec{R}_{\mathcal{F}} \left[Y_{l'm'}^R \right] \right\rangle_{\mathcal{F}} = \delta_{ll'} D_{mm'}^{l,R}(\vec{R}), \quad (13)$$

where $D_{mm'}^{l,R}$ is named real Wigner D-functions, and it can be evaluated simply as a linear combination of complex-valued $D_{\pm m, \pm m'}^l$ (Supplemental Equation (60)). See Supplemental Section 2.4 for more details.

For computational efficiency, most existing computer graphics works use real SH. However, both real and complex SH should be considered for our polarized SH, which will be introduced in Section 6.

4.2 Polarization and Mueller Calculus

Given a local frame $\vec{F} = [\hat{x}, \hat{y}, \hat{z}]$, the intensity of a polarized ray along the propagation direction \hat{z} is characterized by the four Stokes parameters $\mathbf{s} = [s_0, s_1, s_2, s_3]^T$. Here, each component s_0 to s_3 indicates total intensity, linear polarization in horizontal/vertical direction, linear polarization in diagonal/anti-diagonal direction, and circular polarization, respectively. We refer novice readers to Supplemental Section 3.1 for more introduction.

When taking another local frame $\vec{F}' = \vec{R}_z(\vartheta) \vec{F}$, obtained by rotating \vec{F} by ϑ along its z axis, the Stokes parameters with respect to the new frame \vec{F}' is evaluated as

$$\mathbf{s}' = \mathbf{C}_{\vec{F} \rightarrow \vec{F}'} \mathbf{s} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\vartheta & \sin 2\vartheta & 0 \\ 0 & -\sin 2\vartheta & \cos 2\vartheta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{s}. \quad (14)$$

We can observe here that s_0 and s_3 behave as *scalars*, which are measured independent of local frames. On the other hand, s_1 and s_2 are neither scalars nor coordinates of an ordinary vector, which must have ϑ rather than 2ϑ in Equation (14). This twice rotation property of s_1 and s_2 under coordinate conversion will be dealt as *spin-2 functions* in Section 5.

4.2.1 Stokes vectors in numeric vs. geometric quantities. As discussed before, dealing with polarized radiance needs careful attention for whether focusing on a ray itself as a physical object or

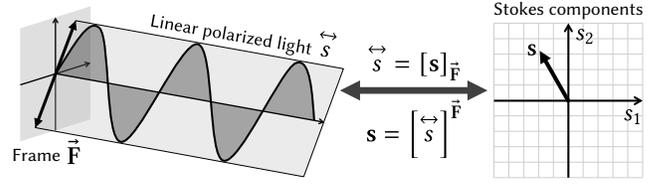


Fig. 2. Intensity of a polarized ray visualized in the left is characterized by a Stokes vector \vec{s} . While \vec{s} is defined without any measurement frame, it can be measured into a Stokes component vector \mathbf{s} under such a frame.

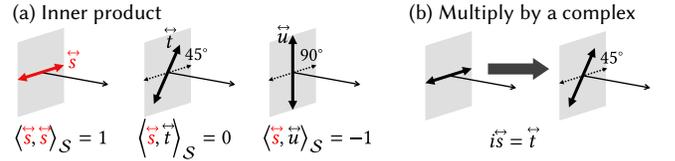


Fig. 3. Additional basic operations on Stokes vectors are defined in (a) Equation (18) and (b) Equation (20).

Stokes parameter values s_0, \dots, s_3 , only defined relative to a measurement frame associated with the ray. Note that we distinguish *numeric* and *geometric* quantities in this paper. Due to the twice rotation property described in Equation (14), the polarized intensity of a ray should be considered as a novel type of geometric quantity, named *Stokes vector* and denoted by \vec{s} . Note that \mathbf{s} and \mathbf{s}' in Equation (14) are numeric quantities and not geometric ones themselves since they depend on observing local frames. Combining data of \mathbf{s} and \vec{F} yields the geometric quantity \vec{s} , but it is not a matrix-vector product as ordinary vectors. Thus, we write it in a novel notation as

$$\vec{s} = [\mathbf{s}]_{\vec{F}} = [\mathbf{s}']_{\vec{F}'}. \quad (15)$$

In addition, we call such numeric vector \mathbf{s} , the Stokes parameters observed under a certain frame, as *Stokes component vector*². We also define the notation that evaluates the Stokes component vector of a given Stokes vector and the frame as

$$\mathbf{s} = [\vec{s}]_{\vec{F}}. \quad (16)$$

Figure 2 visualizes it where the two-sided arrow in the left indicates the actual oscillation direction of a polarized ray characterized by a Stokes vector and the right plot shows the Stokes component vector under a local frame.

We also denote $\mathcal{S}_{\hat{\omega}} = \{[\mathbf{s}]_{\vec{F}} \mid \vec{F} \in \mathbb{F}^3, \vec{F}[:, 3] = \hat{\omega}\}$ as the *Stokes space*, the set of all Stokes vector of rays along direction $\hat{\omega}$, where $\vec{F}[:, 1]$, $\vec{F}[:, 2]$, and $\vec{F}[:, 3]$ indicate the local x , y , and z axes of given frame \vec{F} , respectively.

Stokes vector operations. Binary operations on two Stokes vectors $\vec{s} = [\mathbf{s}]_{\vec{F}_1}$ and $\vec{t} = [\mathbf{t}]_{\vec{F}_2}$ are defined only if they belong to the identical Stokes space. i.e., the ray directions are same ($\vec{F}_1[:, 3] = \vec{F}_2[:, 3]$). If so, the addition and the inner product are defined by converting the Stokes vectors to the same frame as

²Note that we try to distinguish terminologies *Stokes vectors* and *Stokes components* as geometric and numeric quantities, respectively, so this distinction is not common in other literature. See also Supplemental Figure 1.

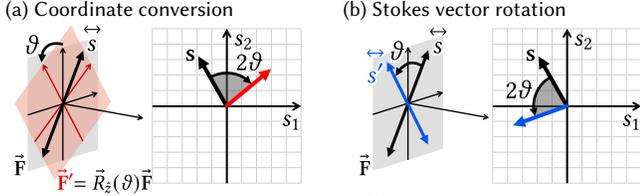


Fig. 4. (a) When we fix the Stokes vector \vec{s} and rotate the frame by ϑ , the (numeric) Stokes components of \vec{s} rotate by -2ϑ . (b) Rotating the (geometric) Stokes vector itself by ϑ is equivalent to rotating its Stokes components by 2ϑ with a fixed frame.

$$\vec{s} + \vec{t} := \left[\mathbf{s} + \begin{bmatrix} \vec{t} \\ \vec{t} \end{bmatrix}_{\vec{F}_1} \right]_{\vec{F}_1}, \quad (17) \quad \langle \vec{s}, \vec{t} \rangle_{\mathcal{S}_{\hat{\omega}}} := \mathbf{s} \cdot \begin{bmatrix} \vec{t} \\ \vec{t} \end{bmatrix}_{\vec{F}_1}, \quad (18)$$

respectively. We also define the rotation of the underlying polarized ray of \vec{s} itself. For a rotation $\vec{R} \in \vec{SO}(3)$, to avoid confusion, we denote \vec{R}_S as Stokes vector version of \vec{R} . Then for any Stokes vector $\vec{s} \in \mathcal{S}_{\hat{\omega}}$, \vec{R}_S acts as

$$\vec{R}_S \vec{s} = \left[\begin{bmatrix} \vec{s} \\ \vec{s} \end{bmatrix}_{\vec{R}\vec{F}} \right]_{\vec{R}\vec{F}} \in \mathcal{S}_{\vec{R}\hat{\omega}}, \quad (19)$$

where $\vec{F}[:, 3] = \hat{\omega}$. Figure 4 visualizes the difference between a coordinate conversion and a rotation around $\hat{\omega}$. Note that these are frame-independently well-defined.

Spin-2 vs. full Stokes vectors. Suppose that we have a Stokes vector $\vec{s} = [\mathbf{s}]_{\vec{F}}$ with $\mathbf{s} = [s_0, s_1, s_2, s_3]^T$. To handle the special behaviors of linear polarization components s_1 and s_2 , we sometimes need to process only these two separately from the four components. To do so, we define a *spin-2* Stokes vector (in *spin-2* Stokes space) as $\left[[s_1, s_2]^T \right]_{\vec{F}} \in \mathcal{S}_{\hat{\omega}}^2$ in a similar way to Equation (15). Then the original Stokes vector, also called a *full Stokes vector* to be clear and written as $\vec{s} = s_0 \oplus \left[[s_1, s_2]^T \right]_{\vec{F}} \oplus s_3$, where \oplus symbol indicates the direct sum in linear algebra, which also can be considered as vector concatenation in numerical programming tools.

Now a spin-2 Stokes vector can also be written with a complex component as $\left[[s_1, s_2]^T \right]_{\vec{F}} = [s_1 + is_2]_{\vec{F}}$. With this representation, multiplication by a complex number:

$$z [s_1 + is_2]_{\vec{F}} = [z(s_1 + is_2)]_{\vec{F}}, \quad (20)$$

is well defined independent of choice of the frame \vec{F} , which indicates scaling by $|z|$ followed by rotating $\arg z/2$ around its ray direction, as illustrated in Figure 3(b). Note that while other operations such as addition (Equation (17)) and inner product (Equation (18)) are defined both for spin-2 and full Stokes vectors, the complex multiple is only defined for spin-2 Stokes vectors.

4.2.2 Mueller transform in numeric vs. geometric quantities. Linear maps from Stokes vectors along $\hat{\omega}_i$ to Stokes vectors along $\hat{\omega}_o$, such as polarimetric BRDF and other polarized light interactions, are called *Mueller transforms*. The set of these Mueller transforms is called a *Mueller space* and written as

$$\mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o} := \left\{ \vec{M} : \mathcal{S}_{\hat{\omega}_i} \rightarrow \mathcal{S}_{\hat{\omega}_o} \mid \vec{M} (a\vec{s} + b\vec{t}) = a\vec{M}\vec{s} + b\vec{M}\vec{t} \right. \\ \left. \text{for any } a, b \in \mathbb{R} \text{ and } \vec{s}, \vec{t} \in \mathcal{S}_{\hat{\omega}_i} \right\}. \quad (21)$$

Similar to Stokes vector, a Mueller transform $\vec{M} \in \mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o}$ is a geometric quantity, and it can be measured into a numeric matrix $\mathbf{M} \in \mathbb{R}^{4 \times 4}$, named *Mueller matrix*³ with respect to observing local frames. Here, we need two frames \vec{F}_i and \vec{F}_o with $\vec{F}_i[:, 3] = \hat{\omega}_i$ and $\vec{F}_o[:, 3] = \hat{\omega}_o$ and relations between \vec{M} and \mathbf{M} is notated as follows:

$$\vec{M} = [\mathbf{M}]_{\vec{F}_1 \rightarrow \vec{F}_2}, \quad \mathbf{M} = [\vec{M}]_{\vec{F}_1 \rightarrow \vec{F}_2}, \quad (22)$$

similar to Equations (15) and (16).

5 CHALLENGES OF STOKES VECTOR FIELDS IN ANGULAR DOMAIN

Stokes vector radiance as a function on an angular domain, called a *Stokes vector field*, is a fundamental quantity to describe polarized transport. It has been the subject of previous work such as polarized environment illumination, including the sky dome [Riviere et al. 2017; Wilkie et al. 2004, 2021] and polarized perspective images in all the existing polarization renderers. However, the challenges of dealing with Stokes vector fields have rarely been discussed. In this section, we introduce such challenges in terms of different continuity conditions from scalar fields in Section 5.1. It raises the necessity of novel basis functions rather than scalar SH for frequency domain methods of polarized light. In Section 5.2, we additionally define basic operations on Stokes vector fields, which are required for frequency domain analysis.

5.1 Continuity of Scalar vs. Stokes Vector Fields

A Stokes vector field on the unit sphere⁴ is formulated as $\vec{f}: \mathbb{S}^2 \rightarrow \mathcal{S}_{\hat{\omega}}$. Here we can observe that, unlike scalar radiance, the value of the Stokes vector field at each direction $\hat{\omega}$ lies on the different Stokes space, i.e., $\vec{f}(\hat{\omega}) \in \mathcal{S}_{\hat{\omega}}$, depending on the direction $\hat{\omega}$.

The simple way to measure a Stokes vector field is to assign local frames for each direction $\hat{\omega}$. We call this type of function the function from directions $\hat{\omega} \in \mathbb{S}^2$ to local frames $\vec{F}(\hat{\omega}) \in \mathbb{F}^3$ with $\vec{F}(\hat{\omega})[:, 3] = \hat{\omega}$ as a *frame field*.

Among the various choices of frame field, one choice is a $\theta\phi$ frame field $\vec{F}_{\theta\phi}$, defined by aligning local x and y axes along longitudinal and latitudinal directions as shown in Figure 5(b) and Supplemental Equation (91).

Note that \vec{f} can be visualized as double-sided arrows (following Figure 2) on tangent planes of the sphere, as shown in Figure 5(a). After choosing the frame field, \vec{f} can be converted into four scalar fields on the sphere based on the numeric-geometric conversion notation we defined in Equation (15) as

$$\left[\vec{f}(\hat{\omega}) \right]_{\vec{F}_{\theta\phi}(\hat{\omega})} = [f_0(\theta, \phi) \quad f_1(\theta, \phi) \quad f_2(\theta, \phi) \quad f_3(\theta, \phi)]^T. \quad (23)$$

³Similar to Stokes vector and Stokes component vectors, we distinguish terminologies *Mueller transforms* and *Mueller matrices*.

⁴Rigorously, it should be written as $\left\{ \vec{f}: \mathbb{S}^2 \rightarrow \cup_{\hat{\omega} \in \mathbb{S}^2} \mathcal{S}_{\hat{\omega}} \mid \forall \hat{\omega} \in \mathbb{S}^2, \vec{f}(\hat{\omega}) \in \mathcal{S}_{\hat{\omega}} \right\}$, but we write as the main text for the sake of simplicity and better intuition.

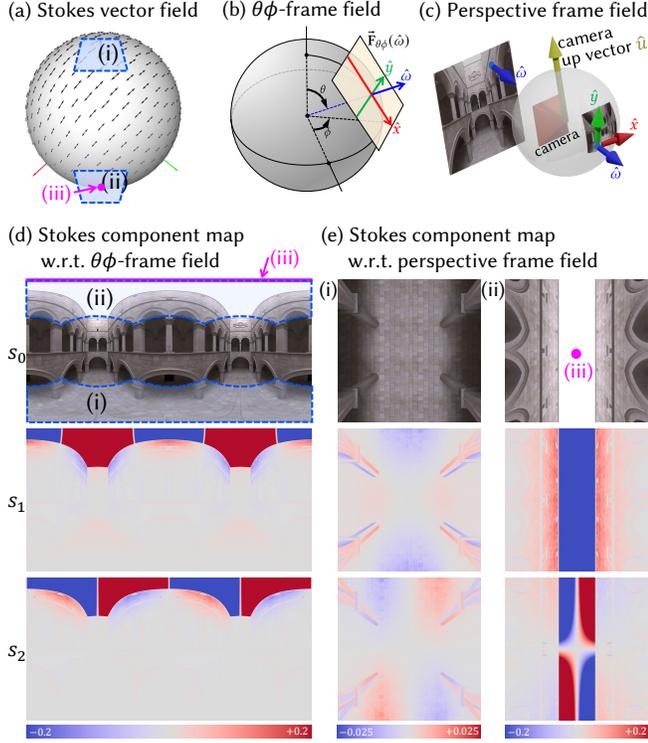


Fig. 5. Visualizing a Stokes vector field (polarized environment map) depends on the choice of frame fields. Taking Stokes components of Stokes vector field (a) with respect to a typical $\theta\phi$ -frame field (b) yields equirectangular images shown in (d). Using a perspective frame field used in Mitsuba 3 renderer, several perspective images are visualized as (e). Note that while the s_1 component (ii) in (e) at the sky, especially (iii), has consistent signs of values, and the component in (d) under a different frame field has a different trend of values.

We can visualize scalar fields of each component f_i as equirectangular images by unwrapping the spherical domain into the rectangle of spherical coordinates θ and ϕ as Figure 5(d).

However, there is an issue that *any* frame field always has a *singularity*, which means a local frame cannot be continuously defined due to the Hairy Ball Theorem [Nash and Sen 1983]. For example, $\vec{F}_{\theta\phi}(\hat{\omega})$ has two singularities⁵ at \hat{z}_g ($\theta = 0$) and $-\hat{z}_g$ ($\theta = \pi$). In the rectangle domain, the top (and bottom, respectively) edge indicates just a single point \hat{z}_g ($-\hat{z}_g$, respectively) but has different local frames that rotate one turn in counterclockwise (clockwise, respectively) as ϕ increases. It yields different continuity conditions for scalar and Stokes vector fields. While scalar fields (e.g., scalar radiance, f_0 or f_3) have constant values at those top and bottom edges, a two-dimensional numeric vector $[f_1(0, \phi), f_2(0, \phi)]$ rotates twice in clockwise as ϕ increases from 0 to 2π due to rotation of the frame field $\vec{F}_{\theta\phi}(0, \phi)$, and similarly for $\theta = \pi$. These difference are illustrated in Figure 6.

