Frame-Free Representation of Polarized Light for Resolving Stokes Vector Singularities

SHINYOUNG YI, KAIST, South Korea and Kyung Hee University, South Korea
JIWOONG NA, KAIST, South Korea
SEUNGMIN HWANG, KAIST, South Korea
INSEUNG HWANG, KAIST, South Korea
MIN H. KIM, KAIST, South Korea

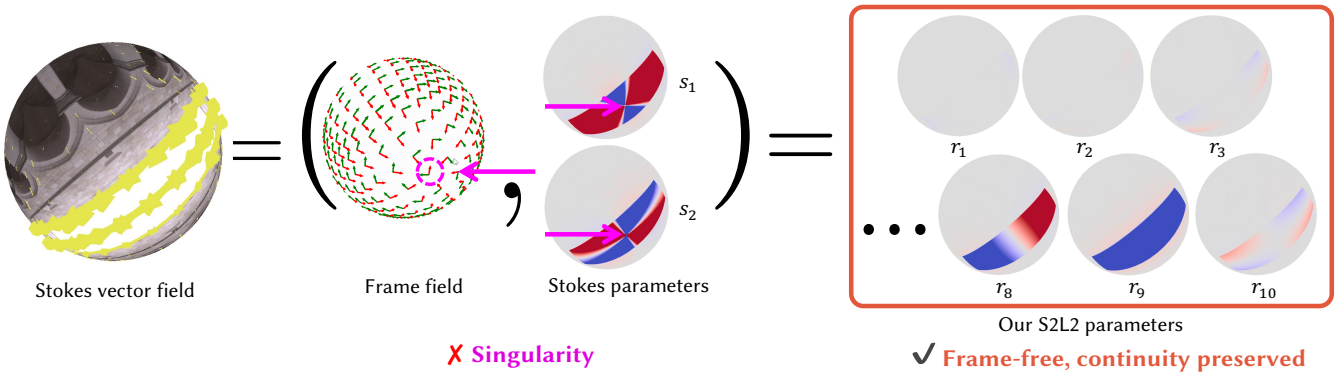

Fig. 1. We propose a novel representation for the polarized intensity of rays, named the *S2L2 representation*. While the conventional Stokes parameter representation suffers from nondeterministic choices of local frames and a singularity problem in the directional domain, our representation provides the first frame-free representation of polarized rays, preserving continuity as actual physical quantities. Such guaranteed properties of our S2L2 representation lead to defining distance and interpolation between polarized rays in different directions, as well as singularity-free polarized radiance fields.

Stokes parameters are the standard representation of polarized light intensity in Mueller calculus and are widely used in polarization-aware computer graphics. However, their reliance on local frames–aligned with ray propagation directions–introduces a fundamental limitation: numerical discontinuities in Stokes vectors despite physically continuous fields of polarized light. This issue originates from the Hairy Ball Theorem, which guarantees unavoidable singularities in any frame-dependent function defined over spherical directional domains. In this paper, we overcome this long-standing challenge by introducing the first frame-free representation of Stokes vectors. Our key idea is to reinterpret a Stokes vector as a Dirac delta function over the directional domain and project it onto spin-2 spherical harmonics, retaining only the lowest-frequency coefficients. This compact representation supports coordinate-invariant interpolation and distance computation between Stokes vectors across varying ray directions–without relying on local frames. We demonstrate the advantages of our approach in two representative applications: spherical resampling of polarized environment maps (e.g., between cube map and equirectangular formats), and view synthesis from polarized radiance fields. In both cases, conventional frame-dependent methods produce singularity artifacts. In contrast, our frame-free representation eliminates these artifacts, improves numerical robustness, and simplifies implementation by decoupling polarization encoding from local frames.

CCS Concepts: • **Applied computing** → **Physics**; • **Computing methodologies** → **Computer graphics**; *Computer vision representations*; • **Human-centered computing** → Visualization; • **Mathematics of computing**;

Additional Key Words and Phrases: polarization in computer graphics, polarized spherical harmonics

ACM Reference Format:
Shinyoung Yi, Jiwoong Na, Seungmin Hwang, Inseung Hwang, and Min H. Kim. 2025. Frame-Free Representation of Polarized Light for Resolving Stokes Vector Singularities. *ACM Trans. Graph.* 44, 6, Article 177 (December 2025), 16 pages. https://doi.org/10.1145/3763276

1 INTRODUCTION

Polarization encodes rich physical information about light that cannot be captured by scalar intensity alone. It enables a range of scene understanding tasks, such as estimating surface normals [Kadambi et al. 2015], separating diffuse and specular components [Ma et al. 2007], and retrieving refractive indices [Baek et al. 2018]. And thus, incorporating polarization into graphics and vision pipelines has gained increasing attention for modeling light–matter interactions more accurately.

While previous methods have modeled polarization using scalar-valued radiance and linear polarizing assumptions [Kadambi et al. 2015; Ma et al. 2007], such formulations remain limited in generality.

Authors' Contact Information: Shinyoung Yi, KAIST, South Korea and Kyung Hee University, South Korea, syyi@khu.ac.kr; Jiwoong Na, KAIST, South Korea, jwna@vclab.kaist.ac.kr; Seungmin Hwang, KAIST, South Korea, smhwang@vclab.kaist.ac.kr; Inseung Hwang, KAIST, South Korea, ishwang@vclab.kaist.ac.kr; Min H. Kim, KAIST, South Korea, minhkim@vclab.kaist.ac.kr.

ACM 1557-7368/2025/12-ART177
https://doi.org/10.1145/3763276

They often rely on specific material models or controlled polarization elements, making them unsuitable for describing general, partially polarized light undergoing complex scattering and multi-bounce interactions. To properly handle these phenomena, radiance must be expressed not as a scalar, but as a Stokes vector–a four-component physical quantity that encodes the full state of polarized light, including linear, circular, and partially polarized forms.

However, using Stokes vectors introduces a fundamental mathematical challenge. Unlike scalar quantities, Stokes vectors exhibit spin-2 symmetry [Yi et al. 2024], meaning their representation depends on a local measurement frame whose z-axis aligns with the light ray direction. This dependence leads to a nontrivial issue: there is no continuous, singularity-free way to define measurement frames over a spherical domain, as guaranteed by the Hairy Ball Theorem [Poincaré 1885]. In practice, this makes it impossible to represent a smoothly varying polarized field across directions using continuous parameters.

This issue has profound implications for practical applications. For instance, interpolating Stokes vectors across ray directions results in discontinuities when crossing frame field singularities. Even simple tasks such as converting between cubemap and equirectangular environment maps suffer from visible artifacts. Furthermore, view synthesis methods based on radiance fields also exhibit persistent sign-flipping artifacts near singularities in the directional domain, independent of camera viewpoint, scene content, or network architecture.

Although Mueller calculus has been widely used for over eight decades in optics, this discontinuity problem has received limited attention—likely because traditional optics focuses on narrow-angle, coherent setups (e.g., lasers), where directionality is well-defined. In contrast, computer graphics and vision inherently depend on directional representations such as environment lighting, radiance fields, and BRDFs, all governed by the rendering equation, which integrates over the hemisphere. As such, the need for a frame-free, continuous representation of polarized light is critical in this domain.

In this work, we address this long-standing limitation by introducing the first frame-free representation of Stokes vectors, termed the S2L2 representation. Our approach eliminates the need for arbitrary frame selection and ensures that continuous changes in polarized radiance correspond to continuous parameter variation in both directions. The S2L2 representation is derived by projecting Dirac delta–like Stokes vectors onto the lowest-order spin-2 spherical harmonics basis [Yi et al. 2024], resulting in a compact and rotation-consistent formulation.

We demonstrate the utility of our representation across several applications. First, we apply S2L2 to interpolate polarized environment maps without introducing artifacts or singularities, enabling stable conversions between cubemap, equirectangular, and other spherical formats (Section 5.1). Second, we address the ambiguity caused by inconsistent conventions in visualizing polarized environment maps by proposing a unified and modular framework (Sections 5.2 and 5.3). Third, we integrate S2L2 into a neural polarimetric radiance field pipeline [Kim et al. 2023], which previously suffered from frame-related artifacts; our method eliminates these issues and improves reconstruction accuracy (Section 6). By resolving discontinuities in directional polarization representation, our approach enables more robust and physically meaningful modeling of light–matter interaction in graphics and vision.

2 RELATED WORK

2.1 Polarization Formulations

Polarized light can be mathematically described using either Jones calculus or Mueller calculus. Jones calculus, which models electric field amplitudes and phase, is suitable for coherent monochromatic light such as lasers. However, it fails to account for incoherent superpositions, which are prevalent in natural illumination and scattering scenarios. In contrast, Mueller calculus represents light intensity using Stokes vectors, making it capable of handling fully and partially polarized light under incoherent conditions. Since the rendering equation in computer graphics–and its analog in vision–integrates radiance over multiple directions, it naturally produces incoherent mixtures. Consequently, most polarization-aware methods in graphics and vision have adopted the Mueller calculus formulation. Prior works on polarization rendering [Jarabo and Arellano 2018; Mojzík et al. 2016], inverse rendering [Baek et al. 2018; Hwang et al. 2022], and material acquisition [Baek et al. 2020] have all represented light as Stokes vectors, treating each ray independently.

2.2 Directional Continuity

A long-standing limitation of the Mueller calculus is its frame-dependence: Stokes vectors must be defined relative to a local measurement frame whose z-axis aligns with the light ray direction. Due to this requirement—and the topological constraint imposed by the Hairy Ball Theorem—it is mathematically impossible to define a globally continuous field of measurement frames on the sphere [Poincaré 1885]. As a result, even if polarized radiance varies smoothly across directions, its representation using Stokes vectors cannot. Because of this, existing methods have largely sidestepped the problem. Most evaluate Stokes vectors only at fixed ray directions [Baek et al. 2018; Hwang et al. 2022; Jarabo and Arellano 2018; Mojzík et al. 2016], avoiding interpolation or continuity across rays. The problem of interpolating polarization in the directional domain remains largely unaddressed.

Kim et al. [2023] took a step toward addressing this by training a neural network to represent Stokes parameters as a function of direction. While this allows novel view synthesis of polarized scenes, it does not resolve the singularity issue: their representation still exhibits sign-flipping artifacts at ray directions corresponding to discontinuities in the underlying frame field. These artifacts stem from the unavoidable need to choose a reference orientation for polarization and persist across viewpoints and network architectures.

2.3 Alternative Representations

Other parameterizations–such as total intensity, degree of polarization, angle of linear polarization, and ellipticity–have been explored [Kim et al. 2023], but they are mathematically equivalent to Stokes vectors and inherit the same frame-dependence and discontinuity. Steinberg and Yan [2021] introduced generalized Stokes

parameters to address partial coherence, but their work does not resolve the topological constraints on directional continuity.

2.4 Spin-Weighted and Polarized Spherical Harmonics

Attempts to properly handle the spin-2 nature of Stokes vectors have been explored in the context of frequency-domain analysis. Stokes vector fields defined over the directional domain can be decomposed in the frequency domain using spin-weighted spherical harmonics [Goldberg et al. 1967; Newman and Penrose 1966]–specifically, spin-0 and spin-2 harmonics [Zaldarriaga and Seljak 1997]. This representation enables smooth and low-frequency fields to be expressed compactly using a small number of coefficients, while also satisfying desirable properties such as rotational invariance.

Spin-weighted spherical harmonics have also been referred to as generalized spherical harmonics [Garcia and Siewert 1986; Kuščer and Ribarič 1959; Tapimo et al. 2018]. More recently, Yi et al. [2024] unified the notation and terminology for the spin-0 and spin-2 cases under the name polarized spherical harmonics (PSH) and demonstrated their application in the context of computer graphics. However, although (polarized) SH-based frequency domain approaches have various applications, including efficient low-frequency shading [Ramamoorthi and Hanrahan 2002; Sloan et al. 2002; Wang and Ramamoorthi 2018; Yi et al. 2024] and BRDF acquisition [Ghosh et al. 2007], their framework does not provide a way to represent individual Stokes vectors in a per-ray and lossless manner.

Our work directly tackles the core challenge of representing Stokes vectors for each ray individually. Building on the insight of the frequency-domain framework, we introduce a novel idea: reinterpreting each light ray as a Dirac delta distribution over directions.

