

Polarimetric Helmholtz Stereopsis

Yuqi Ding¹ Yu Ji² Mingyuan Zhou² Sing Bing Kang³ Jinwei Ye¹
¹Louisiana State University ²DGene U.S.A. ³Zillow Group

Abstract

Helmholtz stereopsis (HS) exploits the reciprocity principle of light propagation (i.e., the Helmholtz reciprocity) for 3D reconstruction of surfaces with arbitrary reflectance. In this paper, we present the polarimetric Helmholtz stereopsis (polar-HS), which extends the classical HS by considering the polarization state of light in the reciprocal paths. With the additional phase information from polarization, polar-HS requires only one reciprocal image pair. We formulate new reciprocity and diffuse/specular polarimetric constraints to recover surface depths and normals using an optimization framework. Using a hardware prototype, we show that our approach produces high-quality 3D reconstruction for different types of surfaces, ranging from diffuse to highly specular.

1. Introduction

Reconstructing 3D surfaces from 2D images is a long-standing ill-posed problem in computer vision. The complex surface reflectance properties of real-world objects make the problem highly challenging. All existing methods are limited to certain types of surface reflectance. For example, passive techniques examine the optical appearance of a surface under non-tightly focused illumination (e.g., a distant light source) and often assume photo-consistency or Lambertian reflectance model for 3D reconstruction.

Helmholtz Stereopsis (HS) [55] is a 3D reconstruction technique that can recover surfaces with arbitrary and unknown reflectance. HS exploits the symmetry of surface reflectance; this is accomplished by using reciprocal image pairs (minimal three) that are captured with exchanged camera and light source positions. The reciprocity property guarantees that the relationship between the intensities at corresponding pixels depends only on the surface shape, and is independent of surface reflectance.

In this paper, we present a novel method we call polarimetric Helmholtz Stereopsis (polar-HS), which extends the classical HS by considering the polarization state of light in the reciprocal paths. We investigate the reciprocity relationship when the polarization states of incident and outgoing

light are unrestricted (in which case the original Helmholtz reciprocity property cannot be directly applied). We derive a transpositional reciprocity relationship based on the Stokes-Mueller formalism, and formulate a reciprocity constraint for depth and normal estimation. We also exploit polarimetric cues under different types of reflections. We propose a new polarimetric image decomposition method that allows us to apply the polarimetric constraints under different circumstances. By combining the reciprocity and polarimetric constraints, our method can recover the surface depth and normal with only one reciprocity pair, which greatly simplifies the capture process. Unlike other one-pair HS methods [56, 48] that assume continuous parametric depth functions, polar-HS works for discontinuous depth and does not require priors on surface geometry and material properties.

We validate our method with both synthetic and real experiments. We build a real polar-HS acquisition system with a rotating wheel to allow exchange of camera and light source positions. We perform experiments on objects with various shapes and reflectances, and on different composite scenes. Results show that our method is state-of-the-art.

Contributions. Our key contribution is the derivation of the transpositional reciprocity relationship when unrestricted polarization states are being considered. The original Helmholtz reciprocity property has strong restrictions on the polarization states of the light beams in the reciprocal paths, which is hard to satisfy in practice. The transpositional reciprocity we derive allows the classical HS to be extended to the polarimetric case.

We propose a new image decomposition formulation that consists of three components: polarized-specular, polarized-diffuse, and unpolarized-diffuse. The decomposition provides a more accurate estimation of the angle of polarization for regularizing the surface normal, because it separates diffuse and specular (the angles of polarization under these two cases have a 90° shift). Shape-from-polarization methods usually use the overall angle, and assume a dominant type of reflection. This decomposition also provides a specular map that allows us to apply the polarimetric cues under different types of reflection.

Polar-HS reduces the minimal number of image pairs to

only one, without imposing any surface prior. Compared to other shape-from-polarization methods, polar-HS does not have the problem of angular ambiguity and does not require the refractive index of the surface to be known.