⁵We let axis symbols without subscripts such as \hat{x} and \hat{y} denote values of a frame field, which are used to measure Stokes vectors along each direction, while those with subscript g such as \hat{z}_g denote a fixed global frame which is used to assign spherical coordinates on a sphere.

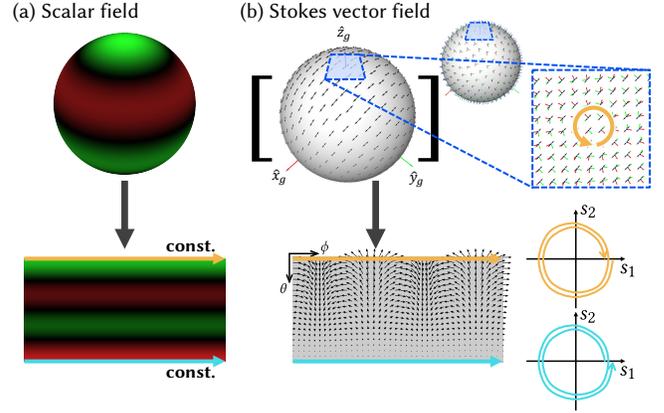


Fig. 6. Analyzing continuity and smoothness for Stokes vector fields in $\theta\phi$ domain. (a) The visualization of a Stokes vector field. As a *geometric* quantity, Stokes vector fields are continuous and smooth on the entire sphere, including the zenith. (b) To make the *geometric* Stokes vector fields to *numeric* Stokes components, we can assign the specific frame field, named $\theta\phi$ -frame field.

Note that such pair of spherical functions with the continuity condition of twice rotation such as s_1 and s_2 Stokes components are called *spin-2* functions, and scalar functions with the constant continuity condition are called *spin-0* functions.

To construct a frequency domain method similar to ones based on scalar SH in scalar rendering, one may consider a naive approach to apply scalar SH combined with the $\theta\phi$ -frame field as a basis of Stokes vector fields. However, this approach raises the singularity problem due to the different continuity conditions between scalar and Stokes vector fields. In Figure 7, the s_1 and s_2 components of the original Stokes vector field are nearly flat around $\pm\hat{z}_g$ (views (i) and (ii)), but its projection onto the basis obtained by the naive approach yields (i) too high-frequency change at (b) or (ii) singularity at (b). This is a fundamentally different feature from how the conventional SH behaved on scalar fields, which always converts finite coefficients to continuous functions and has a smoothing role. We also point out that this singularity problem also implies a violation of rotation invariance. We refer to Figure 11, which is described in Section 6.2.2, and Supplemental Section 4 for more discussion.

In summary, the different continuity conditions are an essential difference in the nature of Stokes vector fields. Although we only show the case of the $\theta\phi$ frame field here, Stokes vector fields *always* have different properties in terms of continuity regardless of which frame field is used.

5.2 Stokes Vector Fields Operations

To discuss bases for Stokes vector fields, we should define several operations on Stokes vector fields. It can be done by generalizing scalar field operations in Section 4.1, based on Stokes vectors operations in Section 4.2. The inner product of two Stokes vector fields \vec{f} and $\vec{g}: \hat{\mathbb{S}}^2 \rightarrow \mathcal{S}_{\hat{\omega}}$ is defined as follows:

$$\langle \vec{f}, \vec{g} \rangle_{\mathcal{F}} := \int_{\hat{\mathbb{S}}^2} \langle \vec{f}(\hat{\omega}), \vec{g}(\hat{\omega}) \rangle_{\mathcal{S}} d\hat{\omega}. \quad (24)$$

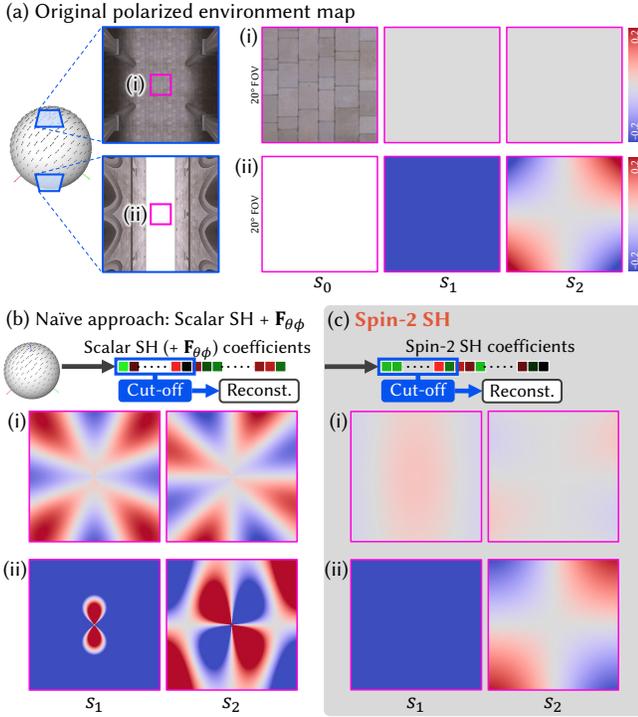


Fig. 7. We propose a frequency-domain analysis framework of Stokes vector fields, which is represented by a polarized environment map here. Then, we need spin-2 spherical harmonics rather than conventional ones to avoid the singularity problem. See Figure 11 for rotation invariance of spin-2 SH, and see Figure 10 and Equation (36) for how the coefficient matrix of rotation under conventional SH (Wigner D-functions) can be utilized to spin-2 SH.

In addition, the rotation acting on Stokes vector fields by $\vec{R} \in \vec{SO}(3)$ is defined by

$$\vec{R}_{\mathcal{F}}[\vec{f}](\hat{\omega}) = \vec{R}_{\mathcal{S}}(\vec{f}(\vec{R}^{-1}\hat{\omega})), \quad (25)$$

for any \vec{f} . We summarize inner products and rotations on different types of quantities in Table 2.

6 POLARIZED SPHERICAL HARMONICS

To overcome the challenges described in Section 5 and bring benefits of frequency-domain framework to polarized radiance functions, we need a novel set of basis functions, *polarized spherical harmonics*. In Section 6.1, we introduce spin-weighted SH and show how it plays a role in the basis functions for polarized light transport in computer graphics. Although spin-weighted SH are an existing theory in physics [Goldberg et al. 1967; Newman and Penrose 1966], it has never been used in rendering pipelines to describe full Stokes vectors and Mueller transforms. In Section 6.2, we introduce our polarized spherical harmonics, combining spin-0 (scalar) SH and spin-2 SH, which can fully describe Stokes vector fields for polarized frequency domain analysis.

Moreover, we will also show how to perform rotation (Section 6.2), linear operators (e.g., general pBRDFs and radiance transfer, Section 6.3), and convolution (Section 6.4) in the PSH domain, which are inevitable operations in frequency-domain analysis. These three main operations are not only the theoretical foundation but also the

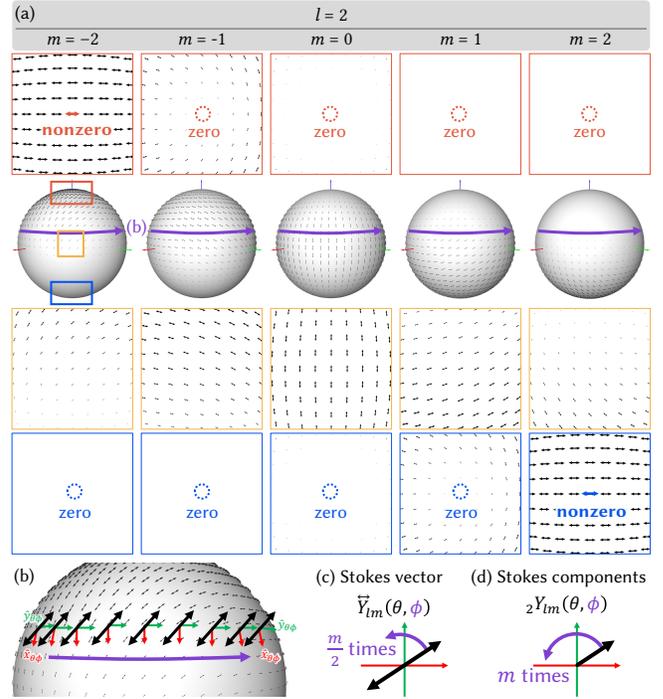


Fig. 8. Visualization of the first order ($l = 2$) of spin-2 spherical harmonics, defined in Eqs. (26a) and (28), which are the basis functions of the space of linear polarization as functions of directions (spin-2 Stokes vector fields). (a) The first, third, and fourth rows show the closeup of the region indicated in the second row. Note that spin-2 SH only have nonzero values at the north pole for $m = -2$ and the south pole for $m = 2$, ($\hat{\omega} = \pm \hat{z}_g$, i.e., $\theta = 0$ and π) respectively. In addition, tracing a line with fixed θ by increasing ϕ , as the blue curves in (a) and (b), can be seen as (c) a Stokes vector rotating $\frac{m}{2}$ times and (d) Stokes components ${}_2Y_{lm}$ rotating m times.

main building blocks of our PSH rendering pipeline. See Section 7 for our real-time polarized rendering results based on our theory described in overall Section 6.

6.1 Spin-Weighted Spherical Harmonics

The spin-weighted spherical harmonics ${}_sY_{lm}$ are the basis for spin- s functions on the sphere, and they have continuity conditions depicted in Figure 6 by replacing the double rotation by s times rotation [Goldberg et al. 1967; Newman and Penrose 1966]. As a brief introduction, SWSH can be derived from the basis for functions on higher dimensional space, rotation transforms $\vec{SO}(3)$, and introducing appropriate constraints that make these higher dimensional functions equivalent to spin- s functions on the sphere.

For more motivation and derivation of SWSH, refer to Supplemental Section 5.1, and here we focus on the usage of SWSH.

To handle Stokes vectors, we focus on spin $s = 0$ and $s = 2$. With $s = 0$, SWSH are exactly the same as conventional SH (${}_0Y_{lm} = Y_{lm}$), so SWSH can be considered as a generalization of SH. When $s = 2$, spin-2 SH (${}_2Y_{lm}$) become an orthonormal basis for spin-2 functions such as Stokes vector fields. While there are several types of formulae to evaluate spin-2 SH, we introduce a way by utilizing scalar (spin-0) SH as follows:

Spin-2 SH (w.r.t. $\theta\phi$ -frame field)

$${}_2Y_{lm}(\theta, \phi) = \sqrt{\frac{(l-2)!}{(l+2)!}} \left[\alpha_{lm}(\theta) Y_{lm}(\theta, \phi) + \beta_{lm}(\theta) Y_{l-1,m}(\theta, \phi) \right], \quad (26a)$$

$$\alpha_{lm}(\theta) = \frac{2m^2 - l(l+1)}{\sin^2 \theta} - 2m(l-1) \frac{\cot \theta}{\sin \theta} + l(l-1) \cot^2 \theta, \quad (26b)$$

$$\beta_{lm}(\theta) = 2\sqrt{\frac{2l+1}{2l-1}} (l^2 - m^2) \left(\frac{m}{\sin^2 \theta} + \frac{\cot \theta}{\sin \theta} \right). \quad (26c)$$

Note that ${}_2Y_{lm}$ here is complex-valued Stokes *components* of a basis for spin-2 Stokes *vector* fields with respect to the $\theta\phi$ frame field $\vec{F}_{\theta\phi}$. Thus, ${}_2Y_{lm}$ satisfies the following condition, which indicates the double rotation at the north and south poles as visualized in Figure 6(b) as follows:

$$\begin{aligned} {}_2Y_{lm}(0, \phi) &= 0, & \text{if } m \neq -2 \\ {}_2Y_{l,-2}(0, \phi) &= e^{-2i\phi} \cdot \text{const.} \neq 0, \\ {}_2Y_{lm}(\pi, \phi) &= 0, & \text{if } m \neq 2 \\ {}_2Y_{l,2}(\pi, \phi) &= e^{2i\phi} \cdot \text{const.} \neq 0. \end{aligned} \quad (27)$$

Now, using the numeric-geometric conversion (Equation (15)), we can define a (geometric) Stokes vector version of spin-2 SH as

$$\vec{Y}_{lm}(\hat{\omega}) := [{}_2Y_{lm}(\theta, \phi)]_{\vec{F}_{\theta\phi}(\theta, \phi)}. \quad (28)$$

The first order $l = 2$ of spin-2 SH \vec{Y}_{lm} are visualized in Figure 8. Note that due to the nature of spin-2 functions, there are no orders $l = 0$ and $l = 1$.

Both terms $Y_{lm}(\theta, \phi)$ and $Y_{l-1,m}(\theta, \phi)$ in Equation (26a) have $e^{im\phi}$ terms originated from Equation (1). From the $e^{im\phi}$ term, we can observe that, following the circle formulated by some fixed θ on the sphere, the geometric Stokes vector (double-sided arrow) rotates $\frac{m}{2}$ times. In contrast, numeric Stokes components rotate m times, as shown in Figure 8(b).

Additionally, when comparing spin-2 SH with scalar SH, spin-2 SH are similar in that they have azimuthal symmetry at $m = 0$. However, one difference is the condition of non-zero value at $\hat{\omega} = \pm \hat{z}_g$; for spin-2 SH, it occurs at $m = -2$ or $m = 2$ (Figure 8), while for scalar SH, it occurs at $m = 0$.

6.2 Polarized Spherical Harmonics

Now we combine Stokes components spin-0 functions s_0 (total intensity) and s_3 (circular polarization) with spin-2 functions s_1 and s_2 (linear polarization). Then we define the orthonormal basis *polarized spherical harmonics*, which span the full Stokes vectors fields $\mathcal{F}(\hat{\mathbb{S}}^2, \mathcal{S}_{\hat{\omega}})$ over real coefficients. By using the additional index $p = 0, 1, 2, 3$ that indicates the index of polarization components s_0, s_1, s_2, s_3 respectively, the PSH \vec{Y}_{lmp} are defined by

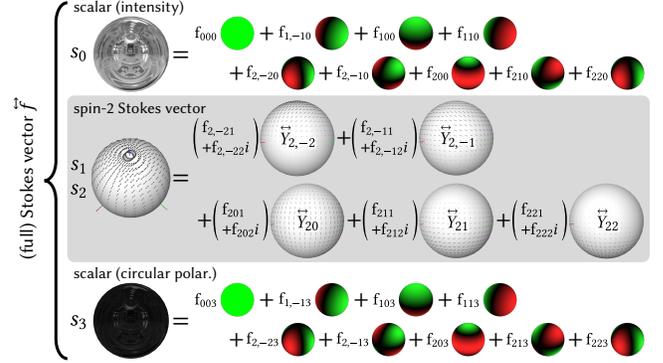


Fig. 9. Any full Stokes vector field \vec{f} can be linearly decomposed using scalar (spin-0) SH Y_{lm}^R for s_0 (intensity) and s_3 (circular polarization) and spin-2 SH \vec{Y}_{lm} (or \vec{Y}_{lm1} and \vec{Y}_{lm2} , equivalently) for s_1 and s_2 (linear polarization) as described in Equations (29) and (30). The coefficient vector, which consists of such f_{lmp} , becomes our frequency domain representation of given \vec{f} . Note that while s_0 and s_3 components in the original angular domain completely correspond to coefficients with $p = 0$ and $p = 3$, respectively, s_1 and s_2 components do not exactly correspond to $p = 1$ and $p = 2$, respectively, due to their values depend on the choice of frame fields.

Polarized spherical harmonics

$$\begin{aligned} \vec{Y}_{lm0}(\hat{\omega}) &= \begin{bmatrix} Y_{lm}^R(\hat{\omega}) \\ 0 \\ 0 \\ 0 \end{bmatrix}_{\vec{F}_{\theta\phi}(\hat{\omega})}, \quad \vec{Y}_{lm1}(\hat{\omega}) = \begin{bmatrix} 0 \\ \Re [{}_2Y_{lm}(\hat{\omega})] \\ \Im [{}_2Y_{lm}(\hat{\omega})] \\ 0 \end{bmatrix}_{\vec{F}_{\theta\phi}(\hat{\omega})}, \\ \vec{Y}_{lm2}(\hat{\omega}) &= \begin{bmatrix} 0 \\ -\Im [{}_2Y_{lm}(\hat{\omega})] \\ \Re [{}_2Y_{lm}(\hat{\omega})] \\ 0 \end{bmatrix}_{\vec{F}_{\theta\phi}(\hat{\omega})}, \quad \vec{Y}_{lm3}(\hat{\omega}) = \begin{bmatrix} 0 \\ 0 \\ 0 \\ Y_{lm}^R(\hat{\omega}) \end{bmatrix}_{\vec{F}_{\theta\phi}(\hat{\omega})}. \end{aligned} \quad (29)$$

Here, $\Re[\cdot]$ and $\Im[\cdot]$ indicate real and imaginary part of some scalar complex number z where $z = \Re[z] + i\Im[z]$. Using these bases, any Stokes vector field \vec{f} can be written as

$$\vec{f}(\hat{\omega}) = \sum_{(l,m,p) \in I_{PSH}} f_{lmp} \vec{Y}_{lmp}, \quad (30)$$

where I_{PSH} denotes the set of the indices l, m , and p :

$$\begin{aligned} I_{PSH} &= \{(l, m, p) \in \mathbb{Z}^3 \mid l \geq 0, |m| \leq l, p = \{0, 3\}\} \\ &\cup \{(l, m, p) \in \mathbb{Z}^3 \mid l \geq 2, |m| \leq l, p = \{1, 2\}\}, \end{aligned} \quad (31)$$

and the coefficient f_{lmp} can be computed as

$$f_{lmp} = \left\langle \vec{Y}_{lmp}, \vec{f} \right\rangle_{\mathcal{F}}. \quad (32)$$

By using PSH, the decomposition example is illustrated in Figure 9, when f is a polarized environment map.