3 BACKGROUND

3.1 Polarization in Mueller Calculus

In Mueller calculus, representing the intensity of a polarized light ray requires the specification of a *local frame* (also referred to as a *measurement frame* or *reference frame*) associated with the ray. This frame is denoted by $\vec{\mathbf{F}} = \begin{bmatrix}\hat{x} & \hat{y} & \hat{z}\end{bmatrix}$, where the $\hat{z}$ axis is aligned with the ray's propagation direction. In this paper local frames are additionally assumed to be orthonormal and right-handed. Within this local frame, the ray's intensity is expressed using a Stokes vector $\mathbf{s} := \begin{bmatrix}s_0 & s_1 & s_2 & s_3\end{bmatrix}^T$. The four Stokes parameters describe the following components: total intensity (s_0), linear polarization along horizontal ($+s_1$) and vertical ($-s_1$) directions, linear polarization along diagonal ($+s_2$) and antidiagonal ($-s_2$) directions, and circular polarization (s_3). For further background and motivation on polarized light, we refer to Collett [2005], Wilkie and Weidlich [2012], and Baek et al. [2023].

When the ray is expressed in a different local frame, denoted as $\vec{\mathbf{G}}$, a conversion of the Stokes parameters is necessary. Since $\vec{\mathbf{G}}$ shares the same z-axis with $\vec{\mathbf{F}}$, their difference can be fully described by a rotation angle ψ, such that $\vec{\mathbf{G}} = \vec{R}_{\hat{z}}(\psi)\,\vec{\mathbf{F}}$, where $\vec{R}_{\hat{z}}(\psi)$ denotes a rotation about the $\hat{z}$ axis by angle ψ. The Stokes parameters in the new frame, $\mathbf{s}' := \begin{bmatrix}s'_0 & s'_1 & s'_2 & s'_3\end{bmatrix}^T$, can then be computed using the following transformation formula:

$$\begin{bmatrix}s'_0\\ s'_1\\ s'_2\\ s_3\end{bmatrix} = \begin{bmatrix}1 & 0 & 0 & 0\\ 0 & \cos 2\psi & \sin 2\psi & 0\\ 0 & -\sin 2\psi & \cos 2\psi & 0\\ 0 & 0 & 0 & 1\end{bmatrix}\begin{bmatrix}s_0\\ s_1\\ s_2\\ s_3\end{bmatrix}. \tag{1}$$

Here, we observe that s_0 and s_3 behave as scalars and remain invariant under frame rotation, whereas s_1 and s_2, which represent linear polarization, constitute the components of a *spin-2 vector* [Newman and Penrose 1966; Yi et al. 2024]. These components follow a distinctive transformation rule that involves twice the rotation angle ψ, a behavior not found in other physical quantities commonly used in light transport or computer graphics.

To distinguish the physical phenomenon of polarized intensity from its numerical representation relative to a specific frame, we adopt the formulation of Yi et al. [2024]. The actual physical quantity of polarized intensity for a given ray is described by a pair consisting of the Stokes parameters $\mathbf{s}$ and the associated local frame $\vec{\mathbf{F}}$, and is denoted as $[\mathbf{s}]_{\vec{\mathbf{F}}}$.

This type of quantity is referred to as a *Stokes vector*, and different representations of the same Stokes vector under different frames are considered equivalent. Specifically, we write:

$$[\mathbf{s}]_{\vec{\mathbf{F}}} = \left[\mathbf{s}'\right]_{\vec{\mathbf{G}}} =: \overleftrightarrow{s} \in \mathcal{S}_{\hat{z}}. \tag{2}$$

The symbol $\overleftrightarrow{s}$, shown with a double-sided arrow, represents the Stokes vector itself, which is invariant under changes of local frames. The notation $\mathcal{S}_{\hat{z}}$ denotes the space of all Stokes vectors associated with the ray direction $\hat{z}$. Conversely, evaluating the Stokes parameters of a given Stokes vector with respect to a specific local frame $\vec{\mathbf{F}}$ is written as:

$$\mathbf{s} = \left[\overleftrightarrow{s}\right]^{\vec{\mathbf{F}}}. \tag{3}$$

Formally, a Stokes vector can be defined as an equivalence class of ordered pairs consisting of Stokes parameters and their associated local frames [Mojzík et al. 2016]. This abstraction is particularly useful for defining physically consistent operations. For example, the addition and inner product of two Stokes vectors $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ associated with the same ray direction can be defined as:

$$\overleftrightarrow{s} + \overleftrightarrow{t} := \left[\left[\overleftrightarrow{s}\right]^{\vec{\mathbf{F}}} + \left[\overleftrightarrow{t}\right]^{\vec{\mathbf{F}}}\right]_{\vec{\mathbf{F}}} \tag{4}$$

$$\left\langle \overleftrightarrow{s}, \overleftrightarrow{t} \right\rangle_{\mathcal{S}} := \left(\left[\overleftrightarrow{s}\right]^{\vec{\mathbf{F}}}\right)^T \left[\overleftrightarrow{t}\right]^{\vec{\mathbf{F}}}, \tag{5}$$

for any choice of local frame $\vec{\mathbf{F}}$. Choosing a different frame yields identical results, as guaranteed by Equations (1) and (2).

Additionally, the rotation of a Stokes vector $\overleftrightarrow{s}$ by a rotation matrix $\vec{R}$ is defined as:

$$\vec{R}_{\mathcal{F}}\left(\overleftrightarrow{s}\right) := \left[\left[\overleftrightarrow{s}\right]^{\vec{\mathbf{F}}}\right]_{\vec{R}\vec{\mathbf{F}}}. \tag{6}$$

To properly account for the local frame dependency of Stokes vectors, it is sometimes useful to treat the scalar and spin-2 vector components separately. The Stokes vector $\overleftrightarrow{s}$ can be decomposed as:

$$\underbrace{\overleftrightarrow{s}}_{\text{(full) Stokes vector}} = s_0 \oplus \underbrace{\overleftrightarrow{s}_{1:2}}_{\text{spin-2 Stokes vector}} \oplus s_3, \tag{7}$$

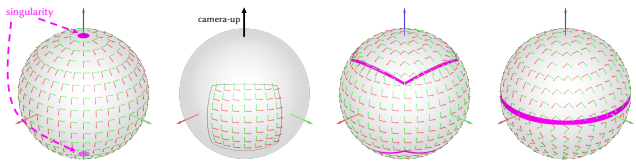

(a) $\theta\phi$-frame field (b) Mitsuba-perspective frame field (c) Cubemap frame field (e) Duff's frame field

Fig. 2. Several frame fields represent Stokes parameters, which are compared with our proposed S2L2 representation in Sections 4 to 6. For detailed definitions and additional frame field variants, refer to Supplemental Figure 4.

where $\overleftrightarrow{s}_{1:2}$, referred to as the *spin-2 Stokes vector*, isolates the directional, frame-dependent portion of $\overleftrightarrow{s}$ by excluding the scalar components s_0 and s_3.[1]

It is important to note that expressions such as $\overleftrightarrow{s}_1$ or $\overleftrightarrow{s}_2$ are invalid, since setting either s_1 or s_2 to zero independently leads to frame-dependent inconsistencies. Such substitutions would result in physically incorrect representations that vary with the choice of local frame.

To prevent ambiguity, the term *full Stokes vector* is sometimes used to explicitly distinguish $\overleftrightarrow{s}$ from its spin-2 component $\overleftrightarrow{s}_{1:2}$. The operations defined in Equations (2), (3), (5), and (6) are also valid for spin-2 Stokes vectors, with similar interpretations.

We refer the reader to Yi et al. [2024] for the construction of a physically well-defined formalism of Mueller calculus. This formalism enables consistent treatment of operations such as addition, inner product, and rotation for Stokes vectors. However, it still relies on a nondeterministic choice of a local frame, and any computation involving a Stokes vector $\overleftrightarrow{s}$ requires decomposing it into the form $[\mathbf{s}]_{\vec{\mathbf{F}}}$ relative to a chosen frame.

The limitations associated with local frames become more pronounced when Stokes vectors are considered as functions over a directional domain. In computer graphics and vision, light intensity is often treated as a function of direction–for instance, in environment lighting, radiance terms in the rendering equation, and neural radiance fields. When polarization is considered, such directional functions must be represented by a Stokes vector field over the unit sphere $\mathbb{S}^2 \coloneqq \left\{\hat{\omega} \in \mathbb{R}^3 \mid \|\hat{\omega}\| = 1\right\}$, defined as $\overleftrightarrow{f} : \mathbb{S}^2 \to \mathcal{S}_{\hat{\omega}}$.

To measure or represent such fields, a local frame must be assigned at each point on $\mathbb{S}^2$, forming what is known as a *frame field*. However, due to the Hairy Ball Theorem [Poincaré 1885], any continuous frame field over the sphere must contain singularities. As a result, it has been fundamentally impossible to represent a continuous Stokes vector field using continuous spherical functions of the four Stokes parameters.

Figure 2 illustrates several common frame fields used in prior work [Jakob et al. 2022; Kim et al. 2023; Yi et al. 2024], along with their associated singularities. For a detailed discussion of the exact definitions, visualization conventions, and additional variants, we refer the reader to Supplemental Figure 4.

[1]Precisely, s_3 depends on handedness of local frames up to the sign. However, we only focus on right-handed frames in this paper.

3.2 Spin-Weighted and Polarized Spherical Harmonics

The challenges introduced by frame field singularities, as discussed in Section 3.1, have recently drawn increased attention in the computer graphics community. Yi et al. [2024] proposed a robust frequency-domain framework based on rotation-invariant basis functions to mitigate this issue. Although the frequency-domain framework itself lies beyond the scope of this work, we briefly summarize their approach, as our frame-free representation of Stokes vectors in the directional domain (introduced in Section 4) is derived from the basis functions proposed by Yi et al. [2024].

Polarized spherical harmonics $\overleftrightarrow{Y}_{lmp} : \mathbb{S}^2 \to \mathcal{S}_{\hat{\omega}}$, parameterized by order l, degree m, and polarization index p, form a basis for the space of Stokes vector fields. The triplet index (l, m, p) belongs to the following set:

$$\begin{aligned}&\left\{(l,m,p) \in \mathbb{Z}^3 \mid l \geq 0,\ |m| \leq l,\ p \in \{0,3\}\right\} \\ \cup\ &\left\{(l,m,p) \in \mathbb{Z}^3 \mid l \geq 2,\ |m| \leq l,\ p \in \{1,2\}\right\}.\end{aligned} \tag{8}$$

$\overleftrightarrow{Y}_{lm0}$ (resp. $\overleftrightarrow{Y}_{lm3}$) denotes a Stokes vector field whose s_0 (resp. s_3) component corresponds to the values of a traditional spherical harmonic (also known as spin-0 SH), while all other components–s_1, s_2, and s_3 (resp. s_0, s_1, s_2)–are identically zero[2].

In contrast, $\overleftrightarrow{Y}_{lm1}$ and $\overleftrightarrow{Y}_{lm2}$ have nonzero values only in the s_1 and s_2 components, and they are constructed from spin-2 spherical harmonics. These basis functions span the space of spin-2 Stokes vector fields.

Since the polarized spherical harmonics form a complete basis, any Stokes vector field $\overleftrightarrow{f} : \mathbb{S}^2 \to \mathcal{S}_{\hat{\omega}}$ can be expressed as an infinite linear combination of them with coefficients f_{lmp}:

$$\overleftrightarrow{f}(\hat{\omega}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{p=0}^{4} \mathrm{f}_{lmp} \overleftrightarrow{Y}_{lmp}(\hat{\omega}). \tag{9}$$

In this expression, basis functions such as $\overleftrightarrow{Y}_{0m1}$, $\overleftrightarrow{Y}_{0m2}$, $\overleftrightarrow{Y}_{1m1}$, and $\overleftrightarrow{Y}_{1m2}$ are defined to be zero, consistent with the index constraints given in Equation (8).