2. Related Work

We first briefly review physics-based methods for 3D shape recovery before focusing on two specific classes of methods that are most relevant to our work: reciprocity-based and polarization-based. Table 1 summarizes a comparison of our method (polar-HS) with classical methods.

Physics-based Shape Recovery. Physics-based techniques examine the optical appearance of a surface under certain illumination mode and often assume photo-consistency or Lambertian reflectance model for 3D reconstruction. We can categorize the techniques as passive or active based on the illumination mode. Passive methods use unknown and non-tightly focused illumination (*e.g.*, a distant light source). Notable examples include multi-view stereo [41, 42, 46] and structure-from-motion [49, 1, 13]. As passive methods heavily rely on the object’s intrinsic appearance for feature matching, they are ineffective on textureless surfaces. Active techniques use known and controlled illumination as a probe; examples include photometric stereo [53, 9, 23], time-of-flight [15, 22, 35], and structured light [30, 20, 19, 21]. These methods can produce dense 3D reconstruction, but are usually sensitive to view-dependent specularities and the inter-reflection caused by concave surfaces. All these methods have limitations due to the complex reflectance of real-world surfaces.

Helmholtz Stereopsis (HS). The method is first introduced by Zickler *et al.* [55]. It is an active approach that is capable of recovering surfaces with arbitrary reflectance. Much progress has subsequently been made to improve the original HS. Tu and Mendonça [48] solve HS with a single pair by assuming a piece-wise linear curve constraint. Zickler *et al.* [56] formulate a PDE constraint by assuming C^1 continuity in depth, so as to perform HS under a binocular setting. Jankó *et al.* [25] introduce a general radiometric calibration method for HS. Delaunoy *et al.* [17] extend HS to full-body scanning by using variational approach to optimize over the entire surface. Weinmann *et al.* [51] combine HS with a structured light technique to improve the reconstruction accuracy. Mori *et al.* [33] introduce an integration-based Helmholtz condition which reduces the noise sensitivity of HS. Roubtsova and Guillemaut [39, 40] derive a Bayesian framework for HS optimization, and use color multiplexing to simultaneously capture the reciprocal pair in order to handle dynamic scenes. Our method extends HS by expanding the reciprocity constraint to polarimetric reflectance and incorporating the polarimetric cues for more accurate 3D reconstruction.

Method	Min # Inputs	Surface Assumption	Accuracy
MVS	2	Lambertian	Moderate
PS	3	Lambertian	High
SL	> 10	Arbitrary	High
SfP	3	Dielectric	Low
HS	6	Arbitrary	Moderate
polar-HS	4	Arbitrary	High

Table 1. A comparison of polarimetric Helmholtz Stereopsis (polar-HS) with classical 3D reconstruction methods. Note: MVS - multi-view stereo; PS - photometric stereo; SL - structured light; SfP - shape-from-polarization; HS - Helmholtz stereopsis.

Shape-from-Polarization (SfP). This class of methods model the surface normal using the degree or angle of polarization. The surface’s refractive index is usually assumed to be known. Miyazaki *et al.* [29] and Atkinson and Hancock [4] leverage the diffuse polarization for shape estimation. Rahmann and Canterakis [38] propose a specular polarization model and apply it on reflective surfaces. SfP methods usually suffer from the azimuth angle ambiguity which may cause the normal estimation being flipped. To resolve this ambiguity, additional shape priors or visual cues (such as convexity prior [29, 31], boundary normal prior [4], shading cues [5], photometric cues [18, 37], and multi-spectral measurements [24]) are combined with the polarization model. Smith *et al.* [44, 45] use SfP to solve for surface height to mitigate the angular ambiguity. Many works use SfP to recover fine surface details given a coarse shape estimated from another technique, such as multi-view stereo [3, 12, 54], photometric stereo [2, 47], space carving [28], structure-from-motion [16], or RGB-D sensors [26, 27]. Beak *et al.* [7] jointly estimate the polarimetric reflectance and the surface geometry. Ba *et al.* [6] propose a data-driven approach that estimates the surface shape from polarimetric images with a deep neural network.