Real coefficient formulation. One important adaption from spin-2 SH to our PSH is the separation of the complex part to make the coefficient a real number. Suppose we have a spin-2 Stokes vector

field \vec{f} , which only considers the linear polarization part. Generally, using spin-2 SH, we can write frequency domain representation with complex number coefficient as

$$\vec{f} = \sum_{l,m} \underbrace{(f_{lm1} + f_{lm2}i)}_{\text{complex coeff.}} \underbrace{\vec{Y}_{lm}}_{\text{basis}}. \quad (33)$$

In contrast, using PSH, we can write the real number coefficient using the form as

$$\vec{f} = \sum_{l,m} \underbrace{f_{lm1}}_{\text{real coeff.}} \underbrace{\vec{Y}_{lm1}}_{\text{basis}} + \underbrace{f_{lm2}}_{\text{real coeff.}} \underbrace{\vec{Y}_{lm2}}_{\text{basis}}. \quad (34)$$

Although it looks trivially identical, there are some reasons why this real-valued adaptation is important. First, since the real-world quantities (Stokes vectors and Mueller transforms) have a real value, using real-valued representation allows us to easily manage the consistency when computing such quantities in the frequency domain. Second, the formulation in Equation (33) actually loses the information for representing Mueller transforms, while Equation (34) does not. This will be introduced in later Section 6.3. For spin-0 components, we use real SH Y_{lm}^R rather than complex SH Y_{lm} not only for the consistency to the angular domain but also to take algebraic closedness of induced coefficients matrices into account, which is discussed in Supplemental Section 5.3.1. Hence, we choose the real-valued formulation to build a solid theory for our PSH, except for some intermediate representations for efficient derivations that do not violate the reasons for choosing the real-valued formulation.

6.2.1 Rotation invariance of PSH. Since PSH are an orthonormal basis, the PSH coefficient rotation can also be done with the coefficient matrix similar to scalar SH (Equation (9)). For given Stokes vector field \vec{f} and rotation $\vec{R}_{\mathcal{F}}$, the rotated coefficient f'_{lmp} can be computed as

$$f'_{l_o m_o p_o} = \sum_{l_i, m_i, p_i} \underbrace{\left\langle \vec{Y}_{l_o m_o p_o}, \vec{R}_{\mathcal{F}} \left[\vec{Y}_{l_i m_i p_i} \right] \right\rangle_{\mathcal{F}}}_{\text{coefficient matrix}} f_{l_i m_i p_i}, \quad (35)$$

where subscript o at the indices notes output (rotated) and subscript i at the indices notes input. By using the definitions, the coefficient matrix at p_o -th row and p_i -th column can be calculated as

Coefficient matrix of the rotation in PSH

$$\left[\left\langle \vec{Y}_{l_o m_o p_o}, \vec{R}_{\mathcal{F}} \left[\vec{Y}_{l_i m_i p_i} \right] \right\rangle_{p_o, p_i} \right] =$$

$$\delta_{l_i l_o} \begin{bmatrix} \underbrace{D_{m_o m_i}^{L,R}(\vec{R})}_{p_o=0, p_i=0} & & 0 \\ 0 & \underbrace{\mathbb{R}^{2 \times 2} \left(D_{m_o m_i}^{L,C}(\vec{R}) \right)}_{p_o=\{1,2\}, p_i=\{1,2\}} & 0 \\ 0 & & \underbrace{D_{m_o m_i}^{L,R}(\vec{R})}_{p_o=3, p_i=3} \end{bmatrix}. \quad (36)$$

Note that $D_{m_o m_i}^{L,C}$ and $D_{m_o m_i}^{L,R}$ are complex and real Wigner-D functions defined in Equations (9) and (13) respectively, and $\mathbb{R}^{2 \times 2}$ indicates an operator that convert complex numbers to 2×2 real matrices as

$$\mathbb{R}^{2 \times 2} (x + yi) := \begin{bmatrix} x & -y \\ y & x \end{bmatrix}. \quad (37)$$

As a result, we can observe that the resulting coefficient matrix of rotation on Stokes vector fields in Equation (36) only computes the same order l for f and f' . This means the resulting matrix is block diagonal, and PSH satisfy the rotation invariance property. Moreover, since Equation (36) consists of existing Wigner D-functions, another advantage is that we can utilize existing formulas and computation methods from scalar SH rotation. For more details and derivations of proving rotation invariance, refer to Supplemental Section 5.4.

6.2.2 Rotation invariance validation.

Numerical validation. So far, we have shown the theoretical guarantee of rotation invariance of PSH; here, we will show it numerically. For the given rotation transform, we can compute the corresponding coefficient matrix that rotates some physical quantity with respect to some basis function. Then, we can validate rotation invariance by checking the block diagonal behavior of the computed coefficient matrix. Figure 10 shows such false-color magnitude visualization of the complex-numbered coefficient matrix, with rotation transform $\vec{R} = \vec{R}_{\hat{u}}(\theta)$ with $\theta \hat{u} = \vec{F}_g [10, 0.1, 0.2]^T$. Figure 10(a) shows the coefficient matrix of rotating scalar radiance projected on the scalar SH. It can be clearly shown that the coefficient matrix is block diagonal. Figures 10(b) and 10(c) shows the case of rotating spin-2 part of Stokes vector projected on the scalar SH and spin-2 SH, respectively. Note that here we use complex-valued representation (Equation (33)) to compare with the scalar radiance case in Figure 10(a). Also, we use the $\theta\phi$ -frame field for scalar SH projection of Stokes vectors since we need to specify the frame field as described in Equation (23). As shown in Figure 10(b), using scalar SH on Stokes vectors never becomes block diagonal so that implies no rotation invariance. In contrast, as shown in Figure 10(c), the coefficient matrix of spin-2 SH on Stokes vector is block diagonal, which implies the rotation invariance and even computed value are the same as scalar radiance case in Figure 10(a).

Polarized environment map reconstruction. We also validate the rotation invariance with the polarized environment map, as shown in Figure 11. Similar to the numerical validation above, we only show the spin-2 part of the Stokes vector. For the given polarized environment map, we initially project it to the basis function such as scalar SH (with $\theta\phi$ -frame field) or spin-2 SH, and *cut-off* the coefficient vector to take the finite coefficient vector. First, we reconstruct that finite coefficient vector into an angular domain, which yields a band-limited environment map. On the other hand, we rotate that finite coefficient vector with the given rotation transform \vec{R} and perform reconstruction with the *rotated* basis with the same rotation transform \vec{R} . Since we rotate both coefficients and basis with the same rotation transform, the reconstruction result should be the same as the vanilla cut-off reconstruction case. As a result, the naïve approach using scalar SH with $\theta\phi$ -frame field shows inconsistent

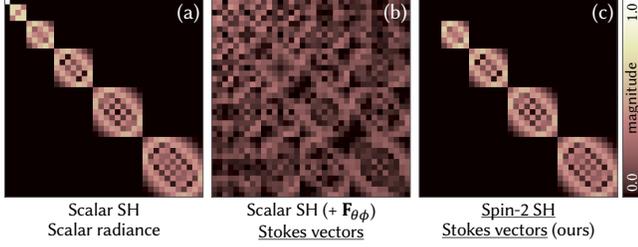


Fig. 10. Comparison of coefficient matrices of a particular rotation with respect to each basis. Each matrix indicates the coefficient matrix of the rotation with respect to scalar SH (Equation (8), also visualized in Supplemental Figure 4), scalar SH with $\theta\phi$ -frame field, and spin-2 SH (Equation (36)) for (a), (b), and (c), respectively. Each row and column indicates enumerated pairs of (l, m) indices. (a) In the case of scalar SH and scalar radiance, the coefficient matrix (Wigner D-function) shows block-diagonal behavior, and it yields the rotation invariance. (b) However, if we naively apply rotation using scalar SH with $\theta\phi$ -frame field to Stokes vectors, the rotation invariance does not hold anymore. (c) By changing the basis to spin-2 SH, the rotation invariance holds on Stokes vectors. Note that spin-2 SH starts from $l = 2$, so the first two block diagonals are empty in (c).

behavior (Figure 11(a)), while using spin-2 SH results between two reconstructions is identical (Figure 11(b)).

6.3 Coefficient Matrices for pBRDF and Radiance Transfer

Beyond coefficient vector representation of polarized environment map, our PSH also provide frequency domain representation for polarized light interaction such as pBRDF or radiance transfer operator into *coefficient matrices*. Here we derive a general formulation of PSH coefficient matrices that extends scalar quantities described in Section 4.1, Equations (4) to (7). Recall that the coefficient matrix generally represents linear operators on Stokes vector fields in the angular domain. Hence, they can be characterized as a *Mueller transform field* \vec{P} , which is a function from given two directions $\hat{\omega}_i$ and $\hat{\omega}_o$ to a Mueller transform as

$$\vec{P}: \hat{\mathbb{S}}^2 \times \hat{\mathbb{S}}^2 \rightarrow \mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o}. \quad (38)$$

Note that a Mueller transform field can be considered as a (cosine-weighted) pBRDF, that can act on a Stokes vector field \vec{f} as a linear operator as

$$\vec{P}_{\mathcal{F}}[\vec{f}](\hat{\omega}_o) = \int_{\hat{\mathbb{S}}^2} \vec{P}(\hat{\omega}_i, \hat{\omega}_o) \vec{f}(\hat{\omega}_i) d\hat{\omega}_i. \quad (39)$$

As a result, by using the appropriate type of inner product described in Equations (18) and (24), the PSH coefficients of the Mueller transform can be directly extended from scalar SH coefficients (Equation (4)) as

$$\begin{aligned} P_{l_o m_o p_o, l_i m_i p_i} &= \left\langle \vec{Y}_{l_o m_o p_o}, \vec{P}_{\mathcal{F}} \left[\vec{Y}_{l_i m_i p_i} \right] \right\rangle_{\mathcal{F}} \\ &= \int_{\hat{\mathbb{S}}^2 \times \hat{\mathbb{S}}^2} \left\langle \vec{Y}_{l_o m_o p_o}(\hat{\omega}_o), \vec{P}(\hat{\omega}_i, \hat{\omega}_o) \vec{Y}_{l_i m_i p_i}(\hat{\omega}_i) \right\rangle_{\mathcal{S}} d\hat{\omega}_i d\hat{\omega}_o. \end{aligned} \quad (40)$$

Similar to scalar SH, regarding the indices (l_o, m_o, p_o) as rows and (l_i, m_i, p_i) as columns, we can obtain the *coefficient matrix* of \vec{P} . Now

suppose that we have a coefficient vector $f_{l_i m_i p_i}$ from a polarized incident radiance \vec{f} , obtained by Equation (30) and a coefficient matrix $P_{l_o m_o p_o, l_i m_i p_i}$ from a pBRDF \vec{P} . Then, similar to the conventional scalar SH-based rendering pipeline (Equation (5)), the coefficient vector of reflected radiance $\left\langle \vec{Y}_{l_o m_o p_o}, \vec{P}_{\mathcal{F}}[\vec{f}] \right\rangle$ is evaluated by a matrix-vector product as

$$\left\langle \vec{Y}_{l_o m_o p_o}, \vec{P}_{\mathcal{F}}[\vec{f}] \right\rangle = \sum_{(l_i, m_i, p_i) \in I_{PSH}} P_{l_o m_o p_o, l_i m_i p_i} f_{l_i m_i p_i}. \quad (41)$$

6.3.1 Submatrices of Mueller transforms and coefficient matrices.

Due to the nature of Mueller transform, there are additional indices p_o and p_i in Equation (40). Consequently, we have $16 (= 4 \times 4)$ times more coefficients than the coefficient matrices in scalar SH. For further analysis and efficient computation in a constant factor, we can split a Mueller transform and the corresponding coefficient matrix. From the given Mueller transform \vec{P} in the angular domain and a single pair of directions $(\hat{\omega}_i, \hat{\omega}_o)$, we can denote a single Mueller transform $\vec{M} = \vec{P}(\hat{\omega}_i, \hat{\omega}_o)$. By using the numeric-geometric conversion, the numeric Mueller matrix \mathbf{M} can be computed as

$\mathbf{M} = \left[\vec{M} \right]_{\vec{F}_i \rightarrow \vec{F}_o}$. Now recall that s_1, s_2 components are dependent to frame (spin-2), and s_0, s_3 are independent to frame (spin-0). In this context, the Mueller matrix can be split into 9 submatrices according to dependency on the \vec{F}_i and \vec{F}_o as

$$\mathbf{M} = \begin{bmatrix} \begin{matrix} M_{00} \\ M_{10} \\ M_{20} \\ M_{30} \end{matrix} & \begin{matrix} M_{01} & M_{02} \\ M_{11} & M_{12} \\ M_{21} & M_{22} \\ M_{31} & M_{32} \end{matrix} & \begin{matrix} M_{03} \\ M_{13} \\ M_{23} \\ M_{33} \end{matrix} \end{bmatrix}. \quad (42)$$

By following its spin-weights, we call the submatrices *spin 0-to-0*, *spin 2-to-0*, *spin 0-to-2*, and *spin 2-to-2* blocks. This decomposition is also valid to evaluate the *coefficient matrix* of \vec{P} . By fixing indices l_o, m_o, l_i , and m_i for the coefficients defined in Equation (40), we can split the coefficient matrix in the same way as

$$\begin{bmatrix} P_{l_o m_o 0, l_i m_i 0} & P_{l_o m_o 0, l_i m_i 1} & P_{l_o m_o 0, l_i m_i 2} & P_{l_o m_o 0, l_i m_i 3} \\ P_{l_o m_o 1, l_i m_i 0} & P_{l_o m_o 1, l_i m_i 1} & P_{l_o m_o 1, l_i m_i 2} & P_{l_o m_o 1, l_i m_i 3} \\ P_{l_o m_o 2, l_i m_i 0} & P_{l_o m_o 2, l_i m_i 1} & P_{l_o m_o 2, l_i m_i 2} & P_{l_o m_o 2, l_i m_i 3} \\ P_{l_o m_o 3, l_i m_i 0} & P_{l_o m_o 3, l_i m_i 1} & P_{l_o m_o 3, l_i m_i 2} & P_{l_o m_o 3, l_i m_i 3} \end{bmatrix}. \quad (43)$$

What we can observe here is each of the nine submatrices in the Mueller matrix in the angular domain (Equation (42)) only affects the corresponding submatrix in the coefficient matrix in the frequency domain (Equation (43)). This fact allows us to compute the coefficient matrix of each block separately, with less memory requirement for simulating numerical integration for Equation (40). In other words, the matrix product with sizes 1×4 , 4×4 , and 4×1 in the integrand of Equation (40) can be reduced to 1×2 , 2×2 , and 2×1 , respectively.