The coefficients f_{lmp}, referred to as the *PSH coefficients* of $\overleftrightarrow{f}$, are computed by projection:

$$\mathrm{f}_{lmp} \coloneqq \int_{\mathbb{S}^2} \left\langle \overleftrightarrow{Y}_{lmp}(\hat{\omega}), \overleftrightarrow{f}(\hat{\omega}) \right\rangle_{\mathcal{S}} \mathrm{d}\hat{\omega}. \tag{10}$$

These coefficients can be interpreted as a frequency-domain representation of polarized intensity, derived from the directional-domain function $\overleftrightarrow{f}$.

Since Equation (9) is an infinite sum, truncating the series by retaining only a finite number of coefficients f_{lmp} up to a certain order l yields a frequency-limited (i.e., smoothed) approximation of the original Stokes vector field $\overleftrightarrow{f}$, with inevitable loss of high-frequency information.

[2]Note that the individual values of the s_1 and s_2 components of a Stokes vector field cannot, in general, be defined prior to choosing a frame field. However, we can state that both s_1 and s_2 components are identically zero across all frames if the spin-2 part is null.

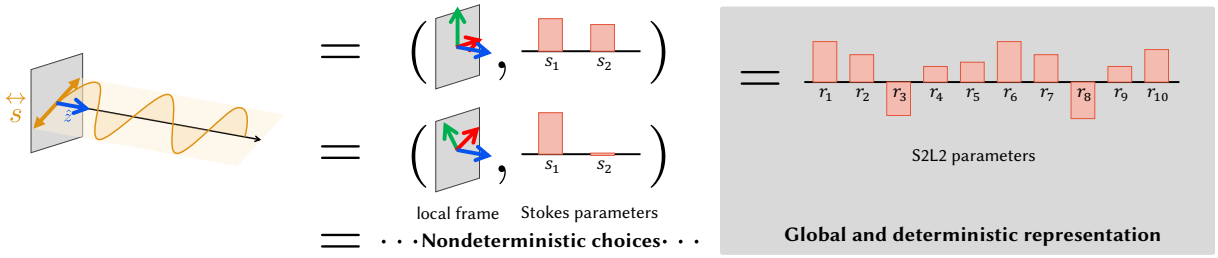

(a) Polarized ray along $\hat{z}$

(b) Conventional representation: Stokes parameters **under a frame**

(c) **Our** representation: S2L2 parameters **without** a frame

Fig. 3. Representation of a polarized ray in numerical parameters. (b) Conventional Stokes parameters require a local frame to be assigned, making them frame-dependent. In contrast, (c) our S2L2 representation eliminates the need for such nondeterministic frame choices. Note that for clarity, this figure omits s_0 and s_3 Stokes parameters to highlight only the frame-dependent components.

A key property of PSH is its *rotation invariance*: smoothing via truncation–by keeping only the PSH coefficients up to a given order l–produces consistent low-frequency approximations regardless of the choice of world coordinates. In linear algebraic terms, this means that the finite-dimensional vector space spanned by $\left\{\overleftrightarrow{Y}_{lmp} \mid |m| \le l,\ 0 \le p < 4\right\}$ is well-defined and invariant under coordinate system rotations.

As with other frequency-domain representations based on harmonic bases–such as traditional spherical harmonics for scalar functions–the PSH representation lacks spatial locality. That is, modifying a single frequency-domain coefficient affects the directional-domain values over the entire sphere.

In Section 4, we depart from the conventional use of harmonic bases and introduce a new approach: the first *frame-free representation* for individual Stokes vectors that preserves directional continuity. Our method leverages the PSH framework, and in particular, exploits the structure of spin-2 spherical harmonics to construct this representation.

4 A FRAME-FREE REPRESENTATION OF STOKES VECTORS

As described in Section 3.1, conventional measurement of Stokes vectors inherently depends on an arbitrary choice of local frames. Even when a frame field is consistently defined over a directional domain, singularities inevitably arise, causing discontinuities between the physically continuous polarization state (the Stokes vector itself) and its scalar components (the Stokes parameters). To address this fundamental limitation, we introduce a novel frame-free representation of Stokes vectors, termed the *S2L2 representation*. Figure 3 illustrates an overview of our method. This representation ensures several desirable properties, notably the preservation of continuity and robustness against singularities. We illustrate the inherent limitations of traditional frame-dependent methods through a simple yet representative application—environment map warping—and demonstrate how our S2L2 representation effectively eliminates the associated artifacts (Section 5). We further demonstrate the utility and generality of our approach by integrating it into a neural polarized radiance field pipeline, which highlights significant improvements in stability and accuracy (Section 6).

We also make our code available on our project website (https://vclab.kaist.ac.kr/siggraphasia2025p2/). To support easy use in future applications, we include in the supplementary material a minimal, stand-alone tutorial Python code that relies only on NumPy and SciPy[3].

4.1 S2L2 representation of Stokes vectors

To construct a frame-free representation of Stokes vectors, we leverage the property that polarized spherical harmonics inherently capture essential characteristics of Stokes vectors, particularly their rotation invariance. As discussed in Section 3.2, PSH is traditionally employed to represent Stokes vector fields distributed over the entire directional domain. In contrast, our objective is to represent individual Stokes vectors associated with distinct ray directions. We observe that any single Stokes vector $\overleftrightarrow{s} \in \mathcal{S}_{\hat{\omega}}$ can be naturally interpreted as a Dirac delta Stokes vector field defined as $\overleftrightarrow{f}(\hat{\omega}') = \overleftrightarrow{s}\delta(\hat{\omega}', \hat{\omega})$. Substituting this into Equation (10), we project this field onto the PSH basis and compute its PSH coefficients up to order $l = 2$:

$$\mathrm{f}_{2mp} = \left\langle \overleftrightarrow{Y}_{2mp}(\hat{\omega}), \overleftrightarrow{s} \right\rangle, \tag{11}$$

for $m = -2$ to 2. Since the s_0 and s_3 Stokes parameters are inherently frame-independent, we can directly include them in our frame-free representation without additional considerations. Thus, our primary focus is restricted to handling cases where $p = 1$ or 2.

Of course, the finite number of coefficients f_{lmp} can only represent frequency-limited information of the original Stokes vector field and thus cannot reconstruct arbitrarily chosen Stokes vector fields exactly. However, our current scenario strongly restricts the original Stokes vector fields to Dirac delta functions. Under this restriction, we observe that the ten parameters f_{2mp} uniquely determine the original single Stokes vector $\overleftrightarrow{s}$. Directly solving Equation (11) to reconstruct $\overleftrightarrow{s}$ from the known f_{2mp} is challenging. Instead, we find that reconstructing a frequency-limited Stokes vector field from f_{2mp} and evaluating it at the direction $\hat{\omega}$ recovers the original vector $\overleftrightarrow{s}$ up to a constant factor. More concretely, denoting this reconstructed frequency-limited Stokes vector field as $\overleftrightarrow{\tilde{f}}$, its value at the original

[3]We refer to `tutorial.ipynb`, or `tutorial.html` if your environment does not support running notebooks.

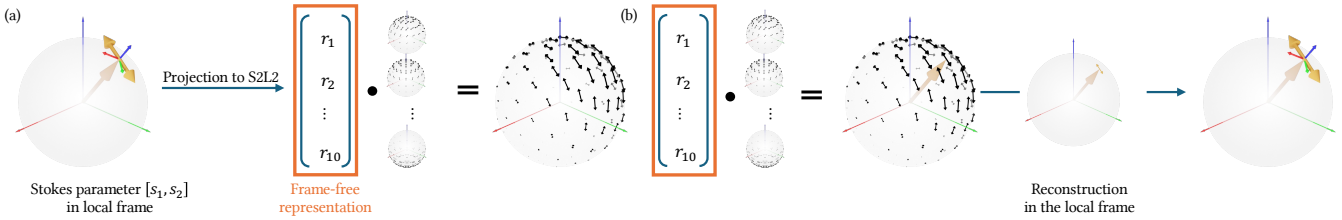

Fig. 4. Our frame-free representation for polarized light intensity. (a) Traditionally, the linear polarization of a ray is described using the Stokes parameters s_1 and s_2, whose values depend on an arbitrary choice of local reference frame. We instead interpret each ray as a Dirac delta Stokes vector *field*, projecting it onto spin-2 spherical harmonics up to order $l = 2$. This yields a set of frame-independent coefficients $r_1, \ldots, r_{10}$ [Yi et al. 2024]. (b) Although these coefficients generally provide a band-limited approximation for arbitrary Stokes vector fields, under the Dirac delta assumption, they allow exact reconstruction of the original s_1 and s_2 values.

ray direction $\hat{\omega}$ can be computed as follows:

$$\begin{aligned}\overset{\leftrightarrow}{f}(\hat{\omega}) &= \sum_{m=-2}^{2}\sum_{p=1}^{2} \mathrm{f}_{lmp}\,\overset{\leftrightarrow}{Y}_{lmp}(\hat{\omega}) \\ &\underset{\text{Eq. (11)}}{=} \sum_{m=-2}^{2}\sum_{p=1}^{2} \left\langle \overset{\leftrightarrow}{Y}_{2mp}, (\hat{\omega}), \overset{\leftrightarrow}{s} \right\rangle_{\mathcal{S}} \overset{\leftrightarrow}{Y}_{2mp}(\hat{\omega}) \\ &\underset{\text{Section C}}{=} \frac{5}{4\pi}\overset{\leftrightarrow}{s}.\end{aligned} \tag{12}$$

Although we omitted several symbolic evaluations in the final step[4], we intuitively understand that the rotation invariance property of spin-2 spherical harmonics ensures this constant factor does not depend on either the direction $\hat{\omega}$ or the vector $\overset{\leftrightarrow}{s}$.

Determine the constant factor. We can now directly treat f_{2mp} defined by Equation (11) as a frame-free representation of a Stokes vector, using Equation (12) as its inverse transform method by simply multiplying the result by the constant factor $\frac{4\pi}{5}$. However, we note that any constant scaling factor c applied as $c\mathrm{f}_{2mp}$ could also define an alternative yet valid convention for this representation. We found that choosing $c = \sqrt{\frac{4\pi}{5}}$ yields a particularly useful property, which will be discussed in Equation (15) in Section 4.2. With this choice, we finalize our frame-free representation, which we name the *S2L2 representation*. Here, "S2" denotes the *spin-2* spherical harmonics, and "L2" indicates the spherical harmonic order $l = 2$.

S2L2 representation of a Stokes vector

$$\mathrm{S2L2}\left(\overset{\leftrightarrow}{s}\right) = \begin{bmatrix} s_0 \\ \sqrt{\frac{4\pi}{5}}\left\langle \overset{\leftrightarrow}{Y}_{2-21}(\hat{\omega}), \overset{\leftrightarrow}{s}\right\rangle_{\mathcal{S}} \\ \sqrt{\frac{4\pi}{5}}\left\langle \overset{\leftrightarrow}{Y}_{2-22}(\hat{\omega}), \overset{\leftrightarrow}{s}\right\rangle_{\mathcal{S}} \\ \vdots \\ \sqrt{\frac{4\pi}{5}}\left\langle \overset{\leftrightarrow}{Y}_{221}(\hat{\omega}), \overset{\leftrightarrow}{s}\right\rangle_{\mathcal{S}} \\ \sqrt{\frac{4\pi}{5}}\left\langle \overset{\leftrightarrow}{Y}_{222}(\hat{\omega}), \overset{\leftrightarrow}{s}\right\rangle_{\mathcal{S}} \\ s_3 \end{bmatrix}. \tag{13}$$

[4]We refer to Section C in our supplemental document for its step-by-step proof.

The 12 parameters resulting from applying S2L2 $\left(\overset{\leftrightarrow}{s}\right)$ to a given Stokes vector $\overset{\leftrightarrow}{s}$ are termed the *S2L2 parameters* of $\overset{\leftrightarrow}{s}$. These parameters are collectively denoted as $\mathbf{r} = \begin{bmatrix} r_0 & \cdots & r_{11} \end{bmatrix}^T$. We use the letter r to denote "representation," subtly indicating its convertibility with the Stokes parameters, represented by the adjacent letter s.

Note that the use of the Dirac delta function serves purely as a conceptual tool for deriving Equation (13). Once the derivation is complete, the equation provides a concrete definition of a per-ray polarization embedding that can be applied directly–without requiring any interpretation involving Dirac deltas. For example, Figure 1 illustrates how a non-Dirac delta distribution of Stokes vectors over the directional domain can be encoded using the S2L2 representation.