3. Helmholtz Stereopsis

Helmholtz stereopsis (HS) [55] works by exploiting the symmetry of surface reflectance. It uses several reciprocal image pairs with exchanged camera and light source positions to estimate surface normal and depth. Let $O_a \in \mathbb{R}^3$ and $O_b \in \mathbb{R}^3$ be two 3D positions. A reciprocal image pair $\mathcal{I} = \{I_a, I_b\}$ is captured by swapping the camera and light source at O_a and O_b (*i.e.*, I_a is captured with the light source at O_a and the camera at O_b ; I_b is captured with the light source at O_b and the camera at O_a). Given a point on the object surface, the goal is to estimate its 3D position P and normal vector $\mathbf{n}$.

Let $f(\mathbf{i}, \mathbf{o})$ be the bidirectional reflectance distribution function (BRDF) of the surface point. f is calculated as the

ratio of the outgoing radiance (along the direction $\mathbf{o}$) and the incident irradiance (along the direction $\mathbf{i}$). The Helmholtz reciprocity indicates that f is symmetric about the incident and outgoing directions, i.e., $f(\mathbf{i}, \mathbf{o}) = f(\mathbf{o}, \mathbf{i})$. Let $\mathbf{v}_a = (O_a - P) / \|O_a - P\|$ and $\mathbf{v}_b = (O_b - P) / \|O_b - P\|$ be two unit directions from P to O_a and from P to O_b . The two intensity images in the reciprocal pair can be formulated as

$$\begin{aligned} I_a &= f(\mathbf{v}_a, \mathbf{v}_b) E \rho_a(\mathbf{v}_a \cdot \mathbf{n}), \\ I_b &= f(\mathbf{v}_b, \mathbf{v}_a) E \rho_b(\mathbf{v}_b \cdot \mathbf{n}), \end{aligned} \quad (1)$$

where E is the light source intensity, $\rho_a = 1 / \|O_a - P\|^2$ and $\rho_b = 1 / \|O_b - P\|^2$ are distance attenuation factors, and $\mathbf{v}_a \cdot \mathbf{n} = \cos \theta_a$ and $\mathbf{v}_b \cdot \mathbf{n} = \cos \theta_b$ are angular fall-off factors.

By dividing the above two equations, we eliminate the light source intensity E and the surface BRDF (noting that $f(\mathbf{v}_a, \mathbf{v}_b) = f(\mathbf{v}_b, \mathbf{v}_a)$). We thus obtain the following reciprocal constraint that regularizes depth and normal with respect to the reciprocal image pair:

$$(I_a \rho_b \mathbf{v}_b^\top - I_b \rho_a \mathbf{v}_a^\top) \mathbf{n} = 0. \quad (2)$$

Given a pre-calibrated camera and light source positions (O_a and O_b), the surface position P and normal $\mathbf{n}$ can be solved with at least three reciprocal pairs. This is because given an estimated P , we need at least three equations to uniquely solve for $\mathbf{n} \in \mathbb{R}^3$.

4. Polarimetric Helmholtz Stereopsis

Our polarimetric Helmholtz stereopsis (polar-HS) uses a linearly polarized light source and a polarization camera to acquire the reciprocal image pair. Our images therefore embed the polarization state of light. Note that the original Helmholtz Reciprocity Principle restricts its applicability to corresponding polarization states for incident and outgoing light [14] (i.e., when the light path is reversed, the polarization states of the two light beams should also be interchanged). It cannot be directly applied to unrestricted representations of the polarization states. We derive the reciprocity relationship under the unrestricted cases, and use it as a constraint for surface reconstruction (Section 4.1).

We introduce polarimetric constraints that are dependent on the surface reflection types (i.e., specular polarized, diffuse polarized, and diffuse unpolarized), and propose a new image separation method to enable the usage of the polarimetric constraints (Section 4.2). An optimization framework is used to jointly estimate the surface normal and depth by combining the reciprocity and polarimetric constraints (Section 4.3).

4.1. Polarimetric Reciprocity

We use the Stokes vector to describe the polarization states. A Stokes vector has four components: $S =$




Figure 1. Configuration of the polarimetric Helmholtz stereopsis.