6.3.2 Complex pair separation of spin 2-to-2 Mueller transform.

In addition to separating the full Mueller transform into nine blocks,

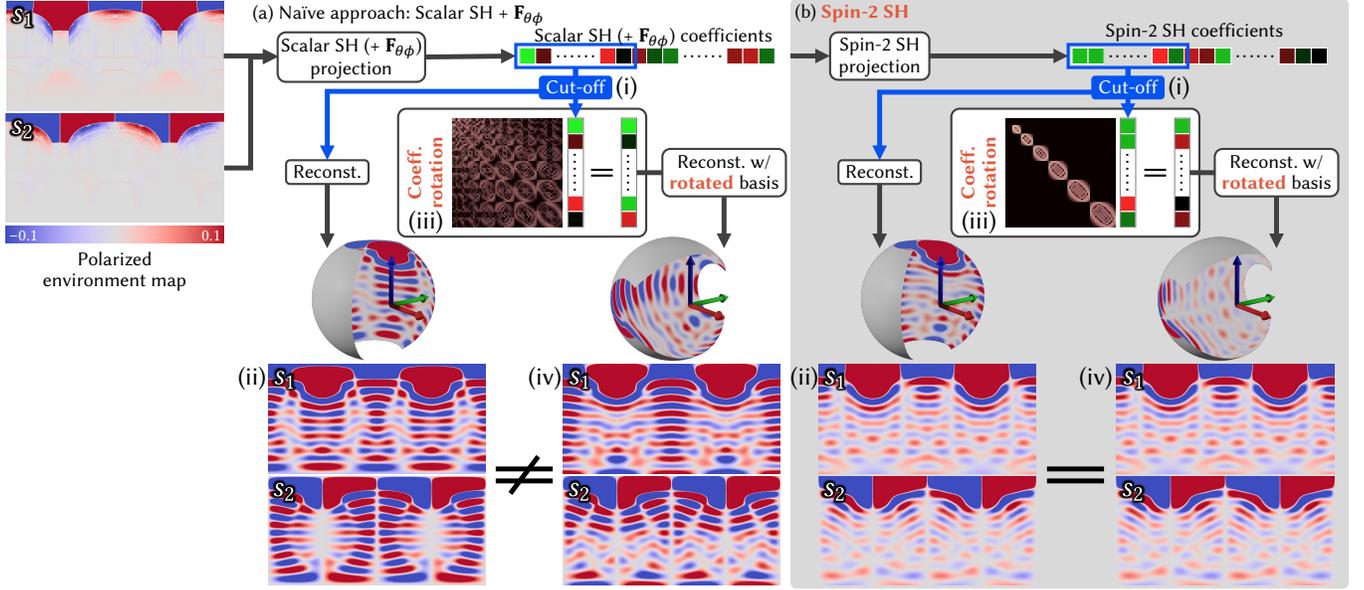
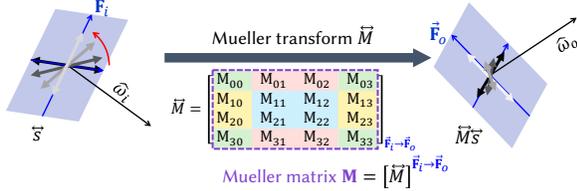


Fig. 11. Numerical validation of rotation invariance. First, get band-limited Stokes vector field (polarized environment map) (ii) from finite coefficients (i) under each basis: (a) a naïve approach that combines scalar SH with the $\theta\phi$ -frame field and (b) spin-2 SH from Equations (26a) and (29). Then (iii) applying rotation transform in the frequency domain, i.e., multiplying a coefficient matrix for a rotation (See Figure 10 for more details). Finally, the rotated coefficient vectors are reconstructed, and the inverse rotation is applied in the final angular domain. Then, while (a) the naïve approach gives inconsistent results, (b) our spin-2 SH give rotation-invariant results.

(a) Mueller transform and Mueller matrix



(b) Complex pair separation for Spin 2-to-2 Mueller matrix

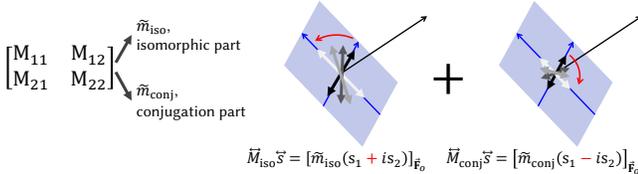


Fig. 12. (a) When polarized light is reflected, the output Stokes vector $\vec{M}\vec{s}$ changes its magnitude and direction, and even direction change is not constant for the general Mueller transform. (b) The spin 2-to-2 block of a Mueller transform from $\hat{\omega}_i$ to $\hat{\omega}_o$ can be represented into two complex numbers: the *isomorphic* part, denoted by M_{iso} , and the *conjugation* part, denoted by M_{conj} . The isomorphic part indicates a Mueller transform, which preserves the self-rotation of the input Stokes vector \vec{s} . In contrast, the conjugation part indicates one which rotates the output $\vec{M}_{\text{conj}}\vec{s}$ CW around $\hat{\omega}_o$ as \vec{s} rotates CCW around $\hat{\omega}_i$.

we find that **spin 2-to-2** part \vec{M} can once more separated into two frame-independent parts \vec{M}_{iso} and \vec{M}_{conj} . However, such separation is not as simple as the full Mueller transform, which splits the matrix into submatrices. For example, taking only M_{11} and replacing

M_{12} , M_{21} , and M_{22} to zero results in different Mueller transforms depending on the choice of frames.

While we find such a separation between theoretical and computational convenience, we also introduce a way to understand it intuitively. Suppose that there is a **spin 2-to-2** Mueller transform $\vec{M} \in \mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o}$, that transforms a spin-2 Stokes vector \vec{s}_i to \vec{s}_o . Imagine the rotation of \vec{s}_i around $\hat{\omega}_i$, as depicted in Figure 12. As shown in Figure 12(a), both magnitude and direction are changed in the output \vec{s}_o . Decomposing it into \vec{M}_{iso} and \vec{M}_{conj} , the output Stokes vectors $\vec{M}_{\text{iso}}\vec{s}_i$ and $\vec{M}_{\text{conj}}\vec{s}_i$ rotate around $\hat{\omega}_o$ in opposite directions without changing their magnitude as shown in Figure 12(b).

To obtain such two parts of the Mueller transform, we define the following conversion functions that convert 2×2 real matrices to complex numbers as

$$\begin{aligned} \mathbb{C}_{\text{iso}}(\mathbf{M}) &:= \frac{M_{11} + M_{22}}{2} + \frac{M_{21} - M_{12}}{2}i, \\ \mathbb{C}_{\text{conj}}(\mathbf{M}) &:= \frac{M_{11} - M_{22}}{2} + \frac{M_{21} + M_{12}}{2}i. \end{aligned} \quad (44)$$

The output pair of complex numbers from this conversion is denoted by $\tilde{m}_{\text{iso}} := \mathbb{C}_{\text{iso}}(\mathbf{M})$ and $\tilde{m}_{\text{conj}} := \mathbb{C}_{\text{conj}}(\mathbf{M})$. Conversely, we can reconstruct to the original 2×2 real matrix as

$$\mathbf{M} = \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{iso}}) + \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{conj}})\mathbf{J}, \quad \text{where } \mathbf{J} := \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (45)$$

Then we can separate the Mueller transform as

$$\begin{aligned} \vec{M}_{\text{iso}} &:= [\mathbb{R}^{2 \times 2}(\tilde{m}_{\text{iso}})]_{\vec{F}_i \rightarrow \vec{F}_o}, \\ \vec{M}_{\text{conj}} &:= [\mathbb{R}^{2 \times 2}(\tilde{m}_{\text{conj}})\mathbf{J}]_{\vec{F}_i \rightarrow \vec{F}_o}. \end{aligned} \quad (46)$$

The key property of this representation is that it converts the product between the matrix \mathbf{M} and the vector $\mathbb{R}^2(z)$ or the matrix $\mathbb{R}^{2 \times 2}(z)$ into complex products as

$$\mathbf{M}\mathbb{R}^2(z) = \mathbb{R}^2(\tilde{m}_{\text{iso}}z + \tilde{m}_{\text{conj}}z^*), \quad (47a)$$

$$\mathbf{M}\mathbb{R}^{2 \times 2}(z) = \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{iso}}z) + \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{conj}}z^*)\mathbf{J}, \quad \forall z \in \mathbb{C} \quad (47b)$$

where \mathbb{R}^2 here denotes the nature conversion from a complex number to a 2-dimensional real vector:

$$\mathbb{R}^2(x + yi) = \begin{bmatrix} x \\ y \end{bmatrix}. \quad (48)$$

Now, we will show that these are well-defined frame-independent quantities. If we rotate the frames $\vec{\mathbf{F}}_i$ and $\vec{\mathbf{F}}_o$ around their z axes by α and β , respectively, the new Mueller matrix under the rotated frames can be evaluated as follows:

$$\begin{aligned} \mathbf{M}' &= \mathbb{R}^{2 \times 2}(e^{-2i\alpha}) \mathbf{M} \mathbb{R}^{2 \times 2}(e^{2i\beta}) \\ &\stackrel{\text{Eq. (45)}}{=} \mathbb{R}^{2 \times 2}(e^{-2i\alpha}) \left[\mathbb{R}^{2 \times 2}(\tilde{m}_{\text{iso}}) + \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{conj}})\mathbf{J} \right] \mathbb{R}^{2 \times 2}(e^{2i\beta}) \\ &\stackrel{\text{Eq. (47b)}}{=} \mathbb{R}^{2 \times 2}(e^{-2i\alpha}) \left[\mathbb{R}^{2 \times 2}(\tilde{m}_{\text{iso}}e^{2i\beta}) + \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{conj}}e^{-2i\beta})\mathbf{J} \right] \\ &= \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{iso}}e^{2i(-\alpha+\beta)}) + \mathbb{R}^{2 \times 2}(\tilde{m}_{\text{conj}}e^{-2i(\alpha+\beta)})\mathbf{J}. \end{aligned} \quad (49)$$

We note here that it is identical to Equation (45) by replacing \tilde{m}_{iso} with $\tilde{m}_{\text{iso}}e^{2i(-\alpha+\beta)}$ and \tilde{m}_{conj} with $\tilde{m}_{\text{conj}}e^{-2i(\alpha+\beta)}$. Since \tilde{m}_{iso} and \tilde{m}_{conj} do not affect each other, this separation is well-defined independent of the choice of frames.

Finally, we obtain the following property:

$$\begin{aligned} \mathbf{M}\mathbb{R}^2(e^{i\vartheta}(s_{i1} + is_{i2})) &= \mathbb{R}^2(e^{i\vartheta}\tilde{m}_{\text{iso}}(s_{i1} + is_{i2})) \\ &\quad + \mathbb{R}^2(e^{-i\vartheta}\tilde{m}_{\text{conj}}(s_{i1} - is_{i2})). \end{aligned} \quad (50)$$

By using this property and Equation (46), it implies that $\vec{\mathbf{M}}_{\text{iso}}$ preserves the rotation direction of the input, and $\vec{\mathbf{M}}_{\text{conj}}$ reverses the rotation direction. We call this **spin 2-to-2** Mueller matrix \mathbf{M} (spin 2-to-2 Mueller transform $\vec{\mathbf{M}}$, respectively) to two complex numbers \tilde{m}_{iso} and \tilde{m}_{conj} ($\vec{\mathbf{M}}_{\text{iso}}$ and $\vec{\mathbf{M}}_{\text{conj}}$, respectively) conversion as *complex pair separation*. And we call each resulting complex number as the *isomorphic part* and *conjugation part*, respectively.

The important property of this separation is that the isomorphic and conjugation parts of the spin 2-to-2 Mueller transform in the angular domain only affect the corresponding spin 2-to-2 submatrix of the coefficient matrix in the frequency domain. Consequently, we can reduce direct 4 integrals in Equation (40) for $p_o, p_i = 1, 2$ into only 2 integrals as separating coefficient matrix as

$$\begin{bmatrix} P_{l_o m_o 1, l_i m_i 1} & P_{l_o m_o 1, l_i m_i 2} \\ P_{l_o m_o 2, l_i m_i 1} & P_{l_o m_o 2, l_i m_i 2} \end{bmatrix} = \mathbb{R}^{2 \times 2} \left(\vec{P}_{l_o m_o, l_i m_i, \text{iso}} \right) + \mathbb{R}^{2 \times 2} \left(\vec{P}_{l_o m_o, l_i m_i, \text{conj}} \right) \mathbf{J}, \quad (51a)$$

$$\vec{P}_{l_o m_o, l_i m_i, \text{iso}} := \int_{\hat{\mathbb{S}}^2 \times \hat{\mathbb{S}}^2} \vec{P}_{\text{iso}}(\hat{\omega}_i, \hat{\omega}_o) {}_2Y_{l_o m_o}^*(\hat{\omega}_o) {}_2Y_{l_i m_i}(\hat{\omega}_i) d\hat{\omega}_i d\hat{\omega}_o, \quad (51b)$$

$$\vec{P}_{l_o m_o, l_i m_i, \text{conj}} := \int_{\hat{\mathbb{S}}^2 \times \hat{\mathbb{S}}^2} \vec{P}_{\text{conj}}(\hat{\omega}_i, \hat{\omega}_o) {}_2Y_{l_o m_o}^*(\hat{\omega}_o) {}_2Y_{l_i m_i}^*(\hat{\omega}_i) d\hat{\omega}_i d\hat{\omega}_o, \quad (51c)$$

where \vec{P}_{iso} and \vec{P}_{conj} denote isomorphic and conjugation parts of $\left[\vec{P}(\hat{\omega}_i, \hat{\omega}_o) \right]_{\vec{\mathbf{F}}_{\theta\phi}(\hat{\omega}_i) \rightarrow \vec{\mathbf{F}}_{\theta\phi}(\hat{\omega}_o)}$, respectively. Note that both ${}_2Y_{l_o m_o}$ is complex conjugated in Equations (51b) and Equation (51c), while ${}_2Y_{l_i m_i}$ is complex conjugated only in Equation (51c). This difference comes from the property of the complex pair separation in Equation (47a).

Based on these formulations, we can now explain the information loss problem in Section 6.2, the reason for using the real coefficient formulation (Equation (34)) rather than complex coefficient formulation (Equation (33)). For fixed order l and degree m , Equation (33) represents a spin-2 Stokes vector field into a single complex coefficient. This implies the spin 2-to-2 block of a Mueller transform field is also represented as a single complex coefficient. Since a single complex number is equivalent to the isomorphic part $\vec{P}_{l_o m_o, l_i m_i, \text{iso}}$ in Equations (51a) and (51b), it only has half the information.

Another further interesting property of the complex pair separation is that we can utilize the commutativity of the complex product, while the original matrix product is non-commutative. It is the main key to proving our polarized spherical convolution theorem, which will be introduced in Section 6.4.

6.3.3 Isotropic pBRDF. Similar to the sparsity condition of isotropic BRDF (Equation (7)), the PSH coefficients of isotropic pBRDF have a sparsity condition. Such constraints can be easily obtained from Equations (51b) and (51c) using $e^{im\phi}$ term in Equation (26a). As a result, the sparsity constraint of the PSH coefficients of isotropic pBRDF can be written as

$$P_{l_o m_o p_o, l_i m_i p_i} = 0, \text{ if } |m_i| \neq |m_o|. \quad (52)$$

Not only the above constraints but there are also additional linear constraints for the **spin 2-to-2** submatrix:

$$\begin{aligned} \vec{P}_{l_o m_o, l_i m_i, \text{iso}} &= 0, \text{ if } m_i \neq m_o, \\ \vec{P}_{l_o m_o, l_i m_i, \text{conj}} &= 0, \text{ if } m_i \neq -m_o. \end{aligned} \quad (53)$$

By using those constraints from isotropy, the complexity of pBRDF coefficient matrix $O(4 \times 4l_{\text{max}}^4)$ reduces to $O(4 \times 4l_{\text{max}}^3)$, which is also similar to the scalar SH coefficients of BRDF.

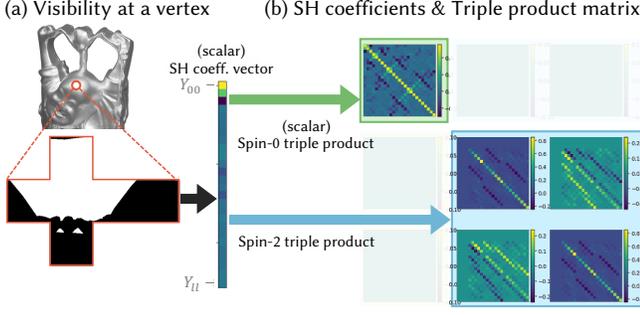


Fig. 13. For precomputation of self-shadow in radiance transfer matrices, (a) a visibility function $V(\hat{\omega})$ at a vertex can be converted into (b) SH coefficients v_{lm} first and then these are expanded to the radiance transfer matrix using the triple product equations. The **spin 0-to-0 submatrix** can be obtained by the conventional triple product described in Equation (55), and the **spin 2-to-2 submatrix** can be obtained by the triple product of spin-0, spin-2, and spin-2 functions described in Equation (56).

6.3.4 Shadowed radiance transfer via triple products. For more realistic rendering, the self-shadow at a vertex of an object can be considered. To compute the self-shadowing radiance transfer coefficients $V_{l_o m_o p_o, l_i m_i p_i}$ on the vertex, we first consider the binary visibility mask $V: \hat{\mathbb{S}}^2 \rightarrow \mathbb{R}$ at the vertex. Then we consider such binary visibility mask as a linear operator $V_{\mathcal{F}}: \mathcal{F}(\hat{\mathbb{S}}^2, \mathcal{S}_{\hat{\omega}}) \rightarrow \mathcal{F}(\hat{\mathbb{S}}^2, \mathcal{S}_{\hat{\omega}})$, which acts on a polarized illumination (Stokes vector field) \vec{f} as $V_{\mathcal{F}}[\vec{f}](\hat{\omega}) := V(\hat{\omega})\vec{f}(\hat{\omega})$. Consequently, the coefficients can be computed as

$$V_{l_o m_o p_o, l_i m_i p_i} = \int_{\hat{\mathbb{S}}^2} \langle \vec{Y}_{l_o m_o p_o}(\hat{\omega}), V(\hat{\omega}) \vec{Y}_{l_i m_i p_i}(\hat{\omega}) \rangle_{\mathcal{S}} d\hat{\omega}. \quad (54)$$

Note that this equation can be also considered as Equation (40) with a Dirac delta Mueller transform.

While Equation (54) can be evaluated in $O(n_{\text{ray}} l_{\text{max}}^4)$ times, it has a useful relationship with the scalar SH coefficients v_{lm} of V , which has $O(n_{\text{ray}} l_{\text{max}}^2)$ complexity. Here note that n_{ray} indicates the number of ray castings in numerical computation (i.e., number of discrete samples for the integrals in Equations (3) and (54)). We can compute the submatrices of $V_{l_o m_o p_o, l_i m_i p_i}$ separately using the identities of triple products of spin-weighted spherical harmonics as depicted in Figure 13.