Conversely, the inverse S2L2 representation, which reconstructs a spin-2 Stokes vector $\overset{\leftrightarrow}{s} \in \mathcal{S}^2_{\hat{\omega}}$ from the given S2L2 parameters $\mathbf{r} \in \mathbb{R}^{12}$, can be defined as follows.

Stokes vector from S2L2 representation

$$\mathrm{S2L2Inv}(\mathbf{r}; \hat{\omega}) = r_0 \oplus \left(\sqrt{\frac{4\pi}{5}}\sum_{m=-2}^{2}\sum_{p=1}^{2} r_{2m+p+4}\,\overset{\leftrightarrow}{Y}_{2mp}\right)(\hat{\omega}) \oplus r_{11}. \tag{14}$$

Recall that $\oplus$ symbol is defined in Equation (7), and note that the index of r_{2m+p+4} covers the range from 1 to 10. Figure 4 depicts the motivation behind and the process of and process of both the forward and inverse S2L2 representations.

We provide concise formulae for the S2L2 representation in Equations (13) and (14), expressed using the PSH basis functions $\overset{\leftrightarrow}{Y}_{2mp}$ defined by Yi et al. [2024]. For self-contained definitions of these basis functions, please refer to Section B in our supplemental document or the stand-alone Python tutorial (`tutorial.ipynb`) included in our supplementary material.

4.2 Properties

The proposed S2L2 representation of Stokes vectors possesses several desirable properties. We first summarize these properties briefly, followed by detailed explanations for each. Detailed, step-by-step

proofs for each property are provided in Section C of the supplemental document. Although each property is established theoretically, we further validate selected properties through numerical experiments described in Section D, confirming the correctness and effectiveness of our derivation and implementation.

Properties of S2L2 representation

(1) *Left invertible*: For $\overleftrightarrow{s} \in \mathcal{S}_{\hat{\omega}}$,
$$\text{S2L2Inv}\left(\text{S2L2}\left(\overleftrightarrow{s}\right); \hat{\omega}\right) = \overleftrightarrow{s}.$$
(2) *Injectivity for nonzero elements*: For two Stokes vectors $\overleftrightarrow{s} \in \mathcal{S}_{\hat{\omega}_1}$ and $\overleftrightarrow{t} \in \mathcal{S}_{\hat{\omega}_2}$, if $\text{S2L2}\left(\overleftrightarrow{s}\right) = \text{S2L2}\left(\overleftrightarrow{t}\right)$ then either they are zeros or $\overleftrightarrow{s} = \overleftrightarrow{t}$, which also implies $\hat{\omega}_1 = \hat{\omega}_2$.
(3) *Norm preserving*:
$$\sum_{i=0}^{4} s_i^2 = \sum_{i=0}^{11} r_i^2, \quad \text{moreover,} \sum_{i=1}^{2} s_i^2 = \sum_{i=1}^{10} r_i^2. \tag{15}$$
(4) *Continuous*: $\text{S2L2} : \mathcal{S} \to \mathbb{R}^{12}$ and $\text{S2L2Inv} : \mathbb{R}^{12} \times \mathbb{S}^2 \to \mathcal{S}$ are continuous functions.
(5) *Induces a rotation invariant metric*:
$$\left\|\text{S2L2}\left(\vec{R}_{\mathcal{S}}\overleftrightarrow{s}\right) - \text{S2L2}\left(\vec{R}_{\mathcal{S}}\overleftrightarrow{t}\right)\right\|_2 = \left\|\text{S2L2}\left(\overleftrightarrow{s}\right) - \text{S2L2}\left(\overleftrightarrow{t}\right)\right\|_2. \tag{16}$$
(6) *Fiberwise*[5] $\mathbb{R}$*-linear*: If $\overleftrightarrow{s}, \overleftrightarrow{t} \in \mathcal{S}_{\hat{\omega}}$, i.e., two Stokes vectors are associated with the same ray direction $\hat{\omega}$,
$$\text{S2L2}\left(a\overleftrightarrow{s} + b\overleftrightarrow{t}\right) = a\text{S2L2}\left(\overleftrightarrow{s}\right) + b\text{S2L2}\left(\overleftrightarrow{t}\right). \tag{17}$$
(7) *Fiberwise* $\mathbb{C}$*-linear*: Restricting the domain of S2L2 to the set of spin-2 Stokes vectors and regarding the codomain as $\mathbb{C}^5$, it preserves a complex multiple of spin-2 Stokes vectors, defined by Yi et al. [2024].

Left invertibility (1) simply states that the original Stokes vector can be exactly reconstructed from the given 12 S2L2 parameters. Conversely, the equality $\text{S2L2}\left(\text{S2L2Inv}\left(\mathbf{r}; \hat{\omega}\right)\right) = \mathbf{r}$ does not generally hold, because the set $\mathbb{R}^{12}$ is significantly larger than the image of the function S2L2, denoted by $\text{S2L2}\left(\mathcal{S}\right)$. If the provided parameters satisfy $\mathbf{r} \notin \text{S2L2}\left(\mathcal{S}\right)$, i.e., lie outside the valid range, then $\text{S2L2Inv}\left(\mathbf{r}; \hat{\omega}\right)$ projects these parameters onto an appropriate Stokes vector.

In addition to left invertibility, the *injectivity for nonzero elements* property (2) indicates that the S2L2 representation distinguishes not only different Stokes parameter values within the same local frame but also uniquely identifies any pair of Stokes vectors associated with different ray directions, without exception. Note that the zero Stokes vector always maps to zero-valued S2L2 parameters, independent of ray direction.

[5]The set of all Stokes vectors can be viewed as a fiber bundle over the differentiable manifold $\mathbb{S}^2$, in the sense of differential geometry. Fiberwise linearity then refers to linearity within each set of Stokes vectors associated with a given ray direction. However, no background in differential geometry is required to read this paper–Equation (17) gives a self-contained definition of this property.

Finally, the *norm preservation* property (3) ensures that the norm of a Stokes vector can be computed directly in the S2L2 representation without requiring conversion. This also implies that the degree of polarization can be directly evaluated by:
$$\text{DoP} = \frac{\sqrt{\sum_{i=1}^{11} r_i^2}}{r_0}, \tag{18}$$
and similarly, the unpolarized intensity by:
$$r_0 - \sqrt{\sum_{i=1}^{11} r_i^2}. \tag{19}$$
Note, however, that Equations (18) and (19) are valid only if $[r_0 \ldots r_{11}]^T$ is within the valid domain $\text{S2L2}\left(\mathcal{S}\right)$.

Continuity property (4) of our S2L2 representation directly inherits from the underlying polarized spherical harmonics. Specifically, small perturbations in the original Stokes vector result in correspondingly small changes in the S2L2 parameters, with no singular points that could cause numerical instability. This ensures stable numerical computations, which are critical for computational applications involving polarized light.

The S2L2 representation also enables defining a meaningful distance between two Stokes vectors, even when they are associated with different ray directions. For any two Stokes vectors $\overleftrightarrow{s} \in \mathcal{S}_{\hat{\omega}_1}$ and $\overleftrightarrow{t} \in \mathcal{S}_{\hat{\omega}_2}$, the distance[6] is naturally defined as:
$$d_{\text{S2L2}}\left(\overleftrightarrow{s}, \overleftrightarrow{t}\right) = \left\|\text{S2L2}\left(\overleftrightarrow{s}\right) - \text{S2L2}\left(\overleftrightarrow{t}\right)\right\|_2. \tag{20}$$

Thanks to the rotation-invariant nature of polarized spherical harmonics, this distance measure is also *rotation invariant* (5), satisfying:
$$d_{\text{S2L2}}\left(\overleftrightarrow{s}, \overleftrightarrow{t}\right) = d_{\text{S2L2}}\left(\vec{R}_{\mathcal{S}}\overleftrightarrow{s}, \vec{R}_{\mathcal{S}}\overleftrightarrow{t}\right), \tag{21}$$
where $\vec{R} \in \overrightarrow{SO}\left(3\right)$ represents an arbitrary rotation. This invariance guarantees that applications built upon the S2L2 representation yield consistent results, independent of the choice of global frames, even when spherical coordinates are defined relative to different north poles $\hat{z}_g$.

Fiberwise linearity properties (6) and (7) show that, although the S2L2 representation is primarily designed for continuity of Stokes vectors across varying ray directions, it also preserves linear operations on Stokes vectors sharing the same ray direction. Specifically, given two S2L2 parameter sets corresponding to two Stokes vectors defined at an identical ray direction, we can directly perform linear operations on their S2L2 parameters without reverting to the original Stokes parameter representation.

4.3 Results for One-Parameter Motions

While Section B in our supplementary material provides explicit formulae for computing S2L2 parameters from given Stokes vectors under various parameterizations (e.g., ZYZ Euler angles or quaternions), here we demonstrate concrete examples of one-parameter motions of Stokes vectors. These examples visually illustrate the

[6]Strictly speaking, d_{S2L2} is a pseudometric rather than a proper metric, because zero Stokes vectors at different ray directions are considered distinct elements in the Stokes space $\mathcal{S}$, yet their distance remains zero.

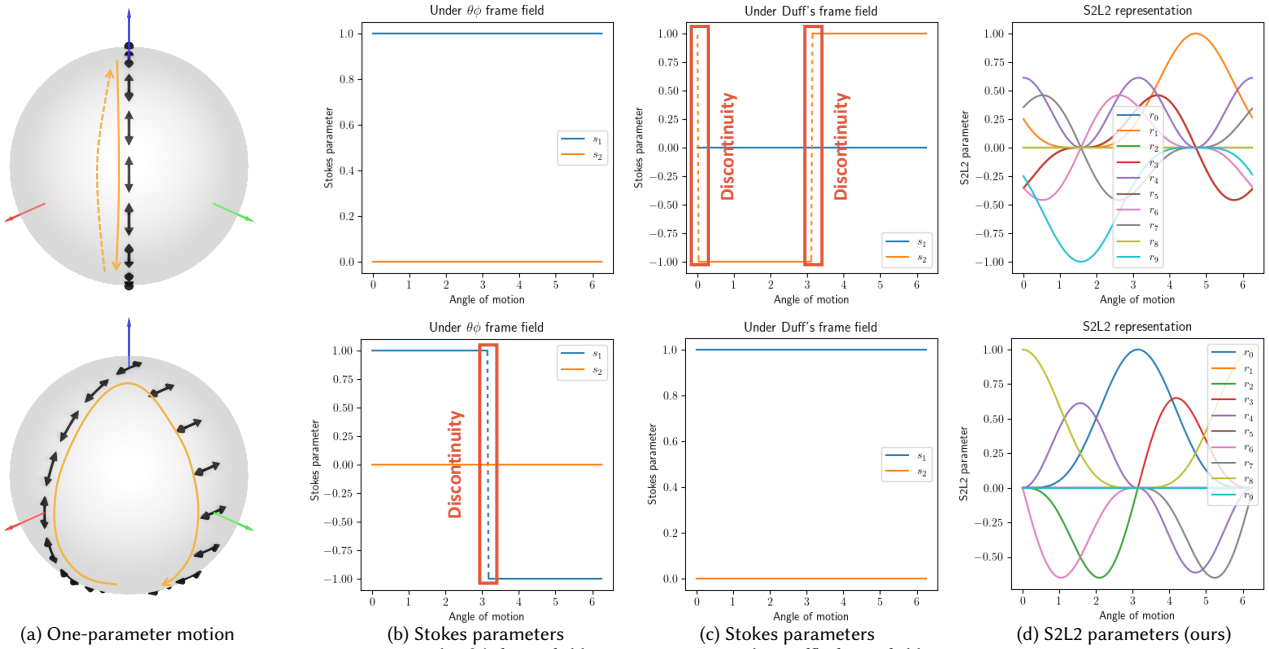

(a) One-parameter motion (b) Stokes parameters under $\theta\phi$-frame field (c) Stokes parameters under Duff's frame field (d) S2L2 parameters (ours)

Fig. 5. Comparison of representations under two continuous motions of polarized rays. (b, c) Frame-dependent Stokes parameters exhibit sign-flipping discontinuities for some motion due to frame field singularities. (d) In contrast, our S2L2 representation remains continuous in both cases.

continuity and robustness of our proposed S2L2 representation, highlighting its significant advantages over conventional frame-dependent representations, which typically exhibit discontinuities.