$[S(0), S(1), S(2), S(3)]^\top$, where $S(0)$ specifies the radiant intensity of light (equivalent to the intensity image), $S(1)$ specifies the preference of horizontal to vertical linear polarization, $S(2)$ specifies the preference of 45° to 135° linear polarization, and $S(3)$ specifies the preference of right to left circular polarization. Additional constraints on Stokes vector values are: 1) $S(0) \in \mathbb{R}^+$, 2) $S(1), S(2), S(3) \in [-S(0), S(0)]$, and 3) $S(0)^2 \geq S(1)^2 + S(2)^2 + S(3)^2$. Note that a Stokes vector is relative to the selection of reference axes (i.e., a two-dimensional orthogonal basis on the wave plane that is perpendicular to the light's propagation direction). Here we assume linearly polarized light.

We use the same configuration as the standard HS, except that the light source is linearly polarized with Stokes vector S_l and the camera is polarization-sensitive so that it can measure the Stokes vectors of the light received. Fig. 1 illustrates our configuration. Given a reciprocal Stokes vector pair $\mathcal{S} = \{S_a, S_b\}$, S_a is measured with the light source at O_a and the camera at O_b , and S_b with swapped light source and camera positions. Similar to the standard HS, we can formulate the Stokes vectors measured by the camera (S_a and S_b) with the polarimetric surface reflectance along with the distant and angular fall-off factors as

$$\begin{aligned} S_a &= M(\mathbf{v}_a, \mathbf{v}_b) S_l \rho_a(\mathbf{v}_a \cdot \mathbf{n}), \\ S_b &= M(\mathbf{v}_b, \mathbf{v}_a) S_l \rho_b(\mathbf{v}_b \cdot \mathbf{n}). \end{aligned} \quad (3)$$

The 4×4 matrix M (Mueller matrix) represents the polarimetric surface reflectance that describes how the Stokes vector is changed after reflection.

Now the question is: *does M still follows the same reciprocal relationship as the BRDF (i.e., $M(\mathbf{v}_a, \mathbf{v}_b) = M(\mathbf{v}_b, \mathbf{v}_a)$) when the Stokes vectors of the incident and reflected light are represented under arbitrary (or unrestricted) reference axes?* The original Helmholtz Reciprocity Principle [50] is as follows:

Theorem 1 (Helmholtz Reciprocity). *Suppose a certain amount of light J leaving the point A in a given direction is polarized in a , and that of this light, the amount K arrives at point B polarized in b . Then, when the light returns over the same path, and the quantity of light J polarized in b proceeds from the point B , the amount of this light that arrives at point A polarized in a will be equal to K .*




Figure 2. We rotate the local reference axes to a global reference so that Theorem 1 can be applied to describe the reciprocity relationship between the two paths.

Theorem 1 has restricted applicability on corresponding polarization states for incident and outgoing light (*i.e.*, when the light path is reversed, the polarization states of light should also be interchanged). In reality, it is hard to acquire the Stokes measurement under this restricted circumstance. For example, both the camera and light source have their own Stokes representation under local reference axes. It is impractical to exchange their Stokes reference axes when their positions are swapped. We therefore derive a new reciprocity relationship when the representation of polarization state is unrestricted (*i.e.*, the polarization states of incident and outgoing light are not interchangeable due to different reference axes). Our reciprocity relationship is stated as follows:

Lemma 1.1. *When the two light beams in a reversible path are represented by Stokes vectors with reference axes that are associated with the light beams, the reciprocity relationship can be expressed as a transposition of the Mueller matrix that correlates the two Stokes vectors.*