The **spin 0-to-0** submatrix of $V_{l_o m_o p_o, l_i m_i p_i}$ can be computed using the triple product of three spin-0 (scalar) SH functions. In other words, it can be computed by the scalar SH coefficient of the point-wise product of two scalar SH bases as

$$\int_{\hat{\mathbb{S}}^2} Y_{l_o m_o}^* Y_{l' m'} Y_{l_i m_i} d\hat{\omega}, \quad (55)$$

which has a known analytic formula. Here, l' and m' corresponds to indices of v_{lm} . For the **spin 2-to-2** part, it can be computed using the triple product of one spin-0 and two spin-2 functions. In other words, it can be computed by the spin-2 SH coefficient of the point-wise product of the scalar SH basis and the spin-2 SH basis functions as

$$\int_{\hat{\mathbb{S}}^2} 2Y_{l_o m_o}^* Y_{l' m'} 2Y_{l_i m_i} d\hat{\omega}. \quad (56)$$

Finally, the **spin 2-to-0**, **spin 0-to-2** parts are zero since the point-wise product between spin-0 and spin-0 functions, and spin-0 and spin-2 functions are spin-0 function and spin-2 function, respectively.

Precomputing the shadowed radiance transfer using Equations (55) and (56) rather than direct computation using Equation (54) requires less computation as the number of ray castings for visibility test increases. This is because expanding the coefficient vector v_{lm} to the coefficient matrix $V_{l_o m_o p_o, l_i m_i p_i}$ does not depend on the number of rays. Moreover, triple product relations in Equation (56) will be used to extend our polarized PRT, which enables the dynamic self-shadowing based on previous techniques [Xin et al. 2021; Zhou et al. 2005].

Note that we do not describe the exact computation of the above triple product integrals here, but we only point out that the spin-0 triple product described in Equation (55) has already been used in existing SH-based methods, including Zhou et al. [2005]. The spin-0 and spin-2 triple product described in Equation (56) can be easily implemented once the implementation of Equation (55) is given. For detailed explanation and computation, refer to Supplemental Section 5.7.

6.4 Polarized Spherical Convolution

A strength of the frequency domain analysis (e.g., Fourier transform, spherical harmonics) is that it converts the convolution between two functions into an element-wise product, allowing efficient computation. However, even though spin-weighted SH themselves have been already invented in physics, the spherical convolution on polarized light has not been defined, analyzed, or discussed. Hence, we will start by defining a polarized spherical convolution operation in Section 6.4.1. After that, we will show how to represent polarized convolution kernels as coefficients in Section 6.4.2, by investigating the subspace of PSH. Finally, we propose the polarized spherical convolution theorem in Section 6.4.3, which is the frequency domain analysis of polarized spherical convolution in PSH. Note that we only introduce the theorem statement and its experimental validation in Section 6.4.3, but the derivation of such a theorem is a core contribution of this paper. The detailed derivation and step-by-step proof can be found in Supplemental Section 5.8.

6.4.1 Definition of spherical convolution on Stokes vector fields.

While scalar spherical convolution (Equation (10)) can be naturally defined without considering its rotation equivariance, extending such definition to Stokes vector fields are not trivial. When we try to extend Equation (10) to Stokes vector fields, a question may be asked: What will be the kernel k ? Will it still be a scalar? Otherwise, will it be a Stokes vector field or a Mueller transform field? Although we can answer the question with heuristic choice, we will build a general and standard definition here. To do so, we will start with the *linearity* and *rotation equivariance*, which also defines scalar spherical convolution as described in Supplemental Section 2.6.

Suppose there is a linear and rotation equivariant operator on Stokes vector fields. Since it is a linear operator, it can be characterized as a Mueller transform field $\vec{K}: \hat{\mathbb{S}}^2 \times \hat{\mathbb{S}}^2 \rightarrow \mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o}$, as discussed in Equations (38) and (39). In the beginning, we simply

write the result of the convolution as

$$\int_{\hat{S}^2} \vec{K}(\hat{\omega}_i, \hat{\omega}_o) \vec{f}(\hat{\omega}_i) d\hat{\omega}_i, \quad (57)$$

where \vec{f} is the input Stokes vector field. On the other hand, the rotation equivariance yields:

$$\vec{K}(\vec{R}\hat{\omega}_i, \vec{R}\hat{\omega}_o) = \vec{R}_M \left[\vec{K}(\hat{\omega}_i, \hat{\omega}_o) \right] \quad \forall \vec{R} \in \vec{SO}(3), \quad (58)$$

where $\vec{R}_M[\cdot]$ is the rotation on Mueller transforms. Here, $\vec{R}_M[\cdot]$ is defined as the composition of three Mueller transforms \vec{R}_S , \vec{K} , and \vec{R}_S^{-1} via matrix multiplication as

$$\vec{R}_M \left[\vec{K} \right] := \vec{R}_S \vec{K} \vec{R}_S^{-1}, \quad (59)$$

where \vec{R}_S is defined in Equation (19).

Now we will define the corresponding *kernel* from the above linear and rotation equivariant operator. Moving back to the scalar spherical convolution, the kernel can be obtained by using Supplemental Equation (75). In particular, the scalar spherical convolution kernel can be obtained from the output of the convolution operation when the input source $f(\hat{\omega}_i)$ is a Dirac delta at the north pole $\delta(\hat{\omega}_i, \hat{z}_g)$. Similarly, we can naturally extend this to the polarization by using a Dirac delta Stokes vector $\vec{f}(\hat{\omega}_i) = \vec{s}_i \delta(\hat{\omega}_i, \hat{z}_g)$ with any Stokes vector $\vec{s}_i \in \mathcal{S}_{\hat{z}_g}$. Substituting this Dirac delta Stokes vector into Equation (57), we can get the resulting Stokes vector field as

$$\vec{K}(\hat{z}_g, \hat{\omega}_o) \vec{s}_i. \quad (60)$$

Here we have two choices to define the kernel, the Stokes vector fields in Equation (60) itself with a fixed \vec{s}_i or a Mueller transform $\vec{K}(\hat{z}_g, \hat{\omega}_o)$ as a function of $\hat{\omega}_o$. Among the choices, it is natural to choose the latter one since the former choice cannot store all the information of the spin **0-to-0**, **0-to-2**, **2-to-0**, and **2-to-2** parts of \vec{K} . Finally, we define the kernel using rotation equivariance in Equation (58) to $\vec{K}(\hat{z}_g, \hat{\omega}_o)$, which reduces $\vec{K}(\hat{z}_g, \hat{\omega}_o)$ into a Mueller matrix as a function of single zenith angle θ . As a result, the kernel is defined as

$$\mathbf{k}(\theta) := \left[\vec{K}(\hat{z}_g, \hat{\omega}_{\text{sph}}(\theta, \phi)) \right]_{\vec{F}_{\theta\phi}(0, \phi) \rightarrow \vec{F}_{\theta\phi}(\theta, \phi)} \in \mathbb{R}^{4 \times 4}, \quad (61)$$

which is independent of the choice of ϕ . Note that when evaluating the above equation to obtain the numeric Mueller matrix from the Mueller transform \vec{K} under the (θ, ϕ) , we have to consider the alignment between incident and outgoing frames. In particular, we have to rotate the incident frame by ϕ to get $\vec{F}_{\theta\phi}(0, \phi)$, so that the incident and outgoing frames are always aligned along the great circle of the constant ϕ , as illustrated in Figure 14(a).

Finally, by using the defined kernel above and reformulating Equation (57), we can define the *polarized spherical convolution*, which is the spherical convolution on Stokes vector fields as follows.

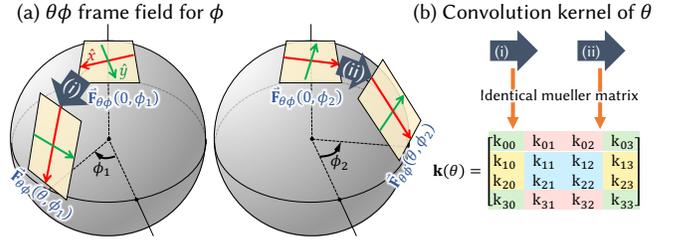


Fig. 14. We propose the concept of spherical convolution of Stokes fields. A convolution kernel is defined as Mueller transform as a function of a single direction $\vec{k}(\hat{\omega}) \in \mathcal{M}_{\hat{z}_g \rightarrow \hat{\omega}}$. Due to rotation equivariance, its numeric Mueller matrix under the $\theta\phi$ -frame field has azimuthal symmetry: $\mathbf{k}(\theta) \in \mathbb{R}^{4 \times 4}$. Concretely, the Mueller matrix under $\vec{F}_{\theta\phi}(0, \phi_1) \rightarrow \vec{F}_{\theta\phi}(\theta, \phi_1)$, illustrated as (a) (i), and the Mueller matrix under $\vec{F}_{\theta\phi}(0, \phi_2) \rightarrow \vec{F}_{\theta\phi}(\theta, \phi_2)$, illustrated as (a) (ii) become an identical matrix, as illustrated in (b). Similar to the general Mueller transform field described in Equation (42), the convolution kernel also can be separated in spin 0-to-0, 0-to-2, 2-to-0, and 2-to-2 submatrices (b).

Polarized spherical convolution

A *spherical convolution kernel for Stokes vector fields* is defined as a function $\vec{k}: \hat{S}^2 \rightarrow \mathcal{M}_{\hat{z}_g \rightarrow \hat{\omega}}$ which maps a single direction to a Mueller transform and has an azimuthal symmetry, i.e., its numeric kernel $\mathbf{k}: [0, \pi] \rightarrow \mathbb{R}^{4 \times 4}$ is defined independent of ϕ as:

$$\mathbf{k}(\theta) := \left[\vec{k}(\theta, \phi) \right]_{\vec{F}_{\theta\phi}(0, \phi) \rightarrow \vec{F}_{\theta\phi}(\theta, \phi)}. \quad (62)$$

Then the convolution of the kernel \vec{k} and a Stokes vector field $\vec{f}: \hat{S}^2 \rightarrow \mathcal{S}_{\hat{\omega}}$ is defined as:

$$\begin{aligned} (\vec{k} * \vec{f})(\hat{\omega}) &= \int_{\hat{S}^2} \left(\vec{R}_{\hat{z}_g \rightarrow \hat{\omega}'} \right)_M \left[\vec{k} \left(\vec{R}_{\hat{z}_g \rightarrow \hat{\omega}'}^{-1} \hat{\omega} \right) \right] \vec{f}(\hat{\omega}') d\hat{\omega}' \\ &= \int_{\hat{S}^2} \left[\mathbf{k} \left(\cos^{-1} \hat{\omega} \cdot \hat{\omega}' \right) \right]_{\vec{F}_i \rightarrow \vec{F}_o} \vec{f}(\hat{\omega}') d\hat{\omega}'. \end{aligned} \quad (63)$$

Here, $\vec{R}_{\hat{z}_g \rightarrow \hat{\omega}'}$ is a rotation transform satisfying $\vec{R}_{\hat{z}_g \rightarrow \hat{\omega}'} \hat{z}_g = \hat{\omega}'$, and choices of $\vec{R}_{\hat{z}_g \rightarrow \hat{\omega}'}$ does not affect on the definition of the convolution. \vec{F}_i and \vec{F}_o are local frames at $\hat{\omega}'$ and $\hat{\omega}$, respectively, such that their x axes are aligned along the common great circle of $\hat{\omega}'$ and $\hat{\omega}$.

6.4.2 Polarized SH coefficients for convolution kernels. Recall that the scalar spherical convolution kernel $k: [0, \pi] \rightarrow \mathbb{R}$, which is an operand for convolution, can be converted to scalar SH coefficients. Similarly, we will show how to convert the polarized spherical convolution kernel \vec{k} to the PSH coefficients. While it requires rigorous derivation steps, we only provide a comprehensive observation-based description. Refer to Supplemental Section 5.8 for the full derivation.

We start by considering the polarized spherical convolution kernel \vec{k} as a function. Then the domain of the \vec{k} is simply $[0, \pi] \ni \theta$, and its PSH coefficient has a single l -index and no m -index similar to conventional kernels as Equation (11). On the other hand, the codomain of the \vec{k} is Mueller transforms, and its PSH coefficient has

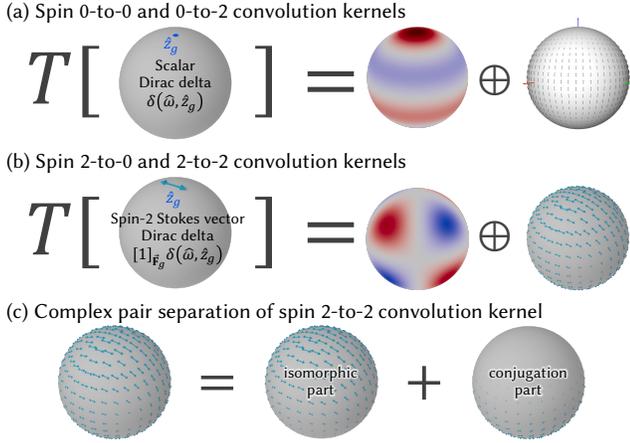


Fig. 15. As conventional convolution kernel operates, (a) spin 0-to-0 and 0-to-2 convolution kernels for polarized spherical convolution can be characterized as the output of a rotation equivariant linear operator on a Dirac delta scalar field. Due to rotation equivariance, output Stokes vector fields are azimuthally symmetric for spin-0 and spin-2 components, i.e., expand with Y_{l0} and ${}_2Y_{l0}$ bases. On the other hand, (b) gives Stokes vector fields with $e^{\pm 2i\phi}$ dependency, i.e., the spin 2-to-0 kernel expands with $Y_{l,-2}^C$ basis. (c) The spin 2-to-2 kernel should even be split into two parts: an isomorphic part and a conjugation part using complex pair separation, which is discussed in Equations (44) and (45).

both p_o and p_i similar to polarized coefficient matrices as described in Section 6.3. Hence, we can write the convolution coefficients in PSH as $k_{l p_o p_i}$, and the remaining question is where these coefficients come from (i.e., from which basis function and which part of the kernel on the angular domain).

We first have a look at the resulting Stokes vector field in Equation (60), which can be written as $\vec{k}(\hat{\omega}) \vec{s}_i$. Taking $\theta\phi$ -frame field, we have:

$$\mathbf{k}(\theta) \mathbf{C}(\phi) \mathbf{s}_i, \quad (64)$$

where \mathbf{C} indicates the frame conversion matrix defined in Equation (14) and $\mathbf{s}_i := [\vec{s}_i]_{\vec{F}_g}$. Note that the frame conversion matrix should be inserted to convert the Stokes component representation \mathbf{s}_i from \vec{F}_g to $\vec{F}_{\theta\phi}(0, \phi)$. Now, similar to the previous derivations, we will consider the spin-0 (s_0, s_3) and the spin-2 (s_1, s_2) parts of the incident Stokes vector \mathbf{s}_i separately.

For the spin-0 Stokes components of \mathbf{s}_i , we can consider $\mathbf{s}_i = [1, 0, 0, 0]^T$ or $[0, 0, 0, 1]^T$. Then the conversion matrix $\mathbf{C}(\phi)$ vanishes so that Equation (64) turns into an azimuthally symmetric full Stokes vector field, which can be expanded by zonal harmonics Y_{l0} (scalar SH kernel coefficients), and a subset of spin-2 SH ${}_2Y_{l0}$. Recall that the ϕ dependency of spin-2 SH ${}_2Y_{lm}$ is characterized as $e^{im\phi}$ in Equation (26a). Thus, similar to scalar SH, spin-2 SH ${}_2Y_{l0}$ with $m = 0$ also have azimuthal symmetry, shown in Figure 8. In summary, Figure 15(a) illustrates the symmetry of the kernel Stokes vector field.

For the spin-2 Stokes components of \mathbf{s}_i , we can consider $\mathbf{s}_i = [0, 1, 0, 0]^T$. Then the s_0 component of Equation (64) turns into:

$$[k_{01}(\theta) \quad k_{02}(\theta)] \mathbb{R}^2 \left(e^{-2i\phi} \right). \quad (65)$$

From the ϕ dependency which comes from $\mathbf{C}(\phi)$, this scalar field can be expanded by $Y_{l,\pm 2}^R$. However, through some derivation details, we find that the best way to describe it is using $Y_{l,-2}^C$. Note that the s_3 component for Equation (64) is also expanded with the same bases. Now considering the spin-2 (s_1, s_2) component of Equation (64), it turns into similarly as follows:

$$\begin{bmatrix} k_{11}(\theta) & k_{12}(\theta) \\ k_{21}(\theta) & k_{22}(\theta) \end{bmatrix} \mathbb{R}^2 \left(e^{-2i\phi} \right). \quad (66)$$

These Stokes vectors in Equations (65) and (66) are also illustrated in Figure 15(b). In addition, we can apply the complex pair separation described in Equations (44) and (45). As a result, we can split Equation (66) into two spin-2 Stokes vector fields:

$$\mathbb{R}^2 \left(\tilde{k}_{\text{iso}}(\theta) e^{-2i\phi} \right) + \mathbb{R}^2 \left(\tilde{k}_{\text{conj}}(\theta) e^{2i\phi} \right), \quad (67)$$

which is also described in Figure 15(c). From the ϕ dependency here, we observe that the isomorphic part expands using ${}_2Y_{l,-2}$ bases, and the conjugation part does using ${}_2Y_{l2}$.