In Figure 5, we evaluate two continuous one-parameter motions (curves) of spin-2 Stokes vectors, defined as $\overleftrightarrow{s}(\lambda) = [\mathbf{s}(\lambda)]\,\vec{\mathbf{F}}(\lambda)$. We compare three different representations: the standard Stokes parameters under the $\theta\phi$-frame field, the frame field proposed by Duff et al. [2017], and our novel S2L2 parameters. The two motions are defined explicitly as follows:

$$\begin{aligned} \mathbf{s}(\lambda) &= \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \\ \vec{\mathbf{F}}_1(\lambda) &= \vec{R}_{\hat{z}_g}\left(\frac{\pi}{4}\right)\vec{R}_{\hat{y}_g}\left(\lambda + \frac{\pi}{2}\right)\vec{\mathbf{F}}_g, \\ \vec{\mathbf{F}}_2(\lambda) &= \begin{cases} \vec{R}_{\hat{y}_g}(\pi - \lambda)\,\vec{\mathbf{F}}_g & \text{if } \lambda \le \pi \\ \vec{R}_{\hat{x}_g}(\pi - \lambda) & \text{if } \lambda \ge \pi \end{cases}, \end{aligned} \tag{22}$$

where the first and second motions are denoted by $[\mathbf{s}(\lambda)]\,\vec{\mathbf{F}}_1(\lambda)$ and $[\mathbf{s}(\lambda)]\,\vec{\mathbf{F}}_2(\lambda)$, respectively. Note that the second motion is also continuous, as the two formulations match exactly at $\lambda = \pi$, i.e., $\vec{R}_{\hat{y}_g}(\pi - \lambda)\,\vec{\mathbf{F}}_g = \vec{R}_{\hat{x}_g}(\pi - \lambda)\,\vec{\mathbf{F}}_g$ when $\lambda = \pi$. The first and second rows of Figure 5(a) visualize these two motions, respectively.

Figure 5(b) plots the Stokes parameters for each motion using the $\theta\phi$-frame field. Although the first motion passes through the singularity at $\hat{z}_g$, no visible discontinuity in s_1 and s_2 occurs. This continuity arises purely by coincidence, because the angular difference between the frame orientations approaching from opposite directions is exactly π radians, resulting in a smooth 2π phase shift of $[s_1, s_2]^T$. However, such continuity is sensitive to the precise path taken across the singularity. Indeed, in the second motion, where the frame orientation difference deviates from exactly π at $\lambda = \pi$, we clearly observe a discontinuity manifesting as a sudden sign flip in the Stokes parameters.

In Figure 5(c), we present the same two motions using the frame field of Duff et al. [2017]. Here, while the second motion remains continuous, the first motion encounters singularities at positions $\vec{\mathbf{F}}_1\left(\frac{\pi}{2}\right) = \frac{1}{\sqrt{2}}(\hat{x}_g + \hat{y}_g)$ and $\vec{\mathbf{F}}_1\left(\frac{3\pi}{2}\right) = -\frac{1}{\sqrt{2}}(\hat{x}_g + \hat{y}_g)$, causing local frames to flip orientation. This again introduces sign-flipping discontinuities in the Stokes parameters.

In contrast, Figure 5(d) illustrates our S2L2 representation. Clearly, our method maintains a smooth and continuous representation of the parameters for all continuous motions, effectively eliminating discontinuities associated with conventional frame-dependent approaches.

5 APPLICATION: IMAGE PROCESSING OF POLARIZED ENVIRONMENT MAPS

In this section we demonstrate the application of our S2L2 representation to image processing tasks involving polarized environment maps. Section 5.1 discusses how to interpolate Stokes vectors across different ray directions. In Section 5.2, we show that interpolating Stokes vectors using our S2L2 representation enables robust spherical resampling of polarized environment maps. Finally, in Section 5.3, we demonstrate our representation allows to compute image gradients without suffering from discontinuities induced by frame fields.

5.1 Stokes Vector Interpolation

Based on the defined S2L2 representation and its inverse, we can now define an interpolation between two Stokes vectors using the S2L2 representation. Consider two Stokes vectors $\overleftrightarrow{s} \in \mathcal{S}_{\hat{\omega}_1}$ and $\overleftrightarrow{t} \in \mathcal{S}_{\hat{\omega}_2}$. Due to the frame-independent nature of the S2L2 representation, interpolation can be performed even when the two Stokes vectors belong to different Stokes spaces, i.e., $\hat{\omega}_1 \neq \hat{\omega}_2$. This interpolation is formulated as:

$$\text{S2L2Inv}\left((1-\alpha)\,\text{S2L2}\left(\overleftrightarrow{s}\right) + \alpha\text{S2L2}\left(\overleftrightarrow{t}\right); \hat{\omega}\,(\alpha)\right), \tag{23}$$

where $\hat{\omega}(\alpha)$ denotes an interpolation between $\hat{\omega}_1$ and $\hat{\omega}_2$, such as linear or spherical linear interpolation.

Property (5), discussed in Section 4.2, implies this interpolation yields consistent results independent of the choice of world coordinates. Properties (6) and (7) imply that if $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ are uniformly scaled or rotated around their respective ray directions, the interpolation result undergoes the equivalent transformation.

Evaluating the *correctness* of our interpolation method is nontrivial since Stokes vectors associated with different ray directions reside in distinct tangent planes (strictly, distinct Stokes spaces), and such interpolation across directions has not been systematically addressed in previous literature. A natural reference method is *parallel transport*, commonly used for interpolating tangent vectors–especially between nearby tangent planes–on curved surfaces. Given two Stokes vectors associated with directions $\hat{\omega}_1$ and $\hat{\omega}_2$, parallel transport defines local frames at each point along the geodesic curve connecting these directions, maintaining consistent angular relationships along the path. Linear interpolation of Stokes parameters within these local frames thus becomes feasible. A simplified example, where both Stokes vectors have the same magnitude and angle relative to the geodesic, illustrates that the interpolated vector retains this consistency, as depicted in Figure 6(b).

However, parallel transport-based interpolation becomes impractical for scenarios involving interpolation across *multiple* directions, such as spherical resampling. Parallel transport requires distinct local frames for each pair of directions, resulting in a pairwise—and hence nondeterministic—representation of Stokes vectors. In spherical resampling, each original pixel contributes to multiple interpolation neighborhoods, each requiring reevaluation under different local frames, causing significant computational overhead and reduced modularity. Consequently, this paper focuses on comparing interpolation methods that represent Stokes vectors independently of such pairwise frame choices.

A common alternative assigns a global frame field per ray direction across the directional domain and interpolates the Stokes parameters with respect to these per-direction frames. However, as shown in Figure 6(a), this approach introduces dependency on the world coordinate system due to frame field construction, resulting in discontinuities and unnatural interpolation behaviors when crossing frame-field singularities.

Figure 6(c) presents interpolation results across 20,000 different world coordinate poses, obtained by rotating 200 spherical Fibonacci directions with 100 uniformly sampled rotation angles. All three tested frame fields demonstrate significant sensitivity to world coordinates rather than actual signal content, as the standard deviations exceed mean values. Moreover, the average values deviate substantially from the parallel transport reference (dashed line), corresponding to $(s_1, s_2) = (1, 0)$.

Conversely, interpolation using the proposed S2L2 representation, illustrated in the fourth column of Figure 6(c), exhibits minimal variance (standard deviation $< 10^{-14}$) across different world coordinate poses. Specifically, the s_2 component consistently matches the parallel transport result. Although the s_1 component deviates slightly from the reference, applying a normalization step by dividing by the interpolated vector norm aligns closely with the parallel transport interpolation, as demonstrated in the final column.

This renormalized interpolation method can be implemented modularly by concatenating the direction vectors as additional channels alongside the RGB and S2L2 parameter channels. Standard interpolation methods can then directly operate on this extended multichannel data, with a subsequent normalization post-processing step. We refer to this variant as *renormalized interpolation*.

5.2 Spherical Resampling

One important application of Stokes vector interpolation is spherical image resampling. As briefly discussed in Section 5.1, spherical resampling involves converting images between different spherical representations, such as cubemaps, equirectangular maps, and perspective views, by interpolating pixel values across spherical grids. In this context, each pixel from the original grid generally contributes to multiple pixels in the resampled grid. Our resampling is implemented via backward mapping from the target image grid and bilinear interpolation on the source image, as described in Figure 7. Specifically, Equation (23) is naturally extended to:

$$\text{S2L2Inv}\left(\sum_{i=1}^{4} \lambda_i \text{S2L2}\left(\overleftrightarrow{s_i}, \hat{\omega}\right)\right), \tag{24}$$

where the weights λ_i (with $\sum_{i=1}^{4} \lambda_i = 1$) correspond to bilinear interpolation weights, and each $\overleftrightarrow{s_i}$ denotes the Stokes vector at a neighboring pixel. This formulation allows smooth, rotation-invariant interpolation across directional changes, even when involving more than two contributing elements. In contrast, using conventional frame-based Stokes parameter representations for interpolation can introduce artifacts, particularly when interpolations span across frame-field singularities, leading to inconsistent or discontinuous results.

Cubemap Upsampling. As an initial experiment, we consider spherical image resampling without changing the representation format–specifically, performing downsampling followed by upsampling on the same spherical map type. Figure 8 shows a cubemap-formatted polarized environment map downsampled by a factor of four and subsequently upsampled back to its original resolution. We compare two interpolation methods: the original cubemap frame field provided by Yi et al. [2024], and our proposed S2L2 representation.

During the upsampling stage—from a lower-resolution grid back to a higher-resolution grid—pixels that are non-adjacent in the 2D cubemap layout may still be adjacent on the spherical domain.

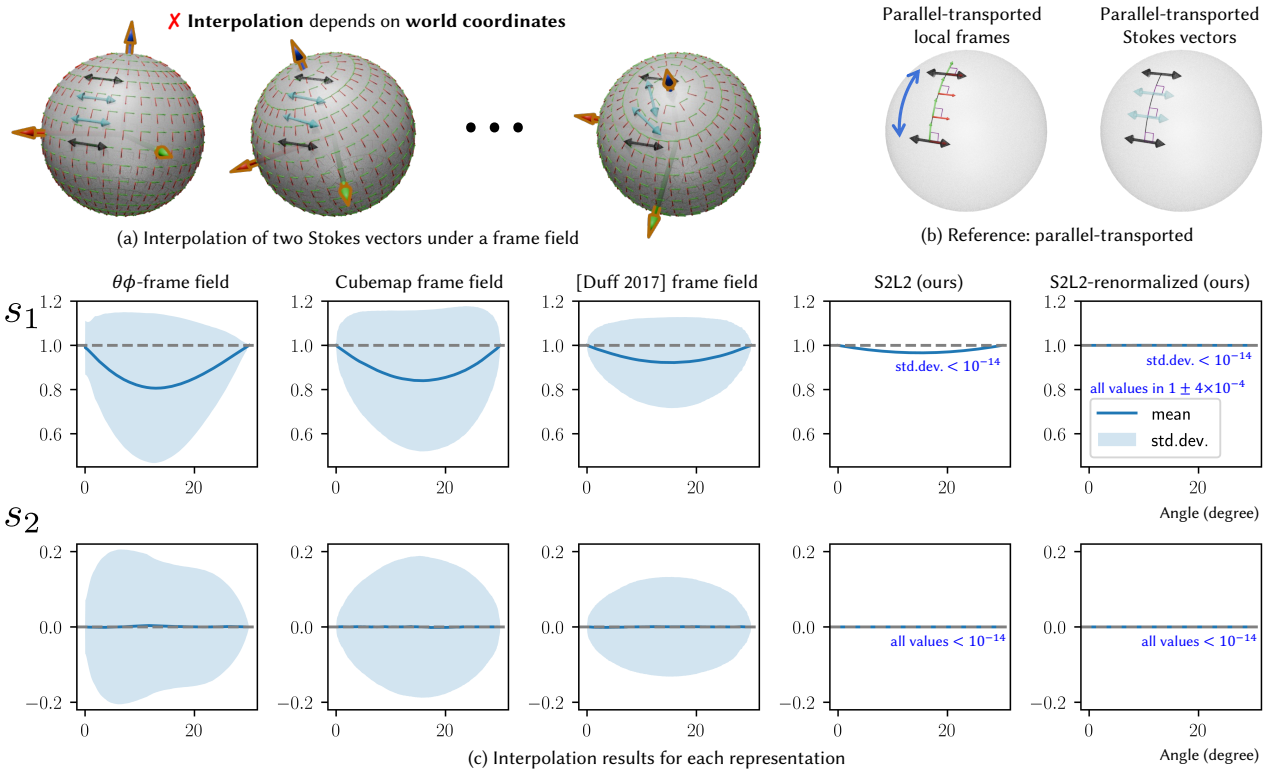

Fig. 6. Comparison of interpolation methods for a pair of Stokes vectors. (a) Conventional frame-dependent interpolation relies on the world coordinate system, resulting in discontinuities due to frame field singularities. (b) Parallel transport-based interpolation serves as a stable geometric reference by maintaining consistent local frames along the interpolation path. While suitable for pairwise interpolation, this approach is not practical for image resampling tasks, as discussed in Section 5.1. (c) Numerical comparison over 20,000 different world coordinate poses. Frame field-based methods exhibit large variations in both s_1 and s_2, depending on the coordinate system. In contrast, our S2L2 interpolation yields consistent results with negligible variance and remains close to the parallel transport reference. The renormalized S2L2 method further improves the match to the parallel transport baseline.