Proof. Consider two light beams with Stokes vectors S_i and S_o that propagate along directions $\mathbf{i}$ and $\mathbf{o}$, respectively. Both S_i and S_o are represented in their local reference axes $\mathbf{b}_i = \{\mathbf{x}_i, \mathbf{y}_i\}$ and $\mathbf{b}_o = \{\mathbf{x}_o, \mathbf{y}_o\}$ (where $\{\mathbf{x}_i, \mathbf{y}_i\}$ and $\{\mathbf{x}_o, \mathbf{y}_o\}$ are two pairs of orthogonal axes on their respective wave plane). Let M be the Mueller matrix that correlates S_i and S_o . We can write the following equations:

$$S_o = M(\mathbf{i}, \mathbf{o})S_i, \quad S_i = M(\mathbf{o}, \mathbf{i})S_o. \quad (4)$$

As $\mathbf{b}_i$ and $\mathbf{b}_o$ are associated with the light beams and are switched when the light path is reversed, Theorem 1 cannot be directly applied to describe the reciprocity relationship between $M(\mathbf{i}, \mathbf{o})$ and $M(\mathbf{o}, \mathbf{i})$.

In order to apply Theorem 1, we define a global reference basis $\mathbf{b}_g = \{\mathbf{x}_g, \mathbf{y}_g\}$ and transform the two Stokes vectors from their local references to the global reference (see Fig. 2). The global reference axes are defined as: $\mathbf{y}_g = \mathbf{i} \times \mathbf{o}$,

$\mathbf{x}_g = \mathbf{y}_g \times \mathbf{i}$ for path $\mathbf{i}$, and $\mathbf{x}_g = \mathbf{y}_g \times \mathbf{o}$ for path $\mathbf{o}$. The global reference axes satisfy the condition of Theorem 1 as it is associated with the paths instead of the light beams.

By multiplying the Mueller matrices that rotate the local references to the global one, we obtain two new Stokes vectors S'_i and S'_o that are represented in the global reference axes: $S'_i = M_r(\Phi_i)S_i$ and $S'_o = M_r(\Phi_o)S_o$. Φ_i and Φ_o are the angles spanned by $\mathbf{y}_i$ and $\mathbf{y}_g$, and $\mathbf{y}_o$ and $\mathbf{y}_g$, respectively. Let M' be the Mueller matrix that correlates S'_i and S'_o . Thus M' satisfies the reciprocity relationship described in Theorem 1: $M'(\mathbf{i}, \mathbf{o}) = M'(\mathbf{o}, \mathbf{i}) = M_f$, leading to

$$S'_o = M_f S'_i, \quad S'_i = M_f S'_o. \quad (5)$$

By substituting S'_i and S'_o with S_i and S_o (respectively) in Eq. 5, we have

$$S_o = M_r^{-1}(\Phi_o)M_f M_r(\Phi_i)S_i, \quad (6)$$

$$S_i = M_r^{-1}(\Phi_i)M_f M_r(\Phi_o)S_o.$$

Using the definitions for S_o and S_i in Eq. 4, Eq. 6 becomes

$$M(\mathbf{i}, \mathbf{o}) = M_r^{-1}(\Phi_o)M_f M_r(\Phi_i), \quad (7)$$

$$M(\mathbf{o}, \mathbf{i}) = M_r^{-1}(\Phi_i)M_f M_r(\Phi_o).$$

Since the rotational Mueller matrices are orthonormal (*i.e.*, $M_r^{-1} = M_r^T$) and M_f is diagonally symmetric (*i.e.*, $M_f = M_f^T$) [11], we can derive the following reciprocity relationship in the form of transposition:

$$M^T(\mathbf{i}, \mathbf{o}) = (M_r^{-1}(\Phi_o)M_f M_r(\Phi_i))^T$$

$$= M_r^{-1}(\Phi_i)M_f M_r(\Phi_o)$$

$$= M(\mathbf{o}, \mathbf{i}). \quad (8)$$

□

Sekera [43] derives a transpositional reciprocity relationship similar to Lemma 1.1 in the scattering processes.

Reciprocity Constraint. As the Stokes vectors of the camera and light source are observed in their local reference axes, the surface reflectance Mueller matrix M follows the transpositional reciprocity according to Lemma 1.1: $M^T(\mathbf{v}_a, \mathbf{v}_b) = M(\mathbf{v}_b, \mathbf{v}_a)$. By substituting this reciprocity relationship into Eq. 3 and eliminating M , we obtain the following reciprocity constraint:

$$(S_a \otimes S_l^g \rho_b \mathbf{v}_b^T - S_l^g \otimes S_b \rho_a \mathbf{v}_a^T) \mathbf{n} = 0. \quad (9)$$

where S_l^g is the transpose of the pseudoinverse of S_l ; $\otimes$ is the Kronecker product.