Eventually, the coefficients of the convolution kernel $\mathbf{k}(\theta)$ with respect to PSH are defined as follows.

Convolution coefficients in polarized spherical harmonics

$$k_{l, \{0,3\}, \{0,3\}} := \int_{\hat{\mathbb{S}}^2} Y_{l0}^*(\hat{\omega}) k_{\{0,3\}, \{0,3\}}(\hat{\omega}) d\hat{\omega} \quad (68a)$$

$$\tilde{k}_{l, \{0,3\}, p} := \int_{\hat{\mathbb{S}}^2} Y_{l,-2}^{C,*}(\hat{\omega}) \tilde{k}_{\{0,3\}, p}(\hat{\omega}) d\hat{\omega} \quad (68b)$$

$$\tilde{k}_{l, p, \{0,3\}} := \int_{\hat{\mathbb{S}}^2} {}_2Y_{l0}^*(\hat{\omega}) \tilde{k}_{p, \{0,3\}}(\hat{\omega}) d\hat{\omega} \quad (68c)$$

$$\tilde{k}_{l, \text{iso}} := \int_{\hat{\mathbb{S}}^2} {}_2Y_{l,-2}^*(\hat{\omega}) \tilde{k}_{\text{iso}}(\hat{\omega}) d\hat{\omega} \quad (68d)$$

$$\tilde{k}_{l, \text{conj}} := \int_{\hat{\mathbb{S}}^2} {}_2Y_{l,2}^*(\hat{\omega}) \tilde{k}_{\text{conj}}(\hat{\omega}) d\hat{\omega} \quad (68e)$$

While $\mathbf{k}(\theta)$ is defined on $[0, \pi]$, Equation (68) considers each part of \mathbf{k} as a function defined on $\hat{\mathbb{S}}^2$. Here, the ϕ dependency of each part can be assumed to vanish the ϕ dependency of the entire integrand. For instance, the ϕ dependency of k_{iso} and k_{conj} is considered as $e^{-2i\phi}$ and $e^{2i\phi}$, respectively, as described in Equation (67).

On both hand sides in Equations (68b) and (68c), the subscript p indicates the collection of the indices 1 and 2 in the Mueller matrix \mathbf{k} and the tilde symbol converts it into a single complex number. In Equation (68c), we explicitly write the superscript C to avoid confusion with the real SH Y_{lm}^R . Recall that we mentioned the convolution coefficient can be written as $k_{l p_o p_i}$. It can be constructed directly from the above five types of complex coefficients by converting them into \mathbb{R}^2 or $\mathbb{R}^{2 \times 2}$. However, we found that the complex-valued forms in Equations (68a) to (68e) are more convenient for evaluating the convolution operation in the polarized SH domain.

6.4.3 Polarized spherical convolution in polarized spherical harmonics. Using the convolution coefficients, we can now perform spherical convolution on a Stokes vector field $\vec{f}: \hat{\mathbb{S}}^2 \rightarrow \mathcal{S}_{\hat{\omega}}$ with PSH coefficients. Recall that scalar spherical convolution is evaluated as

an element-wise product between the kernel coefficient and the coefficient of the input, as described in Equation (11). Similarly, polarized spherical convolution is evaluated by an element-wise product with the coefficients f_{lmp} of \vec{f} and other coefficients obtained by flipping the sign of the m index from f_{lmp} as follows.

Polarized spherical convolution theorem

The PSH coefficients of the convolution of a kernel \vec{k} and a Stokes vector field \vec{f} , denoted by $f'_{lmp} := \langle \vec{y}_{lmp}, \vec{k} * \vec{f} \rangle_{\mathcal{F}}$, is evaluated as

$$f'_{l_{m,\{0,3\}}} = \sqrt{\frac{4\pi}{2l+1}} \left[\sum_{p_i=0,3} \tilde{k}_{l,\{0,3\},p_i} f_{lmp_i} + \sum_{m' \in \{\pm m\}} \Re \left(W_{mm'}^{2 \rightarrow 0, *} \tilde{k}_{l,\{0,3\},p} \tilde{f}_{lmp'} \right) \right], \quad (69a)$$

$$\tilde{f}'_{lmp} = \sqrt{\frac{4\pi}{2l+1}} \left[\sum_{p_i=0,3} \sum_{m' \in \{\pm m\}} W_{mm'}^{0 \rightarrow 2} \tilde{k}_{l,p,p_i} f_{lmp'} + \tilde{k}_{l,iso} \tilde{f}_{lmp} + (-1)^m \tilde{k}_{l,conj} \tilde{f}_{l,-mp}^* \right], \quad (69b)$$

where

$$\tilde{f}_{lmp} := f_{lm1} + f_{lm2}i, \quad (69c)$$

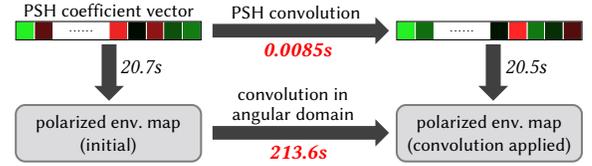
and $W_{mm'}^{2 \rightarrow 0}$ and $W_{mm'}^{0 \rightarrow 2}$ are simple constants taking values of 0, $\pm \frac{1}{\sqrt{2}}$, or $\pm \frac{i}{\sqrt{2}}$ following Supplemental Equations (162) and (156) in Supplemental Section 5.9.

Note that the constant weights $W_{mm'}^{2 \rightarrow 0}$ and $W_{mm'}^{0 \rightarrow 2}$ become zero when $|m| \neq |m'|$, one when $m = m' = 0$, and are evaluated as Supplemental Equations (162) and (156) otherwise. This polarized spherical convolution is nearly an element-wise product for the indices l and m , but similar to 4×4 matrix-vector product for the index p .

Validation between the angular and frequency domains. We here provide a numerical experiment that compares polarized spherical convolution in the angular and frequency domains, and also Supplemental Section 5.9 provides a complete step-by-step derivation to validate our polarized spherical convolution theorem. For the computation in the angular domain, we use an analytic kernel $\mathbf{k}(\theta) = \text{diag}(\pi - \theta)$ for convolution. First, we project the polarized environment map onto the PSH coefficient vector and take the finite (band-limited) coefficient vector for a fair comparison with the computation in the frequency domain. Then, we reconstruct the polarized environment map to the angular domain and perform convolutions on it. For the frequency domain, we first perform convolution on the frequency domain and then reconstruct the polarized environment map. Figure 16(a) depicts the validation pipeline and computation time. The convolution in PSH is significantly faster than the angular domain operation. In addition, the two results are identical, as shown in Figure 16(b).

Validation using rotation averaged pBRDF. Note that the scalar sphere convolution theorem in Equation (11) can be expanded as a

(a) Polarized spherical convolution in angular vs. frequency domain



(b) Resulting images

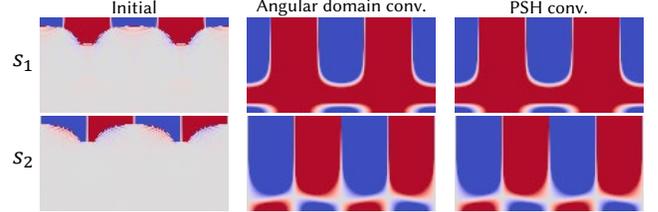


Fig. 16. (a) We compare polarized spherical convolution performed in angular and frequency (polarized SH) domains to validate our polarized spherical convolution theorem in Equation (69). For fair validation, methods on both domains start from the finite (band limited) coefficient vectors of a polarized environment map, which are performed once reconstructing into the angular domain and once the convolution operation is performed. We test frequency levels $l < 100$ and reconstruct pixel numbers $= (l+1)^2$. We observe that convolution in the frequency domain is significantly faster while two operations give identical results (b).

(a) Validation using rotation averaged pBRDF

Average pBRDF over surface normals



(b) Resulting errors

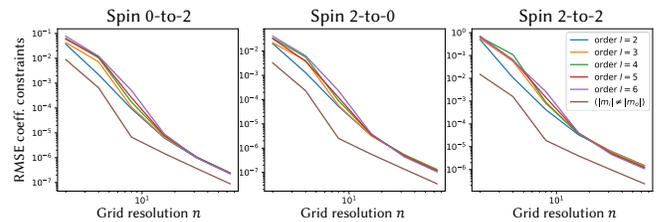


Fig. 17. As scalar sphere convolution theorem in Equation (11) can be expanded as a diagonal coefficient matrix, our polarized spherical convolution theorem in Equations (69a) to (69c) can also be expanded to a coefficient matrix with some linear constraints. Averaging a pBRDF for each normal vector of the material as depicted in (a), we can enforce rotation equivariance to the pBRDF. We validate our convolution theorem by measuring the projection errors of the coefficient matrix of the rotation averaged pBRDF to the linear constraints of convolution operators. (b) provides the error virtually converges to zero as the grid resolution n , which indicates the number of samples of normal vectors, increases. Note that we separate the projection error into each of spin 0-to-2, 2-to-0, and 2-to-2 submatrices (each of three plots), coefficients at $|m_i| = |m_o|$ for each order l (first five curves in the legend), and coefficients at $|m_i| \neq |m_o|$, for better analysis.

coefficient matrix with linear constraints since the entry-wise product of two vectors is equivalent to the product of a diagonal matrix and a vector. Similarly, our polarized spherical convolution theorem from Equations (69a) to (69c) can also be expanded to a coefficient matrix with some linear constraints, described in Supplemental

Section 5.9. We can approximate a pBRDF coefficient matrix into convolution coefficients by averaging each normal vector of the material, as described in Figure 17(a), to ensure rotation equivariance. Figure 17(b) shows the projection error of the rotation averaged pBRDF to the linear constraints of convolution. We can observe that RMS errors virtually converge to zero, which supports our polarized spherical convolution theorem in the frequency domain.

7 PRECOMPUTED POLARIZED RADIANCE TRANSFER

This section presents a real-time rendering pipeline and results of our precomputed polarized radiance transfer, which utilizes PSH and operations in Section 6. How each theoretical component in Section 6 contributes to our rendering pipeline is summarized in Figure 18. Note that our main challenge is related to the linear polarization components s_1 and s_2 ; we omit the circular polarization component s_3 in rendering results since it can be simply processed like total intensity s_0 . We also refer to our supplemental video for real-time rendering results.

Processing polarized environment map. We generate a synthetic polarized environment map using the polarized variant and polarization-aware materials in Mitsuba 3 [Jakob et al. 2022]. In the precomputation stage, we store the PSH coefficient vector of the environment map up to orders (frequency bands) $l \leq l_{\text{high}} = 9$ using Equation (32). Then, in the runtime, these coefficients are rotated to each object frame using Equations (35) and (36) (Section 6.2).

pBRDF projection to PSH coefficients. In the precomputation stage, we also convert data-based isotropic pBRDFs from Baek et al. [2020] into PSH coefficient matrices using Equation (40). When converting and storing the pBRDF coefficient matrix, we utilize the sparsity from the isotropy of pBRDF described in Equations (52) and (53). For the cut-off order, we select $l_{\text{high}} = 9$, same as the environment map (Section 6.3).

Low-high frequency separation. If we increase the order l_{high} , the radiance transfer result will converge to the ground truth. However, the BRDF coefficient matrix requires the complexity of $O(l^4)$, and simply increasing the order by utilizing a full radiance transfer matrix might significantly reduce the computational efficiency. Therefore, we divide the coefficients into low-frequency and high-frequency parts. Then, we apply the $O(l^4)$ radiance transfer matrix only to the low-frequency part rather than utilizing full coefficients. The remaining high-frequency part will be handled in a distinct convolution pipeline. In our implementation, such separation is done in $l_{\text{low}} = 4$, so the low-frequency part contains $0 \leq l \leq 4$ and $5 \leq l \leq 9$ for the high-frequency part.

Radiance transfer using PSH coefficients. Now we rotate the low-frequency part of projected pBRDF coefficients to each vertex normal, yielding the simple unshadowed version of polarized radiance transfer operators. In the runtime, similar to the low-high-frequency separation in pBRDF, the coefficient from the environment map can also be separated by simply splitting the coefficient vector. After that, radiance transfer can be done by a simple matrix-vector product between the radiance transfer operator and the low-frequency part of the environment map, as described in Equation (41) (Section 6.3).

Shadowed transfer using triple product. In the previous paragraph, we propose the unshadowed version of the radiance transfer operator. However, the shadow can also be considered using the triple product as described in Section 6.3.4. To do so, we evaluate visibility for each vertex by casting 2,000 rays from the vertex in the precomputation stage. Then we convert it into the SH coefficient vector and convert it again to a coefficient matrix using SH and PSH triple product in Equations (55) and (56). Finally, applying the matrix product of the projected shadow map coefficient and the unshadowed transfer matrix yields the shadowed transfer matrix that can replace the unshadowed transfer matrix (Section 6.3.4).

Validation with shadowed transfer. We also provide a validation experiment by comparing ours with a physically-based polarization ray tracer, Mitsuba 3. Since our spherical convolution method in the PPRT assumes additional symmetry for pBRDFs and it is already validated in different experiments in Figures 16 and 17, we experiment our shadowed transfer without high-frequency convolution approximation. Figure 19 compares RMSE values between each Stokes component of the rendered images of Mitsuba 3 and our method. We observe that the error for each component decreases close to zero as the cut-off frequency l_{max} increases. Note that the errors will ideally converge to zero when the vertex resolution of the scene additionally increases. We refer to Supplemental Figure 9 for rendered images and difference maps.

Efficient specular appearance using polarized spherical convolution. Now, for the remaining high-frequency part, we project the matrix into convolution coefficients by linear constraints of the convolution following the Sloan et al. [2002]. Note that following Sloan et al. [2002], convolution approximation of a reflected BRDF, which flips the reflected radiance with respect to the surface normal, is preferable to the original BRDF. Thus, we project the product of a reflection operator's coefficient matrix, introduced in Supplemental Section 5.6, and the radiance transfer matrix into a convolution coefficient. This convolution approximation is based on the fact that a specular lobe of a BRDF usually has the peak at the mirror reflection direction so that we can approximate the flipped lobe along the normal as a rotation equivariant one. Then, in runtime, we evaluate PSH values at the reflected direction of the view vector by normals rather than the view vector itself (Section 6.4).

To evaluate the impact of each rendering component, we conduct an ablation study as shown in Figure 20. All experiments are done in the machine with an Intel i9-12900K CPU and an NVIDIA GeForce RTX 4090 GPU. All scenes are rendered in 1024×1024 resolution. We refer to Supplemental Table 2 for detailed specification the scene setups throughout the paper. The *low-order* results only use low-frequency parts with unshadowed radiance transfer. The *+shadow* results use the same order as low-order results, but the shadow is considered. The *+high-order* results use our full pipeline, including the convolution approximation of the high-frequency part. The result shows real-time performance in 102-475 fps for polarized rendering, considering polarized environment lighting. From our shadowed light transport, we can see soft shadows due to environment maps not only in unpolarized s_0 images but also in linear polarization s_1 and s_2 images. When convolution approximation of pBRDF at high order is applied, specular behaviors are enhanced.

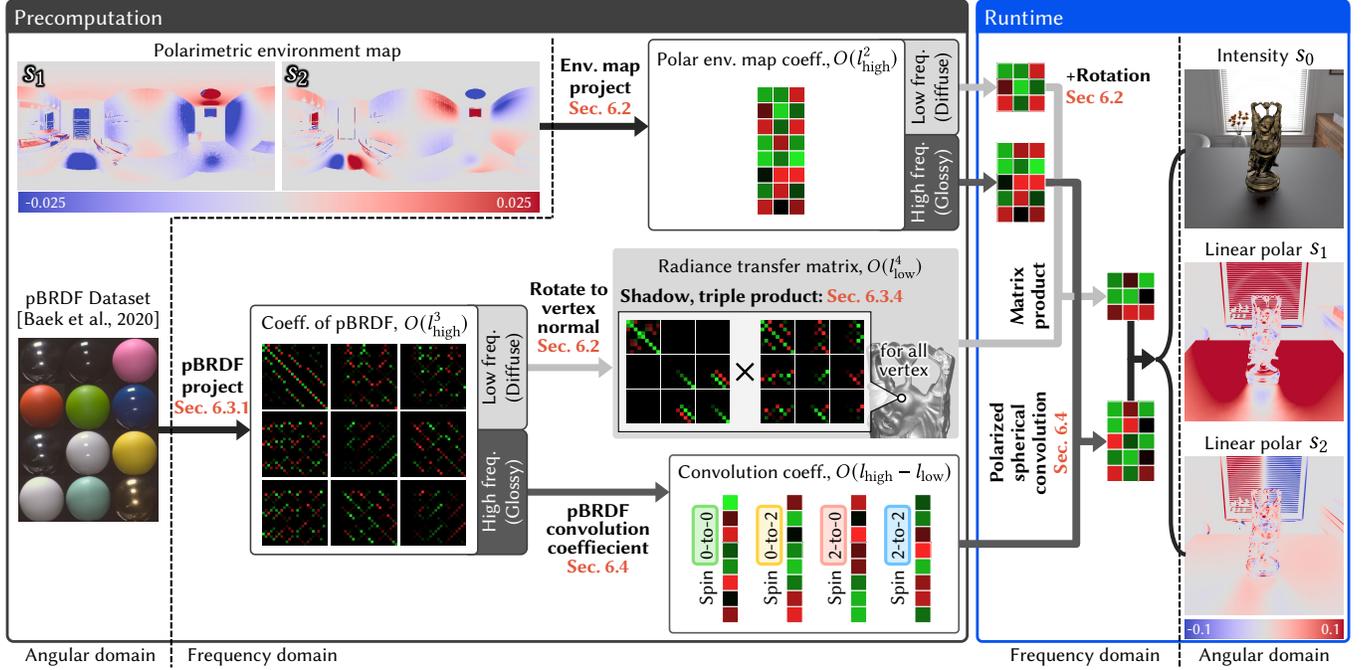


Fig. 18. Our real-time polarized rendering pipeline consists of precomputation and run time parts. Each process described as arrows corresponds to our spin-weighted spherical harmonics operations discussed in Sections 6.2 to 6.4. We use $l_{low} = 4$ and $l_{high} = 9$ for our rendering results.