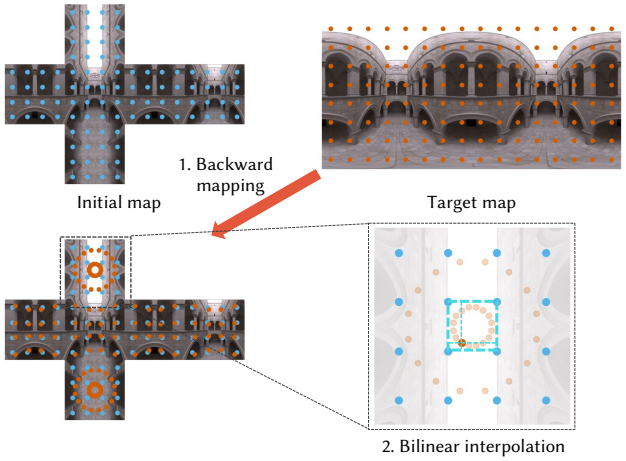

Fig. 7. In Section 5.2, we perform spherical resampling of environment maps by first backward-transforming each target pixel center into the initial map, followed by bilinear interpolation (Equation (24)) of the initial pixel values. For polarized environment maps, the resampling results depend on the representation of Stokes vectors–whether Stokes parameters are defined under different frame fields or expressed via our frame-invariant S2L2 representation.

Consequently, interpolation across such boundaries can encounter frame-field singularities, introducing noticeable artifacts. These artifacts, although not always evident within the cubemap itself, become clearly visible when reprojected into perspective views crossing cubemap boundaries, as illustrated in the second and third columns of Figure 8.

To isolate the effects of downsampling and upsampling, we use the Mitsuba-perspective frame field (Figure 2(b)), which has no singularities within the visible region. Singularities exist globally but are placed outside visible regions to prevent affecting the results. In contrast, interpolation using our S2L2 representation, shown in the fourth and fifth columns of Figure 8, yields consistently smooth and artifact-free outcomes, regardless of resampling steps. Note that artifacts in the cubemap frame field can also appear internally between faces, even in views not explicitly crossing cubemap boundaries, as highlighted in Figure 8(iii).

Conversion from Cubemap to Equirectangular Map. Figure 9 illustrates conversion from cubemap to equirectangular map representations. Highlighted regions (green boxes in Figure 9(b)) demonstrate that our S2L2-based method eliminates interpolation artifacts present in results from the conventional $\theta\phi$-frame field. This figure

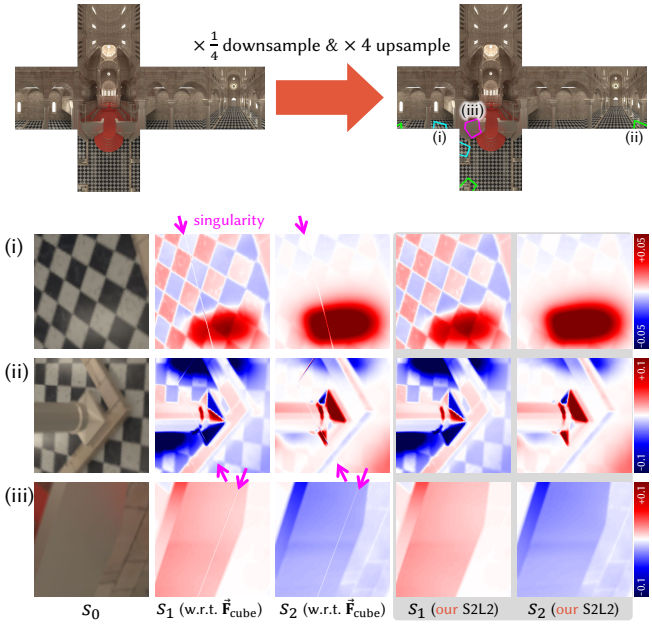

Fig. 8. Effects of spherical resampling (downsampling followed by upsampling) under different Stokes vector representations. The second and third columns visualize the singularities introduced by interpolation using the cubemap frame field, especially when pixels straddle frame field discontinuities. In contrast, interpolation using our S2L2 representation yields smooth, artifact-free results regardless of the resampling process. (i)–(iii) show zoomed-in regions where singularities appear under frame-dependent interpolation but are resolved under our method.

also depicts the experimental setup and provides visual comparisons of resampled Stokes vectors.

A natural question arises: how can we quantitatively evaluate the resulting s_1 and s_2 components in the equirectangular image in Figure 9(b)? As discussed in Section 5.1, defining a ground truth in this context is nontrivial due to the lack of a canonical representation across ray directions. To address this, we propose two approaches for quantitative evaluation.

First, we focus on a specific reference pixel: the center pixel of the upper face in the original cubemap, which corresponds to the top row of the equirectangular image. This pixel has a known Stokes value of $s_1 = (-4.0, -4.0, -5.0)$ and $s_0 = (0, 0, 0)$ across the RGB channels. According to the continuity condition for Stokes vectors under the $\theta\phi$-frame field [Yi et al. 2024], the top row of the equirectangular image should exhibit a smoothly rotating (s_1, s_2) parameters of constant norm, performing two full revolutions along the row.

However, as highlighted by the purple box in Figure 9(b), the resampled result under the $\theta\phi$-frame field incorrectly collapses to zero. In contrast, our S2L2-based result correctly maintains a continuous two-cycle rotation of the vector, with norms nearly identical to the original values: $(3.999, 3.999, 4.999)$ for the RGB channels. Moreover, when using the renormalized interpolation method described in Section 5.1, the resulting values are exactly $(4.000, 4.000, 5.000)$—perfectly matching the original cubemap values.

Second, while defining ground truth for Stokes vector interpolation is inherently nontrivial, establishing ground truth for scalar intensity is straightforward. Based on this, we observe that a linear combination of the s_1 and s_2 components in Figure 9(b) can be interpreted as the result of observing the environment map through a linear polarizer. Specifically, the polarizer is aligned longitudinally with respect to the camera and curves along the image domain, as illustrated on the right of Figure 9(a).

To obtain a ground truth (GT) for this scalar intensity, we re-render the scene using Mitsuba 3 [Jakob et al. 2022] with the same 3D scene configuration of the environment map and the longitudinally aligned polarizer physically modeled in the scene. The rendered high-resolution scalar cubemap is then converted into equirectangular format through spherical resampling to serve as our reference.

Figure 9(c) presents both this ground truth and the scalar intensity results derived from the linear combinations of the s_1 and s_2 components in Figure 9(b). Our method yields the most visually artifact-free result and also achieves the lowest root-mean-square error (RMSE) when compared to the reference.

Another possible approach for spherical resampling that avoids artifacts from frame fields is the frequency-domain method using PSH proposed by Yi et al. [2024]. This method transforms the initial cubemap image into a set of PSH coefficients up to a fixed frequency level, which can then be used to reconstruct the Stokes vectors over the target equirectangular grid. We applied this method to the same experiment shown in Figure 9(c), and compared it against our S2L2 representation in Figure 10. Since frequency-domain approaches approximate the original environment map by only retaining low-frequency components, they are significantly slower and less accurate than ours. Although increasing the frequency level l can reduce the reconstruction error relative to the ground truth, their overall performance remains inferior to that of our method.

Re-converting back to perspective view. Another approach for quantitative evaluation is to reproject the equirectangular images back into the original perspective views contained in the cubemap. This allows for a direct pixel-wise comparison with the original high-resolution cubemap as ground truth. Figure 11 shows such a comparison. Our S2L2-based interpolation yields artifact-free perspective images and achieves lower RMSE compared to the interpolation method based on $\theta\phi$-frame field.

5.3 Image Gradients of Polarized Environment Maps

We analyze gradients of polarized environment maps to visually and quantitatively assess the influence of singularities. Figure 12(a) shows an illustrative example of the gradients computed from an environment map.

When employing the cubemap frame field, the gradients of the s_1 and s_2 components exhibit noticeably higher values at the cubemap boundaries compared to nearby pixels, as highlighted by the red box. These elevated gradient values arise from discontinuities introduced by singularity lines within the cubemap frame field. Conversely, using our S2L2 representation yields gradient values at boundary pixels that are comparable to those of their neighboring pixels, indicating continuity that is invariant to the choice of frame field convention.