We use Eq. 9 to estimate surface depth ($\rho_{a,b}$ and $\mathbf{v}_{a,b}$ are derivable from depth) and normal $\mathbf{n}$ in an iterative way when given reciprocal Stokes vector pair $\mathcal{S} = \{S_a, S_b\}$ and pre-calibrated light source Stokes vector S_l . More details




Figure 3. The proposed decomposition. We first convert the images captured by a polarization camera to the Stokes vectors. We then decompose the Stokes vectors into three components.

on the optimization algorithm can be found in Section 4.3. Given a depth estimation, we need at least three equations to uniquely solve for the $\mathbf{n} \in \mathbb{R}^3$.

The standard HS generates the three equations from three reciprocal pairs. In our polarimetric case, we obtain two equations for one reciprocal pair by changing the polarization state of light source, because our reciprocity constraint (Eq. 9) is dependent on the light source (while the standard HS is not). The maximum number of independent equations we can obtain for one pair is two. This is because two dot products (*i.e.*, $\mathbf{v}_a \cdot \mathbf{n}$ and $\mathbf{v}_b \cdot \mathbf{n}$) reduce the terms that involve the normal to scalars.

In principle, if we use only the reciprocity constraint, we would need at least two reciprocal pairs (*i.e.*, 4 equations) to solve for the depth and normal. However, because the observed polarization states are directly related to the surface geometry and reflectance type, we use the polarimetric constraints (Section 4.2) to formulate one additional independent equation on normal. This means that we can estimate depth and normal *using only one reciprocal pair*.

4.2. Polarimetric Cues

We consider the polarimetric constraints under different types of reflected light. We decompose the measured Stokes vectors into three components of different reflection and polarization characteristics (*i.e.*, specular vs. diffuse, polarized vs. unpolarized), and then derive a specific polarimetric constraint for each type of reflection.

Polarimetric Image Decomposition. The problem has been studied in polarization-based specular removal [52, 36, 34]. An image is often decomposed into a specular component and a diffuse component, depending on the polariza-

tion status. It is commonly assumed that the specular component is polarized while diffuse is unpolarized. However, as shown in [4], diffuse reflection also exhibits useful polarimetric characteristics that regularize the surface normal. Here we propose a new decomposition formulation that separates the observed Stokes vectors (S) into three components: specular-polarized (S_{sp}), diffuse-polarized (S_{dp}), and diffuse-unpolarized (S_{du}):

$$S = S_{sp} + S_{dp} + S_{du}. \quad (10)$$

Suppose we have two linearly polarized light sources with the same intensity but perpendicular angles of polarization. Without loss of generality, we assume their angles of polarization are 0° and 90° . Their Stokes vectors are S_l^0 and S_l^{90} , respectively. Note that these light sources can also provide us the two reciprocity constraints. Let S^0 be the Stokes vector reflected from a surface point and observed by the camera when light source is S_l^0 , and S^{90} is observed under S_l^{90} (see Section 4.1 for composition and properties of the Stokes vector). Here we assume S^0 and S^{90} are linearly polarized (*i.e.*, $S^0(3) = S^{90}(3) = 0$).