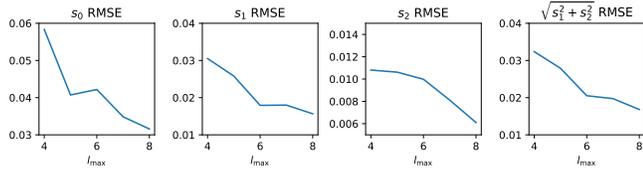


Fig. 19. We validate our real-time polarized rendering with shadowed radiance transfer compared with Mitsuba 3 ray tracer. We report RMSE error, which decreases close to zero as the cut-off frequency l_{max} increases. We refer to Supplemental Figure 9 for resulting rendered images and difference maps.

We also conduct another ablation experiment for convolution approximation, and the result is shown in Figure 21. We provide intensity images through two directions of linear polarizer for better intuition to see specular behavior. If we use only low-order radiance transfer (Figure 21(b)), it is computationally efficient that achieves 480 fps, but it loses some high-frequency appearance. Increasing the order to $0 \leq l \leq 5$ makes the result close to the ground truth, but its performance is degraded to 210 fps. Finally, applying convolution approximation for $5 \leq l \leq 9$ and using full radiative transfer matrix for $0 \leq l \leq 4$ in Figure 21(d) shows a much higher 308 fps but a rich specular appearance than (c), which utilizes orders up to ≤ 5 for the transfer matrix.

8 DISCUSSION

8.1 Choice of PRT Framework

There have been plenty of PRT methods and design choices for the PRT pipeline. For instance, Sloan et al. [2005] store BRDF into SH coefficient (frequency domain) along incident ray direction but

tabulates several outgoing ray directions (angular domain). Sloan et al. [2002] precompute coefficient matrix of self-shadow by directly simulating Equation (54) rather than converting coefficient vector of visibility mask followed by applying SH triple product. However, these choices are totally orthogonal to our main contribution. For a better application of our method to polarization rendering, our PPRT pipeline described in Figure 18 is designed to be aimed to maximize usage of frequency domain operations (theoretical properties for polarized SH). For instance, to build a PPRT method with pBRDF tabulated for each outgoing radiance sample, any method can be plugged in, but to represent pBRDF into a full coefficient matrix, our method is required as described in Section 6.3.

8.2 Physical Constraints

Valid range of Stokes vectors. It is known that the physically valid Stokes vectors should satisfy

$$s_0 \geq \sqrt{s_1^2 + s_2^2 + s_3^2}. \quad (70)$$

There can be many sources of invalidity, such as invalid values in the pBRDF dataset we use [Baek et al. 2020], and the characteristic of the frequency domain method itself. However, since frequency domain analysis decomposes Stokes components into linear factors, such nonlinear inequality is hard to represent in the frequency domain so that the latter source invalidity cannot ideally vanish. We regard Equation (70) as an extended constraint of positivity of radiance in unpolarized radiance transport. Note that SH produce negative values which are related to *ringing* artifacts [Ramamoorthi and Hanrahan 2001b; Sloan et al. 2002], even if the original radiance is positive in any direction. There have been a variety of works

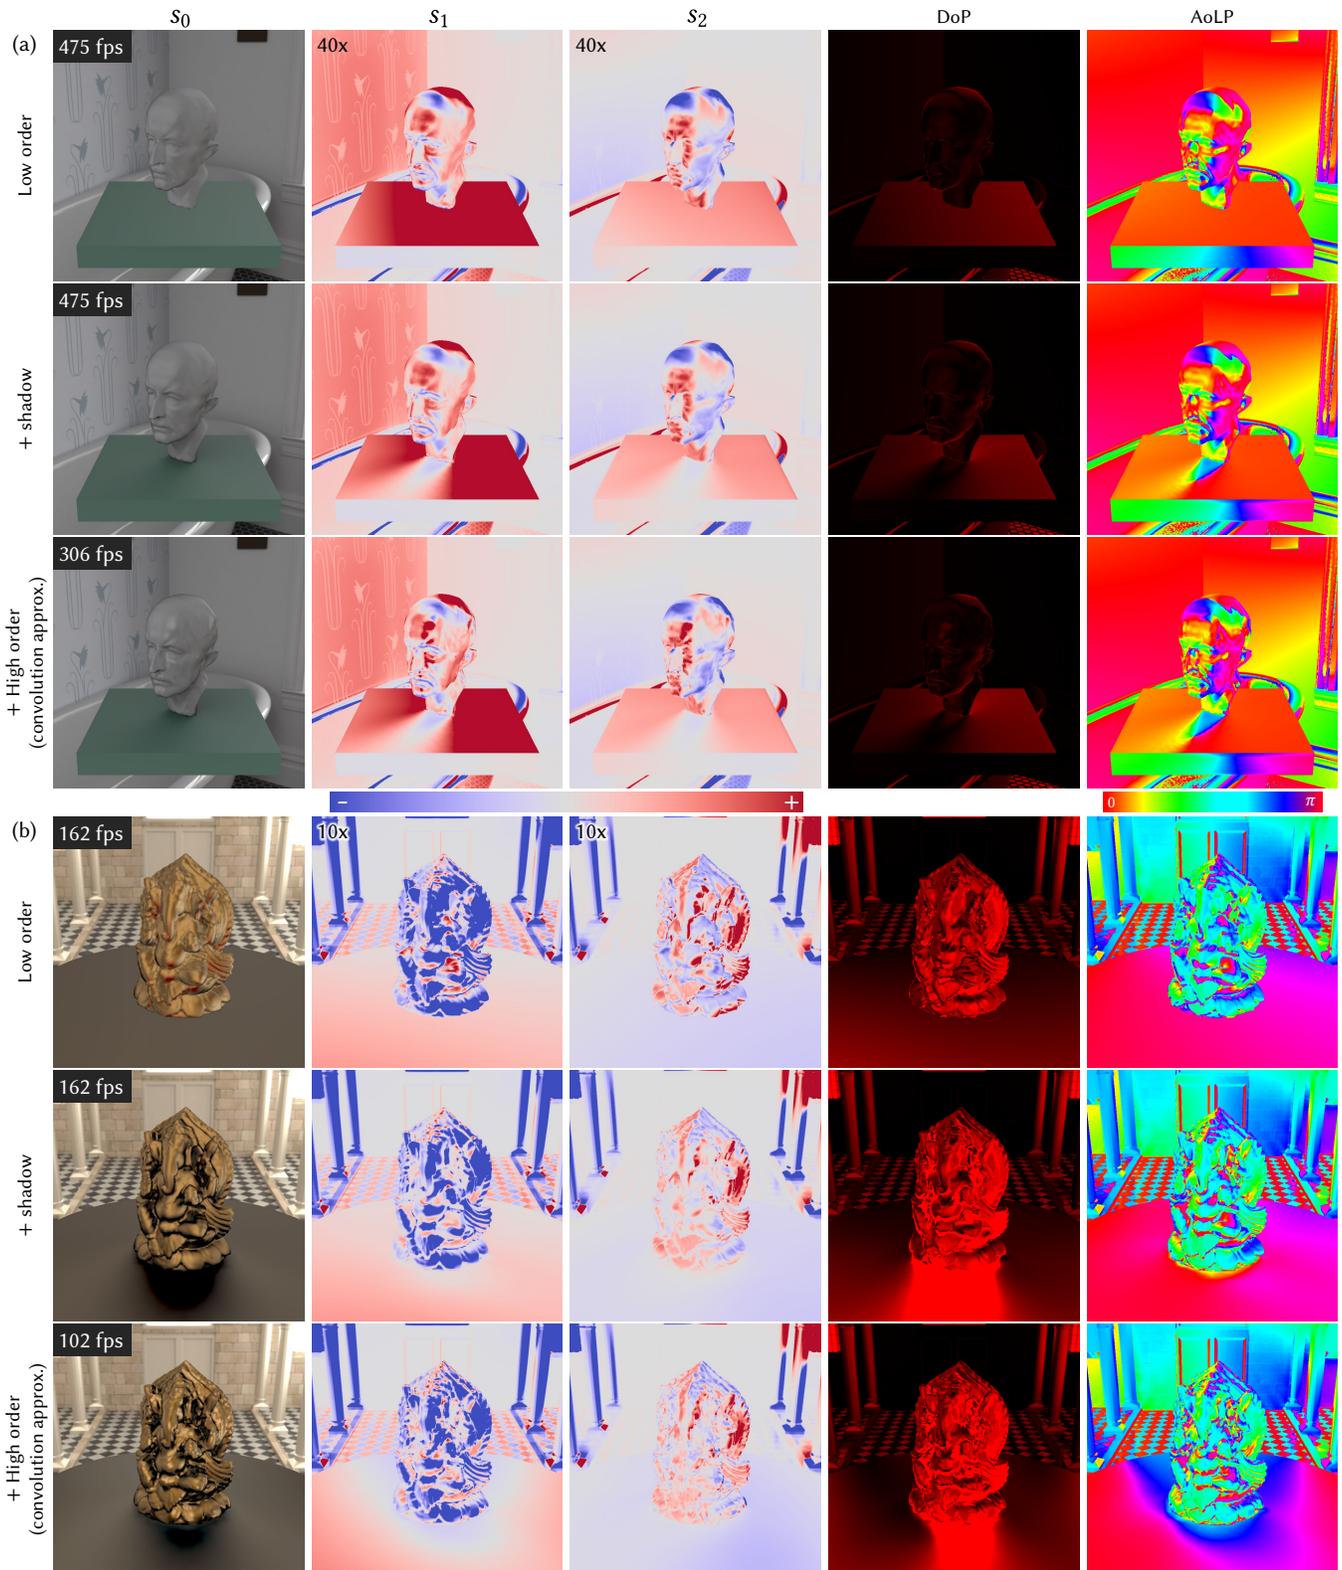


Fig. 20. Ablation study. We provide a real-time frequency-domain environment map lighting method with linear polarization. We use the PSH coefficient vector of a polarized environment map computed in Section 6.2. The coefficient matrix (radiance transfer matrix) of Baek et al. [2020]’s data-based pBRDF computed in Section 6.3, which yields unshadowed transfer shown in the first rows in (a) and (b). We also provide shadowed transfer using Supplemental Section 5.7 and efficient pBRDF approximation for specular appearance using polarized spherical convolution in Section 6.4.

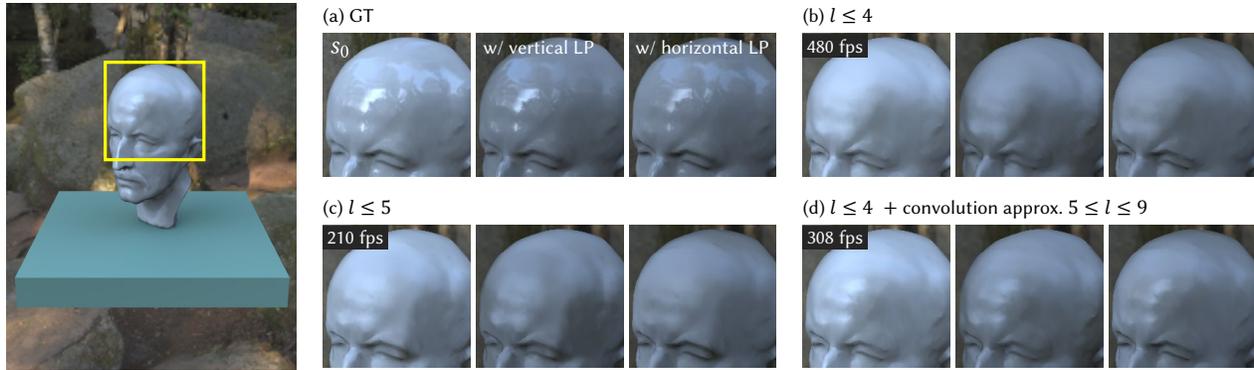


Fig. 21. Although increasing the number of orders l makes the result converge to the actual appearance (a), it suffers from quartic computational complexity as (c) reports less than half fps than (b) even though it uses one more frequency band. Rather than using the full radiance transfer matrix for high orders, (c) we can project the coefficient matrix into convolution coefficients to achieve efficient high order $l < 10$ appearance, which provides much higher performance than the full matrix of $l < 6$ in (b).

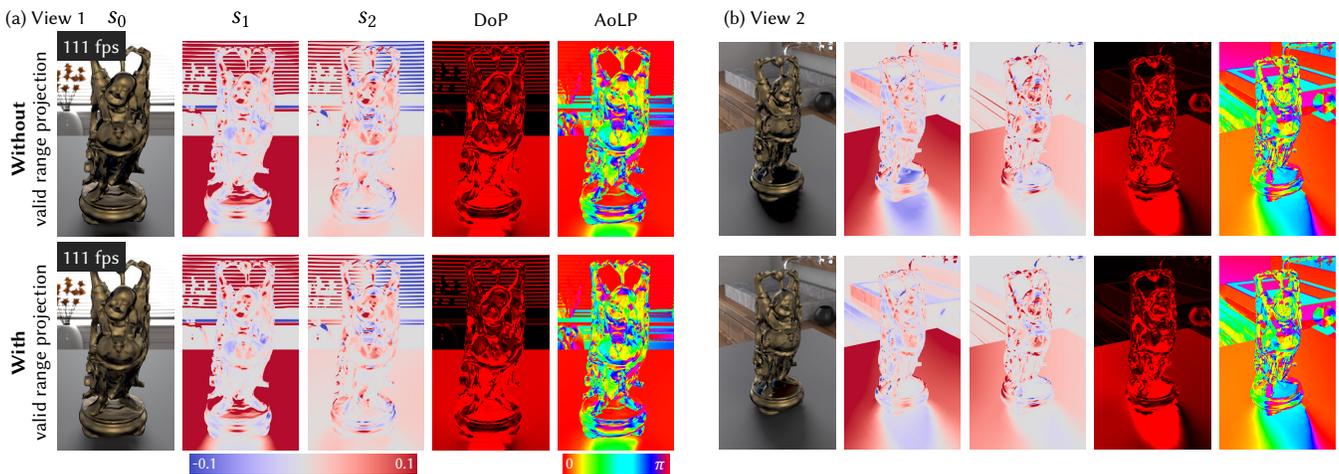


Fig. 22. We can project the rendered image to the physically valid range of Stokes vectors as discussed in Section 8.2, without loss of rendering time. (a) and (b) show two views of the same scene, respectively.

to overcome negativity and ringing artifacts from conventional SH [Berger 2011; Boyd 2001; McClarren et al. 2008; Sloan 2008, 2017]. Extending them for PSH will be an interesting future research direction. For the simplest example, Figure 22 shows the result of our PPRT followed by simply projecting Stokes components to the physically valid range. By enforcing the inequality in Equation (70), s_0 components become slightly brighter, and s_1 and s_2 components do slightly darker, while DoP and AoLP are preserved. Except for this figure, we report our rendering result without this valid range projection to show the direct output of our method.

Constraints for pBRDF. We leave PSH formulation of physical constraints of pBRDF as future work while providing brief discussions. pBRDF should also satisfy energy conservation, but we consider reformulating it into the PSH domain will be a challenging problem since SH and PSH are related to L^2 -norm, but energy conservation is related to L^1 -norm. To the best of our knowledge, energy conservation of SH-projected BRDF is not guaranteed even in unpolarized light transport. Finding the PSH formulation of reciprocity of pBRDF is an interesting problem. Note that flipping the order of direction variables of a pBRDF makes it belong to a different

Mueller space, i.e., $\vec{P}(\hat{\omega}_i, \hat{\omega}_o) \in \mathcal{M}_{\hat{\omega}_i \rightarrow \hat{\omega}_o} \neq \mathcal{M}_{\hat{\omega}_o \rightarrow \hat{\omega}_i} \ni \vec{P}(\hat{\omega}_o, \hat{\omega}_i)$, so that investigating reciprocity of pBRDF requires a solid theoretical foundation. We observe that only a few works address this obscure challenge [Ding et al. 2021; Sekera 1966].