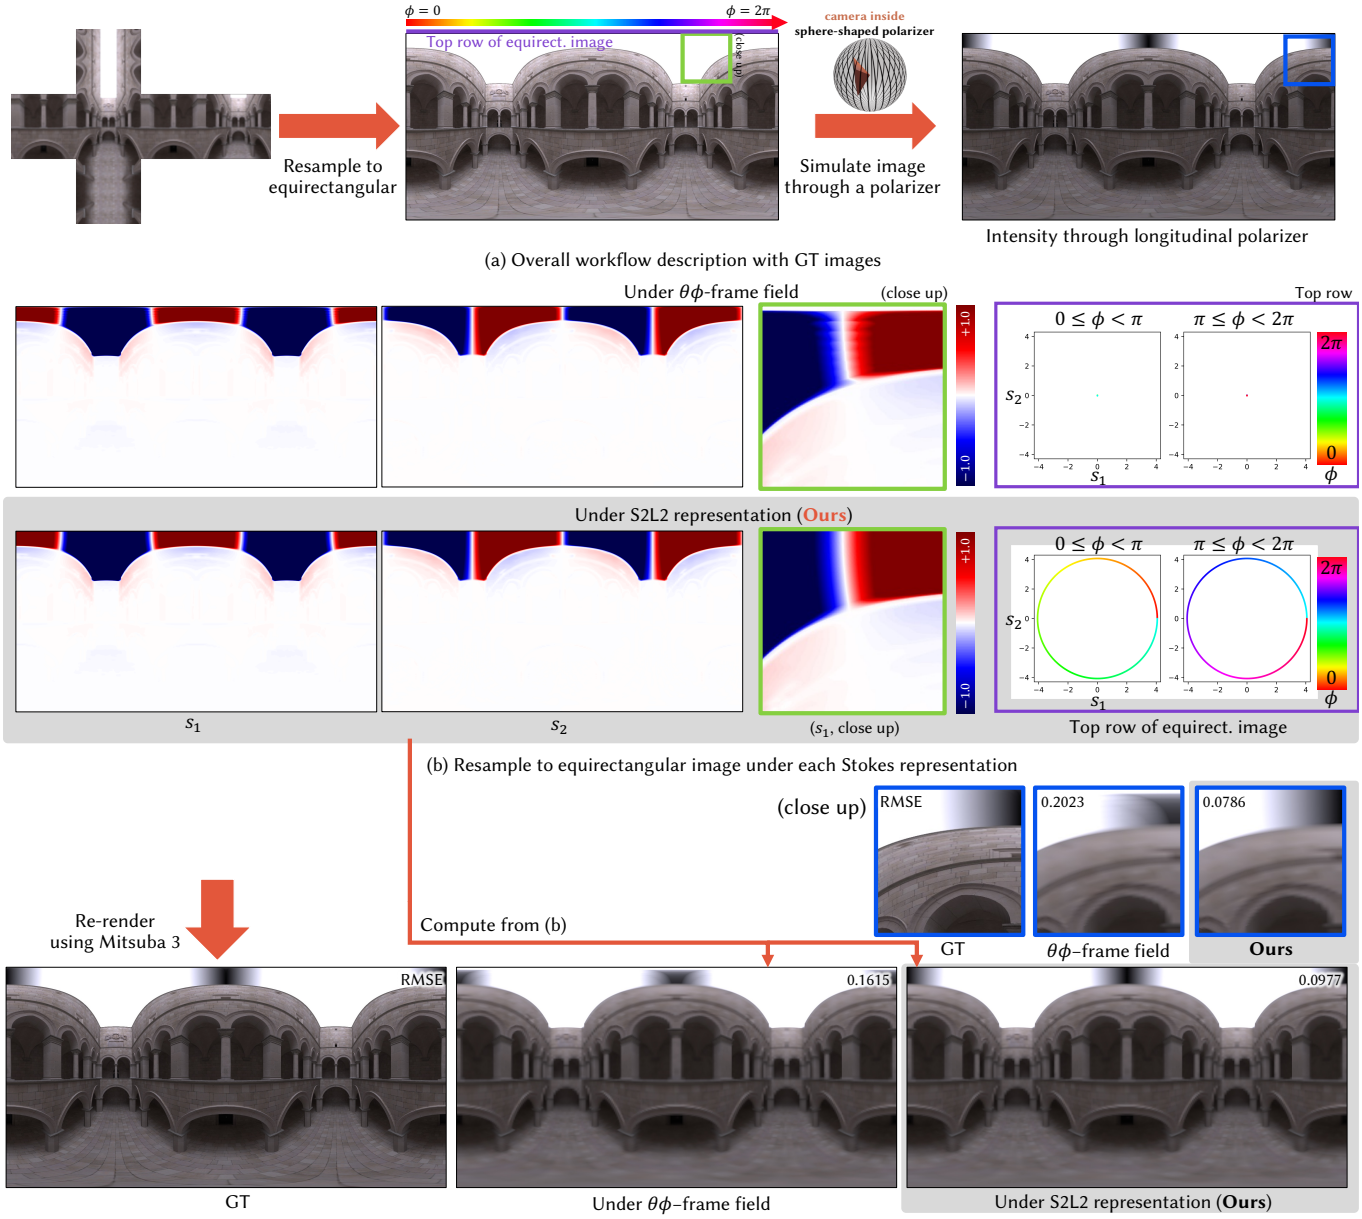

Fig. 9. Our frame-free S2L2 representation for Stokes vectors provides a numerically stable and artifact-free resampling (warping) methods for polarized environment maps. Starting from a cubemap-formatted environment map, we convert it into an equirectangular image using two representations: the conventional $\theta\phi$-frame field and our S2L2 representation. (a) The experimental setup, including a reference scalar intensity image obtained by simulating a longitudinally aligned polarizer. (b) Resampled s_1 and s_2 images show that the $\theta\phi$-frame field introduces artifacts near singularities (green box), while our method preserves continuity and produces artifact-free results. We further validate this by inspecting the top row of the equirectangular image (purple box), where the expected behavior is a continuous two-cycle rotation of the pair of Stokes parameters (s_1, s_2). Our result matches this structure both qualitatively and quantitatively. (c) Scalar intensity images computed by projecting (s_1, s_2) onto the polarizer direction are compared against ground truth rendered with Mitsuba 3. Our method achieves more accurate and artifact-free reconstruction, with lower RMSE.

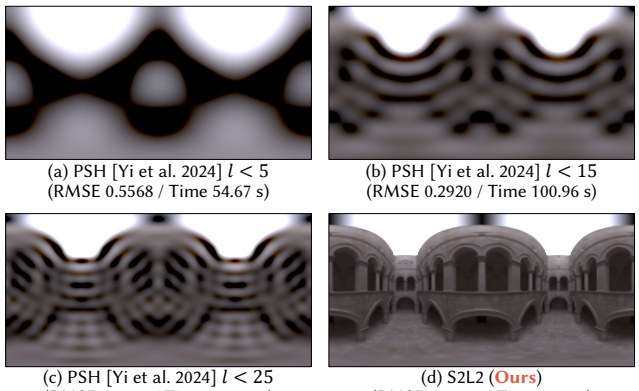

(a) PSH [Yi et al. 2024] $l < 5$ (RMSE 0.5568 / Time 54.67 s)
(b) PSH [Yi et al. 2024] $l < 15$ (RMSE 0.2920 / Time 100.96 s)
(c) PSH [Yi et al. 2024] $l < 25$ (RMSE 0.2162 / Time 230.17 s)
(d) S2L2 (**Ours**) (RMSE **0.0977** / Time **5.87 s**)

Fig. 10. Comparison of spherical resampling from a cubemap to an equirectangular map–under the same setting as Figure 9–between the frequency-domain method of Yi et al. [2024] using PSH and our S2L2 representation. The results show that PSH-based resampling not only incurs substantially higher computation time, but also suffers from greater reconstruction error even with high PSH orders l.

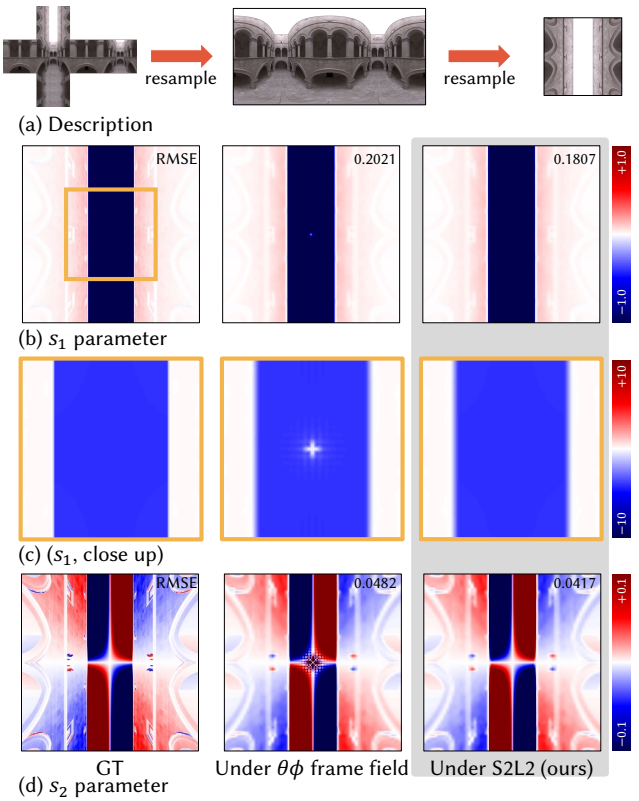

Fig. 11. Quantitative evaluation via re-projection of the equirectangular images obtained in Figure 9. (a) Experimental setup: The original cubemap is converted to equirectangular format and then re-projected back to the original perspective view for comparison. (b)–(d) Comparison of resampled s_1 and s_2 components using the $\theta\phi$-frame field and our S2L2 representation. Our method yields fewer artifacts and lower RMSE.

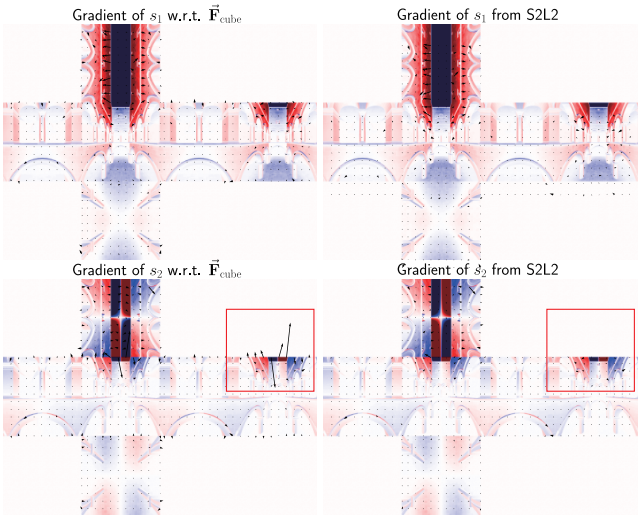

(a) Gradient image of Stokes vector components

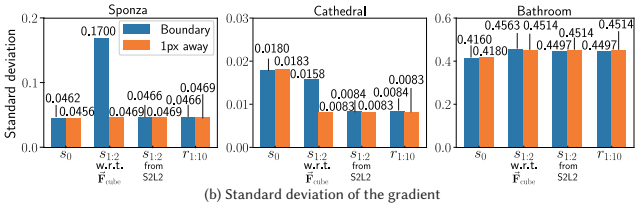

(b) Standard deviation of the gradient

Fig. 12. Analysis of image gradients of polarized environment maps. (a) An example of the gradient of an environment map. The red box highlights high values that deviate from local pixel statistics, caused by discontinuities in the cubemap frame field. (b) Standard deviations of the gradients across pixels. While the results from the cubemap frame field show noticeable variation between boundaries and their adjacent regions, the S2L2-based results remain consistent—similar to the behavior of the s_0 intensity.

We quantitatively evaluate the gradient discontinuities by comparing boundary pixels (where discontinuities are present) with adjacent pixels located one pixel away, which have similar image characteristics but continuous gradient behavior. We calculate the standard deviations of the Frobenius norms for gradients of the s_0 parameter and Jacobian matrices for both the Stokes parameters $s_{1:2}$ and our S2L2 parameters $r_{1:10}$. The quantitative results are shown in Figure 12(b).

For the scalar intensity parameter s_0, the differences between boundary pixels and their adjacent pixels are minimal, as expected for typical image data. However, for the s_1 and s_2 components computed in the cubemap frame field, gradients at discontinuous boundary regions are significantly higher compared to adjacent continuous regions. In contrast, gradients computed directly from our S2L2 representation, as well as those inversely projected from the S2L2 back into Stokes parameters, maintain consistency between boundary pixels and their surroundings. It is important to note that in regions without discontinuities, the Stokes parameters from the cubemap frame field, the inversely projected Stokes parameters from our S2L2 representation, and our S2L2 parameters themselves yield identical standard deviation values for gradient norms.

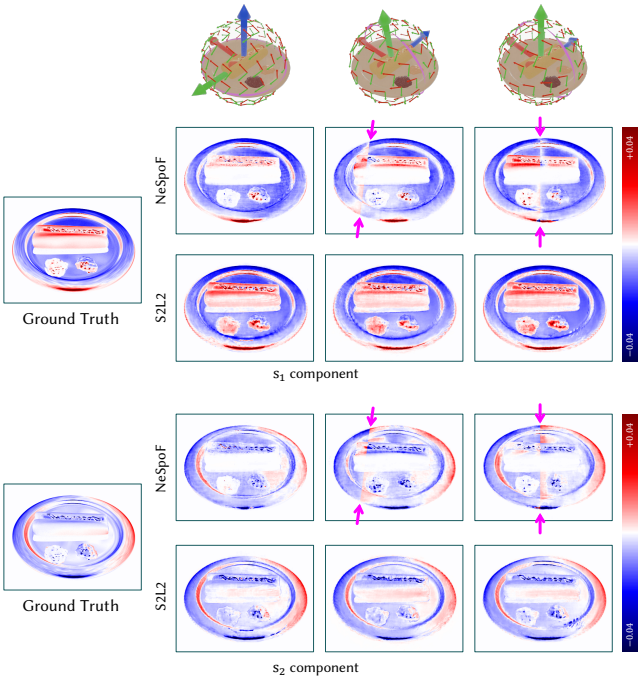

Fig. 13. Rendered images with various choices of world pose. The magenta arrows indicate the singularity.