We now show how S^0 can be decomposed; S^{90} can be similarly decomposed. For notation simplicity, we drop the superscript for degree in the decomposition components (S_{sp} , S_{dp} , and S_{du}). Since the light sources' angles of polarization are crossed by 90° , the polarization parameters in S_l^0 and S_l^{90} have the following relationship:

$$S_l^0(1) + S_l^{90}(1) = 0, \quad S_l^0(2) + S_l^{90}(2) = 0. \quad (11)$$

Since the specular reflection is always fully polarized and its angle of polarization is the same as that for the light source, we use this relationship to cancel out the polarization parameters in S_{sp} by adding S^0 and S^{90} . Since S_{du} is unpolarized, it has only the intensity parameter: $S_{du} = [S_{du}(0), 0, 0, 0]^T$. The polarization parameters in $S^0 + S^{90}$ are then solely related to S_{dp} , yielding

$$S_{dp}(1) = \frac{S^0(1) + S^{90}(1)}{2}, \quad S_{dp}(2) = \frac{S^0(2) + S^{90}(2)}{2}. \quad (12)$$

The polarization parameters in S_{sp} are computed as

$$S_{sp}(1) = S^0(1) - S_{dp}(1), \quad S_{sp}(2) = S^0(2) - S_{dp}(2). \quad (13)$$

We calculate $S(0)$ for both S_{dp} and S_{sp} using the Stokes vector constraint $S(0)^2 = S(1)^2 + S(2)^2 + S(3)^2$, as the two components are fully polarized. Finally, we compute the intensity of S_{du} as

$$S_{du}(0) = S^0(0) - S_{sp}(0) - S_{dp}(0). \quad (14)$$

Using Eqs. 12-14, we can decompose a reflected Stokes vector into three components. Fig. 3 shows an example of our decomposition.

Polarimetric Constraints. Both the specular and diffuse polarized reflections can regularize the surface normal. According to Fresnel’s equations, the specular reflection is dominated by s-polarized light, whose angle of polarization is perpendicular to the reflection plane (*i.e.*, the plane formed by the surface normal and the reflected light). This happens when the incident light is not oblique to the local surface. In the diffuse reflection case, the angle of polarization has a 90° phase shift [4] which means that the vibration direction lies on the reflection plane. Therefore, by projecting the angle of polarization and surface normal onto the image plane, we can formulate the two constraints for the diffuse- and specular-polarized reflections:

$$[\sin(\phi), -\cos(\phi), 0]\mathbf{n} = 0, \quad (15a)$$

$$[\sin(\phi + 90^\circ), -\cos(\phi + 90^\circ), 0]\mathbf{n} = 0, \quad (15b)$$

where $\phi = \arctan(S_{dp}(2)/S_{dp}(1))/2$ is the angle of polarization in the diffuse case. Eq. 15a is the constraint for diffuse-polarized reflection. Eq. 15b is for the specular-polarized case, whose angle of polarization is shifted 90° .

Similar constraints are used in [44, 45]. However, most methods directly use the overall Stokes vector to compute the angle of polarization. We empirically show that we are able to estimate a more accurate angle of polarization by using the diffuse-polarized component.

In order to use the polarimetric constraints, we threshold $S_{sp}(0)$ to a binary mask that indicates the specular pixels. We use Eq. 15b as the additional constraint for the specular pixels, and Eq. 15a for all other pixels considered diffuse.

4.3. Depth and Normal Estimation

By combining the reciprocity and polarimetric constraints, we can form a linear system for the surface normal $\mathbf{n}$, *i.e.*, $\mathbf{W}(d)\mathbf{n} = 0$; this is similar to the standard HS [55]. The coefficient matrix $\mathbf{W}$ is a function of the surface depth d , and we solve for d and $\mathbf{n}$.

We first optimize the depth. Given the true depth d^* , the rank of $\mathbf{W}(d^*)$ should be 2, so that the surface normal $\mathbf{n}$ will be uniquely determined. If the depth value is incorrect, the rank of $\mathbf{W}$ will be greater than 2. This indicates that if we apply SVD on $\mathbf{W}$: $\mathbf{W} = U\Sigma V^\top$, where $\Sigma = \text{diag}(\sigma_1, \sigma_2, \sigma_3)$, $\sigma_1 \geq \sigma_2 \geq \sigma_3$, the ratio σ_2/σ_3 will be infinitely large at the true depth d^* . Thus, we use the exponential decay function proposed in [40] as our data term for depth estimation:

$$E_{data}(d) = \exp(-\mu \frac{\sigma_2(d)}{\sigma_3(d)}), \quad (16)$$

where $\mu = 0.2 \ln(2)$ [40]. We also use a smoothness term to reduce noise in the depth estimation :

$$E_{smooth}(d) = \sum_{(p,q) \in \mathcal{N}} \min(\|d_p - d_q\|, K), \quad (17)$$