8.3 Difference against the Traditional SWSH

The main difference of this work against traditional SWSH theory consists of the SWSH coefficient formulation for linear operators on Stokes vector fields, including pBRDF and polarized spherical convolution, which is generally equivalent to rotation equivariant linear operators both in the angular and frequency domains.

Before discussing the convolution in more detail, we distinguish two senses to extend conventional convolution on Euclidean domains to others. First, let us denote an operation between two quantities as $k * f = g$. One defines the operation $*$ as an extension of convolution by assuming k and f as the same type of quantities, which we call *correlation*. On the other hand, one can define the operation $*$ to have the same kind of input f and the output g , which we call *convolution* here. In the spherical domain, the output of such *correlation* between two Stokes vector fields should be a

function of single angle [Ng and Liu 1999; Zaldarriaga and Seljak 1997] or a function of rotations, which is not compatible with our PPRT pipeline. In the perspective of image processing and computer graphics, we must extend the *convolution* rather than correlation for Stokes vector fields. As discussed in Section 2.3, existing convolution theories for Stokes vector fields are limited to one [Ng and Liu 1999] (spin 0-to-0 and only real part of the isomorphic part of spin 2-to-2) or six [Garcia and Siewert 1986; Tapimo et al. 2018] degrees of freedoms of kernels at each frequency band, which correspond to subsets of our full kernel formulation described in Equation (62) and Equations (68a) to (68e). To the best of our knowledge, we define new spherical convolution on Stokes vector fields so that it is equivalent to rotation equivariant linear operators. We also establish its PSH formulation, which is applicable to pBRDF approximation.

Our main contributions come from two novel technical details that may be hard to recognize at a high level. First, our *real coefficient formulation* discussed in Section 6.2 and Supplemental Section 5.3.1 is a key part of constructing our PSH formulation of linear operators. It includes our discussion about which sense of linearity of Stokes vectors should be chosen to represent general Mueller matrices. The second technical novelty is the *complex pair separation*, introduced in Section 6.3. It is critical to derive our polarized convolution theorem in Equation (69) through Supplemental Equations (166) to (178). We refer to Supplemental Section 6.2 for more detailed discussion.

8.4 Future Work

Wang and Ramamoorthi [2018]’s analytic SH coefficient for polygonal lights can be directly applied to the PPRT method for unpolarized polygonal lights and polarized material. However, finding analytic formulae for polarized polygonal lights is expected to be a further challenging problem. Xin et al. [2021] found a fast triple product method for conventional SH utilizing FFT. While this method cannot be directly applicable to polarized SH, a similar method is expected to be found using a similar idea.

Applying another PRT pipeline to our PSH theory will be an interesting work. For instance, one can tabulate outgoing directions of pBRDFs rather than using full coefficient matrices as Sloan et al. [2005], or compute shadows in runtime following Zhou et al. [2005] utilizing spin-2 SH triple product introduced in Equation (56). Another possible application is combining physically-based ray tracing for polarized environment map lighting. We can use low-order PSH coefficients for polarized environment maps as Monte Carlo control variates.

In subsurface scattering, an analytic solution of the radiative transfer equation (volume rendering equation) for participating media utilizes the SH up to $l = 1$ [Jensen et al. 2001]⁶, and even at higher orders [Zhao et al. 2014]. Similarly, finding an analytic solution to the polarized radiative transfer equation would be interesting for future work. Without limiting forward rendering, our polarized SH can be used to extend various SH-based methods to polarized states such as acquiring pBRDF based on Ghosh et al. [2007] and Tunwattanapong et al. [2013]’s methods for scalar BRDF, constructing novel polarized spherical CNN, or enhancing polarized radiance field methods

based on existing SH-based methods [Sara Fridovich-Keil and Alex Yu et al. 2022; Verbin et al. 2022] and polarized methods without utilizing basis functions [Dave et al. 2022; Kim et al. 2023].

Extending non-harmonic bases such as wavelets and spherical Gaussians to Stokes vector fields would be a completely different approach from this work, but it will be an interesting future work. Even though they are different types of basis functions, properties of Stokes vector fields, including continuity, discussed in Section 5, must be handled properly.

9 CONCLUSION

While spherical harmonics have been a powerful tool in conventional unpolarized light transport, such basis functions that provide frequency domain analysis for polarized light transport have been absent. We have addressed Stokes vector fields’ challenges regarding frame fields’ choices and their singularities. Also, we have presented spin-weighted spherical harmonics, which provide a rotation invariant orthonormal basis for Stokes vector fields. Combining conventional spin-0 SH for s_0 and s_3 Stokes components and spin-2 SH for s_1 and s_2 components, we have provided our polarized spherical harmonics theory, including linear operator formulation for pBRDF and polarized spherical convolution. Also, we have presented the precomputed polarized radiance transfer, which achieves the first real-time polarized rendering, considering environment lighting and shadows. We expect SWSH and our PSH theory to become helpful in understanding the special nature of polarization and to be used in various applications in future work.

ACKNOWLEDGMENTS

Min H. Kim acknowledges the MSIT/IITP of Korea (RS-2022-00155620, RS-2024-00398830, 2022-0-00058, and 2017-0-00072), Microsoft Research Asia, LIG, and Samsung Electronics.

REFERENCES

- Yunhao Ba, Alex Gilbert, Franklin Wang, Jinfa Yang, Rui Chen, Yiqin Wang, Lei Yan, Boxin Shi, and Achuta Kadambi. 2020. Deep shape from polarization. In *Proc. the European Conference on Computer Vision (ECCV)*. Springer, 554–571.
- Seung-Hwan Baek, Daniel S Jeon, Xin Tong, and Min H Kim. 2018. Simultaneous acquisition of polarimetric SVBRDF and normals. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 268–1.
- Seung-Hwan Baek, Ramesh Raskar, Jinwei Ye, Akshat Dave, Achuta Kadambi, and Huaijin Chen. 2023. Polarization-based visual computing. In *ACM SIGGRAPH 2023 Courses*. 1–1.
- Seung-Hwan Baek, Tizian Zeltner, Hyunjin Ku, Inseung Hwang, Xin Tong, Wenzel Jakob, and Min H Kim. 2020. Image-based acquisition and modeling of polarimetric reflectance. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 139.
- Laurent Belcour, Guofu Xie, Christophe Hery, Mark Meyer, Wojciech Jarosz, and Derek Nowrouzezahrai. 2018. Integrating clipped spherical harmonics expansions. *ACM Transactions on Graphics (TOG)* 37, 2 (2018), 1–12.
- Martin Berger. 2011. Approximate importance sampling of functions reconstructed from spherical harmonics. (2011).
- John P Boyd. 2001. *Chebyshev and Fourier spectral methods*. Courier Corporation.
- Michael Boyle. 2013. Angular velocity of gravitational radiation from precessing binaries and the corotating frame. *Physical Review D* 87, 10 (2013), 104006.
- Edward Collett. 2005. Field guide to polarization. Spie Bellingham, WA.
- Akshat Dave, Yongyi Zhao, and Ashok Veeraraghavan. 2022. Pandora: Polarization-aided neural decomposition of radiance. In *European Conference on Computer Vision*. Springer, 538–556.
- Yuqi Ding, Yu Ji, Mingyuan Zhou, Sing Bing Kang, and Jinwei Ye. 2021. Polarimetric helmholtz stereopsis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5037–5046.
- RDM Garcia and CE Siewert. 1986. A generalized spherical harmonics solution for radiative transfer models that include polarization effects. *Journal of Quantitative*

⁶Scalar irradiance and vector irradiance in that paper corresponds to $l = 0$ and $l = 1$ SH expansion.

- Spectroscopy and Radiative Transfer* 36, 5 (1986), 401–423.
- Abhijeet Ghosh, Shruthi Achutha, Wolfgang Heidrich, and Matthew O’Toole. 2007. BRDF acquisition with basis illumination. In *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 1–8.
- Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A Wilson, and Paul Debevec. 2009. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Computer Graphics Forum* 28, 4, 1161–1170.
- Abhijeet Ghosh, Graham Fyffe, Borom Tunwattana, Jay Busch, Xueming Yu, and Paul Debevec. 2011. Multiview face capture using polarized spherical gradient illumination. *Transactions on Graphics (TOG)* 30, 6 (2011), 1–10.
- Joshua N Goldberg, Alan J MacFarlane, Ezra T Newman, Fritz Rohrlach, and EC George Sudarshan. 1967. Spin-s Spherical Harmonics and δ . *J. Math. Phys.* 8, 11 (1967), 2155–2161.
- Inseung Hwang, Daniel S Jeon, Adolfo Muñoz, Diego Gutierrez, Xin Tong, and Min H Kim. 2022. Sparse ellipsometry: portable acquisition of polarimetric SVBRDF and shape with unstructured flash photography. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–14.
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. <https://mitsuba-renderer.org>.
- Adrian Jarabo and Victor Arellano. 2018. Bidirectional rendering of vector light transport. *Computer Graphics Forum* 37, 6, 96–105.
- Henrik Wann Jensen, Stephen R Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proc. ACM SIGGRAPH 2001*. 511–518.
- Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. 2015. Polarized 3d: High-quality depth sensing with polarization cues. In *Proceedings of the IEEE International Conference on Computer Vision*. 3370–3378.
- Jan Kautz, Peter-Pike Sloan, and Jaakko Lehtinen. 2005. Precomputed radiance transfer: theory and practice. In *ACM SIGGRAPH 2005 Courses*. 1–es.
- Phil B Keegstra, GF Smoot, KM Gorski, G Hinshaw, and L Tenorio. 1997. Generalized spherical harmonics for all-sky polarization studies. In *Astronomical Data Analysis Software and Systems VI*, Vol. 125. 198.
- Youngchan Kim, Wonjoon Jin, Sunghyun Cho, and Seung-Hwan Baek. 2023. Neural Spectro-Polarimetric Fields. In *SIGGRAPH Asia 2023 Conference Papers*.
- Ivan Kuščer and Marijan Ribarič. 1959. Matrix formalism in the theory of diffusion of light. *Optica Acta: International Journal of Optics* 6, 1 (1959), 42–51.
- Christian Lessig and Eugene Fiume. 2008. SOHO: Orthogonal and symmetric Haar wavelets on the sphere. *ACM Transactions on Graphics (TOG)* 27, 1 (2008), 1–11.
- Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, Paul E Debevec, et al. 2007. Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination. *Rendering Techniques* 9, 10 (2007), 2.
- Ryan G McClarren, Cory D Hauck, and Robert B Lowrie. 2008. Filtered spherical harmonics methods for transport problems. In *Proceedings of the 2009 international conference on mathematics and computational methods and reactor physics*.
- Michal Mojzík, Tomáš Skřivan, Alexander Wilkie, and Jaroslav Krivánek. 2016. Bidirectional polarised light transport. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. 97–108.
- Charles Nash and Siddhartha Sen. 1983. *Topology and Geometry for Physicists*. Academic Press.
- Ezra T Newman and Roger Penrose. 1966. Note on the bondi-metzner-sachs group. *J. Math. Phys.* 7, 5 (1966), 863–870.
- Kin-Wang Ng and Guo-Chin Liu. 1999. Correlation functions of CMB anisotropy and polarization. *International Journal of Modern Physics D* 8, 01 (1999), 61–83.
- Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. 2003. All-frequency shadows using non-linear wavelet lighting approximation. In *ACM SIGGRAPH 2003 Papers*. 376–381.
- Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. 2004. Triple product wavelet integrals for all-frequency relighting. In *ACM SIGGRAPH 2004 Papers*. 477–487.
- Robert A Phinney and Robert Burridge. 1973. Representation of the elastic-gravitational excitation of a spherical Earth model by generalized spherical harmonics. *Geophysical Journal International* 34, 4 (1973), 451–487.
- Ravi Ramamoorthi et al. 2009. Precomputation-based rendering. *Foundations and Trends® in Computer Graphics and Vision* 3, 4 (2009), 281–369.
- Ravi Ramamoorthi and Pat Hanrahan. 2001a. An efficient representation for irradiance environment maps. In *Proc. ACM SIGGRAPH 2001*. 497–500.
- Ravi Ramamoorthi and Pat Hanrahan. 2001b. A signal-processing framework for inverse rendering. In *Proc. ACM SIGGRAPH 2001*. 117–128.
- Ravi Ramamoorthi and Pat Hanrahan. 2002. Frequency space environment map rendering. In *Proc. ACM SIGGRAPH 2002*. 517–526.
- Tobias Ritschel, Carsten Dachsbacher, Thorsten Grosch, and Jan Kautz. 2012. The state of the art in interactive global illumination. *Computer graphics forum* 31, 1, 160–188.
- Jérémy Riviere, Ilya Reshetouski, Luka Filipi, and Abhijeet Ghosh. 2017. Polarization imaging reflectometry in the wild. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 1–14.
- Vincent Rossetto. 2009. General framework for multiple scattering of polarized waves including anisotropies and Berry phase. *Physical Review E* 80, 5 (2009), 056605.
- Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2022. Plenoxels: Radiance Fields without Neural Networks. In *CVPR*.
- Joseph JG Scanio. 1977. Spin-weighted spherical harmonics and electromagnetic multipole expansions. *American Journal of Physics* 45, 2 (1977), 173–178.
- Zdeněk Sekera. 1966. Scattering matrices and reciprocity relationships for various representations of the state of polarization. *JOSA* 56, 12 (1966), 1732–1740.
- Peter-Pike Sloan. 2008. Stupid spherical harmonics (sh) tricks. In *Game developers conference*, Vol. 9. 42.
- Peter-Pike Sloan. 2017. Deringing spherical harmonics. In *SIGGRAPH Asia 2017 Technical Briefs*. 1–4.
- Peter-Pike Sloan, Jan Kautz, and John Snyder. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *Proc. ACM SIGGRAPH 2002*. 527–536.
- Peter-Pike Sloan, Ben Luna, and John Snyder. 2005. Local, deformable precomputed radiance transfer. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 1216–1224.
- Shlomi Steinberg, Pradeep Sen, and Ling-Qi Yan. 2022. Towards practical physical-optics rendering. *ACM Transactions on Graphics* 41, 4 (2022), 1–13.
- Shlomi Steinberg and Ling-Qi Yan. 2021a. A generic framework for physical light transport. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–20.
- Shlomi Steinberg and Ling-Qi Yan. 2021b. Physical light-matter interaction in hermite-gauss space. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–17.
- Romuald Tapimo, Hervé Thierry Tagne Kamdem, and David Yemele. 2018. A discrete spherical harmonics method for radiative transfer analysis in inhomogeneous polarized planar atmosphere. *Astrophysics and Space Science* 363, 3 (2018), 52.
- Borom Tunwattana, Graham Fyffe, Paul Graham, Jay Busch, Xueming Yu, Abhijeet Ghosh, and Paul Debevec. 2013. Acquiring reflectance and shape from continuous spherical harmonic illumination. *Transactions on graphics (TOG)* 32, 4 (2013), 1–12.
- Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. 2022. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 5481–5490.
- Jingwen Wang and Ravi Ramamoorthi. 2018. Analytic spherical harmonic coefficients for polygonal area lights. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–11.
- Alexander Wilkie, Claudia Ulbricht, Robert F Tobler, Georg Zotti, and Werner Purgathofer. 2004. An analytical model for skylight polarisation. In *Rendering Techniques*. 387–398.
- Alexander Wilkie, Petr Vevoda, Thomas Bashford-Rogers, Lukáš Hošek, Tomáš Iser, Monika Kolářová, Tobias Rittig, and Jaroslav Krivánek. 2021. A fitted radiance and attenuation model for realistic atmospheres. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–14.
- Alexander Wilkie and Andrea Weidlich. 2012. Polarised light in computer graphics. In *SIGGRAPH Asia 2012 Courses*. 1–87.
- Lifan Wu, Guanyan Cai, Shuang Zhao, and Ravi Ramamoorthi. 2020. Analytic spherical harmonic gradients for real-time rendering with many polygonal area lights. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 134–1.
- Hanggao Xin, Zhiqian Zhou, Di An, Ling-Qi Yan, Kun Xu, Shi-Min Hu, and Shing-Tung Yau. 2021. Fast and accurate spherical harmonics products. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–14.
- Zilin Xu, Zheng Zeng, Lifan Wu, Lu Wang, and Ling-Qi Yan. 2022. Lightweight Neural Basis Functions for All-Frequency Shading. In *Proc. ACM SIGGRAPH Asia 2022*. 1–9.
- Matias Zaldarriaga and Uroš Seljak. 1997. All-sky analysis of polarization in the microwave background. *Physical Review D* 55, 4 (1997), 1830.
- Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order similarity relations in radiative transfer. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–12.
- Kun Zhou, Yaohua Hu, Stephen Lin, Baining Guo, and Heung-Yeung Shum. 2005. Precomputed shadow fields for dynamic scenes. In *ACM SIGGRAPH 2005 Papers*. 1196–1201.