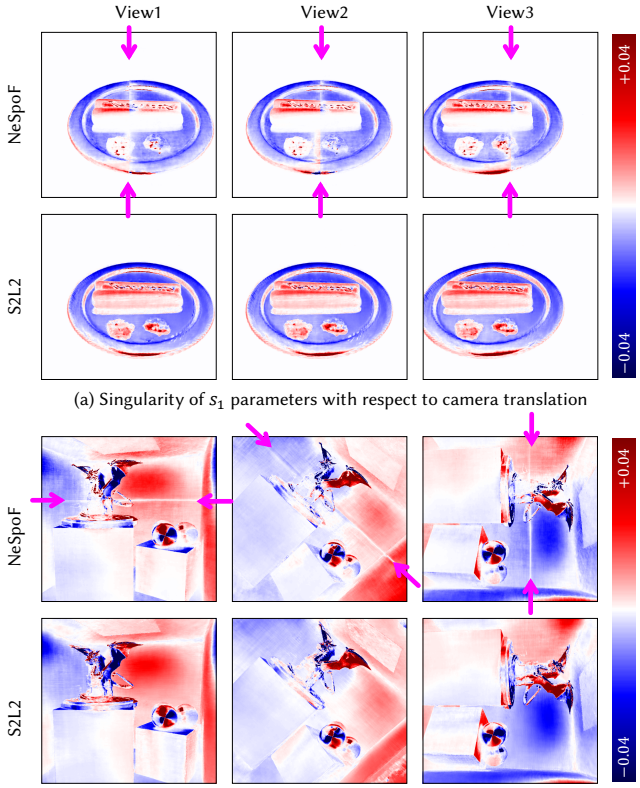

Fig. 14. Rendered images with translation and rotation of the camera pose. The magenta arrows indicate the singularity.

6 APPLICATION: POLARIZED RADIANCE FIELD

Four neural radiance field approaches have been proposed for polarization imaging: PANDORA [Dave et al. 2022], NeISF [Li et al. 2024], NeRSP [Yufei et al. 2024], and NeSpoF [Kim et al. 2023]. Although these methods deal with polarized images, the neural network outputs of the first three frameworks are not directly Stokes vectors. Instead, these methods output diffuse and specular intensities or BRDF components, after which a differentiable renderer computes polarization using a predefined polarization model [Baek et al. 2018] and an assumed refractive index. Consequently, these three methods are unaffected by local ray frames. In contrast, NeSpoF directly predicts polarization components per ray, making its outputs explicitly dependent on local frames.

Specifically, NeSpoF takes as input a 3D position, wavelength, and ray direction, and predicts total intensity s_0, degree of polarization (DoP) ρ, ellipticity χ, and azimuth angle ψ. These predictions are directly converted into Stokes parameters, inherently defined relative to a specific frame. Hence, a predefined frame field aligned to world coordinates is necessary. However, as previously discussed, defining a globally continuous frame field is fundamentally impossible. Given the continuous nature of neural network outputs, unavoidable discontinuities emerge in Stokes components s_1 and s_2 near singular directions. This results in visual artifacts during rendering whenever rays intersect these singular regions.

NeSpoF employs Duff et al. [2017]'s frame field, where singularities occur along the xy-plane. Thus, when the viewing direction aligns with the xy-plane, rendered images exhibit discontinuities. Figure 13 demonstrates how the choice of world coordinates impacts the visibility of singularities. The first column shows a scene aligned with the xy-plane, whereas the third column depicts a scene aligned with the xz-plane. Changing the orientation of the scene alters the camera's viewing direction, thereby affecting the likelihood of capturing singular regions. Specifically, the third column's camera view aligned with the y-axis encounters a singularity line at the image center, whereas the first column, looking downward, does not.

Singularities depend strictly on ray directions; thus, translating the camera does not change their locations. Conversely, principal rotations of the camera shift the pixel positions of rays, causing singularities to rotate correspondingly. Figure 14 illustrates how singular lines remain centered across translated views but rotate with camera principal rotations.

Importantly, such singularities are neither implementation errors nor network design flaws; rather, they represent a fundamental limitation inherent to the frame-dependent Stokes parameter representation. Addressing this issue fundamentally requires a frame-free representation of polarization states.

In contrast to frame-dependent Stokes parameters, our proposed S2L2 representation describes polarization states independently of any local frame. This unique property enables designing neural radiance fields whose outputs can be directly converted to frame-free Stokes parameters on the image plane, eliminating artifacts caused by frame-field discontinuities.

To validate this approach, we modify NeSpoF's output layer dimension from 4 to 12, creating an S2L2-based variant. This 12-dimensional output can be converted directly to camera-frame Stokes parameters without intermediate frame fields. Unlike the original NeSpoF, which converts from world-frame to camera-frame Stokes parameters, our S2L2-based method inherently avoids discontinuities. Figures 13 and 14 illustrate how our S2L2 representation effectively removes singularities arising from frame-field discontinuities, ensuring continuous neural network outputs and artifact-free rendered images.

7 DISCUSSION AND FUTURE WORK

7.1 Alternative Representations and Minimal Dimensionality

A promising avenue for future research is to investigate alternative representations of Stokes vectors that maintain the desirable properties outlined in Section 4.2 (Properties (1) through (7)) but achieve dimensionality lower than the proposed 12-dimensional S2L2 representation. Potential approaches include developing a more compact representation that still fully satisfies these properties or rigorously proving that our 12-dimensional representation is indeed minimal for capturing all desired properties simultaneously. If it turns out that a lower-dimensional representation is fundamentally impossible without sacrificing some properties, it will be valuable to identify which properties can be selectively preserved by reduced-dimensionality representations tailored to specific applications.

7.2 Analogy to Tangent Vector Fields on the Sphere

Though alternative representations with fewer than 12 dimensions might exist, they must still exceed four dimensions to maintain global continuity. Here, an insightful analogy to tangent vector fields on the sphere provides helpful intuition regarding the necessity of a higher-dimensional embedding.

Similar to Stokes vectors, tangent vectors defined on a sphere have two degrees of freedom at each point. Tangent vectors at distinct points reside in different tangent spaces and cannot be directly compared without an appropriate framework. A common practical approach assigns a frame field across the sphere, representing tangent vectors by coefficients with respect to local bases. However, such a scalar field-based representation inherently suffers from discontinuities due to unavoidable singularities in the frame field.

In contrast, a well-known continuous representation exists for tangent vector fields on the sphere by embedding the sphere as a submanifold in $\mathbb{R}^3$. Each tangent vector can then be naturally represented as a vector in this three-dimensional space, thereby achieving continuity at the expense of increasing the dimensionality by one.

This analogy also clarifies the intensity reduction observed in the fourth column of Figure 6. Imagine interpolating two tangent vectors located far apart on the sphere, each aligned along the geodesic connecting their base points. Linear interpolation in $\mathbb{R}^3$ followed by projection back to the tangent space naturally results in shortened vectors due to the discarded component perpendicular to the tangent space. A similar phenomenon explains the intensity decrease observed in our Stokes vector interpolation experiments.

Although this analogy does not constitute a formal mathematical proof, it provides valuable geometric intuition for why achieving continuity in the representation of Stokes vectors may inherently require higher-dimensional embeddings, paralleling established geometric concepts.

7.3 Alternative Approaches for Radiance Fields Application

One can notice that singularity-free frequency-domain representations based on [Yi et al. 2024] are viable for radiance fields, e.g., using Plenoxel-style models, where radiance values across all directions are jointly encoded into a sparse set of (polarized) SH coefficients per spatial location. However, many applications, such as ours in Section 6, rely on NeRF-style radiance fields, which represent angular variation using learned directional encodings rather than a fixed (polarized) SH basis. The global nature of [Yi et al. 2024] makes it incompatible with such models. Our S2L2 is designed for per-ray use and supports frame-invariant encoding and interpolation across ray directions without imposing a strict angular frequency limit like Plenoxel-style methods. We argue that enabling polarization in neural field methods should not be limited to SH-compatible (Plenoxel-style) approaches, especially when modern rendering pipelines favor flexible, local representations.

7.4 Mueller Matrix Representations

Although this paper focuses solely on the representation of the Stokes vector corresponding to light intensity, we believe a frame-free representation of the Mueller matrix–which describes linear transformations of Stokes vectors and characterizes light-matter interactions–is also possible. By applying Equations (13) and (14) respectively to the row and column indices of a 4×4 Mueller matrix, one may obtain a 12×12 frame-free representation. However, detailed proofs, implementation, experiments, and applications are left as future work.

7.5 Future Applications

We expect the S2L2 representation to benefit broader computer vision tasks beyond image processing and radiance field methods. For example, its frame-invariant structure makes it suitable for feature extraction, learning-based methods, and polarization-based scene understanding.

8 CONCLUSION

We have introduced and thoroughly analyzed the S2L2 representation, a novel, frame-free representation for Stokes vectors. This representation effectively resolves the inherent singularities and discontinuities associated with conventional frame-dependent representations of polarization. By employing spin-2 spherical harmonics, the S2L2 method maintains continuity and rotation invariance, offering significant numerical stability improvements. Our extensive numerical experiments demonstrate the robustness of the S2L2 representation across various interpolation and spherical resampling tasks, highlighting its effectiveness in eliminating artifacts

and maintaining consistent behavior independent of world coordinate system choices. Furthermore, we applied our method to a neural polarized radiance field framework, clearly demonstrating that our S2L2-based approach successfully eliminates artifacts caused by singularities in conventional frame fields, providing continuous and artifact-free rendered images. We anticipate that the proposed S2L2 representation sets a new standard for polarization-aware graphics and imaging tasks, especially in applications where the directional variation of light intensity is important, by increasing robustness and visual quality while balancing accuracy and computational efficiency efficiency.

ACKNOWLEDGMENTS

Min H. Kim acknowledges the Samsung Research Funding & Incubation Center (SRFC-IT2402-02), the Korea NRF grant (RS-2024-00357548), the MSIT/IITP of Korea (RS-2022-00155620, RS-2024-00398830, 2022-0-00058, and 2017-0-00072), Microsoft Research Asia, and Samsung Electronics.

REFERENCES

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.

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.

Tom Duff, James Burgess, Per Christensen, Christophe Hery, Andrew Kensler, Max Liani, and Ryusuke Villemin. 2017. Building an orthonormal basis, revisited. *JCGT* 6, 1 (2017).

RDM Garcia and CE Siewert. 1986. A generalized spherical harmonics solution for radiative transfer models that include polarization effects. *Journal of Quantitative 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.

Joshua N Goldberg, Alan J MacFarlane, Ezra T Newman, Fritz Rohrlich, and EC George Sudarshan. 1967. Spin-s Spherical Harmonics and ð. *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.

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.

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.

Chenhao Li, Taishi Ono, Takeshi Uemori, Hajime Mihara, Alexander Gatto, Hajime Nagahara, and Yusuke Moriuchi. 2024. NeISF: Neural Incident Stokes Field for Geometry and Material Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 21434–21445.

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.

Michal Mojzík, Tomáš Skřivan, Alexander Wilkie, and Jaroslav Křivánek. 2016. Bidirectional polarised light transport. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. 97–108.

Ezra T Newman and Roger Penrose. 1966. Note on the bondi-metzner-sachs group. *J. Math. Phys.* 7, 5 (1966), 863–870.

Henri Poincaré. 1885. Sur les courbes définies par les équations différentielles. *Journal de mathématiques pures et appliquées* 1 (1885), 167–244.

Ravi Ramamoorthi and Pat Hanrahan. 2002. Frequency space environment map rendering. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. 517–526.

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.

Shlomi Steinberg and Ling-Qi Yan. 2021. A generic framework for physical light transport. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–20.

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.

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 and Andrea Weidlich. 2012. Polarised light in computer graphics. In *SIGGRAPH Asia 2012 Courses*. 1–87.

Shinyoung Yi, Donggun Kim, Jiwoong Na, Xin Tong, and Min H Kim. 2024. Spin-Weighted Spherical Harmonics for Polarized Light Transport. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–24.

Han Yufei, Guo Heng, Fukai Koki, Santo Hiroaki, Shi Boxin, Okura Fumio, Ma Zhanyu, and Jia Yunpeng. 2024. NeRSP: Neural 3D Reconstruction for Reflective Objects with Sparse Polarized Images.

Matias Zaldarriaga and Uroš Seljak. 1997. All-sky analysis of polarization in the microwave background. *Physical Review D* 55, 4 (1997), 1830.