Figure 4. Normal reconstruction with synthetic data. “White billiard” and “spectralon” are the surface materials.

where p and q are two pixels in the neighborhood $\mathcal{N}$; K is a truncation threshold that prevents the discontinuous depths from being smoothed. The combined cost function for depth estimation is:

$$d^* = \underset{d}{\operatorname{argmin}} \sum (E_{data} + \lambda E_{smooth}), \quad (18)$$

where λ is a balancing weight. In our experiments, we use $\lambda = 0.01$. We use graph-cut [10] to solve the depth as a multi-labeling problem.

Once the depth values are estimated, we can solve the normal $\mathbf{n}$ using $\mathbf{W}(d)\mathbf{n} = 0$. We then iteratively refine the depth and normal with the following steps: 1) we apply Poisson integration on $\mathbf{n}$ to obtain a new set of depth d' ; 2) we use d' as an additional guidance in Eq. 18 to optimize the depth with a finer depth interval; 3) we use the estimated depth to form $\mathbf{W}(d)$ and solve for the normal again. We use the normal difference to decide whether the refinement has converged. In our experiments, it usually converges after two iterations.

5. Experiments

We validate our method with both synthetic and real experiments on scenes with various shapes and reflectance.

5.1. Synthetic Experiments

We use the Mitsuba 2 renderer¹ to simulate polarimetric images as captured by a polarization camera. Specifically, we use the polarized rendering mode to simulate four directional polarization images: I^0 , I^{90} , I^{45} , and I^{135} (each with resolution 500×500). In the polarized rendering mode, the renderer will track the full polarization state of light during simulation. The system configuration mirrors our real experimental setup. We use the real-captured KAIST

¹<https://www.mitsuba-renderer.org>




Figure 5. Synthetic results with respect to different material types. We show rendered color images of the materials, recovered normal maps, and cross-sections of the recovered shapes.

Method \ SNR/dB	10	20	30	40
HS (3-pair)	23.41	13.59	8.05	6.87
polar-HS (1-pair)	14.71	8.05	6.53	6.27
polar-HS (2-pair)	19.16	6.43	2.14	0.95

Table 2. Mean angular error (in degree) of the normal estimations with respect to the noise levels.

pBRDF dataset [8] to model the polarimetric surface reflectance. We test on a variety of 3D models and surface reflectance. Fig. 4 shows our rendered images and recovered normal maps with two reciprocal pairs. We evaluate the reconstruction with per-pixel angular errors and the mean angular error (MAE). More results on normal and surface reconstruction can be found in the supplementary material.

Ablation on Material Types. We use different pBRDFs provided by the KAIST dataset on a sphere object to test our performance with respect to the material types. Fig. 5 shows recovered normal maps and cross-sections of the recovered shapes. Here the results are computed using one reciprocal pair. Results on more materials and two reciprocal pairs can be found in the supplementary material. We can see that the one-pair results are sensitive to the material type as the polarized reflection of some materials (*e.g.*, mint silicone) is weak, which results in the angle of polarization being highly noisy and unreliable. The two-pair results (see supplementary material) are more robust as the reciprocity constraint alone provides sufficient regularization.

Ablation on Noise Levels. We evaluate our method with respect to different levels of noise. In this experiment, we use a sphere object with the “white billiard” material. We add Gaussian white noise to the rendered images and use the signal-noise ratio (SNR) to quantify the noise level (smaller SNR number indicates higher noise level). We evaluate the reconstruction using MAE of the normal estimation. We test our method (polar-HS) using one and two reciprocal pairs respectively, and compare with the standard HS that




Figure 6. Quantitative evaluation of the normal and shape reconstruction on real scenes.

uses three pairs. The results are reported in Table 2. We can see that both our one-pair and two-pair methods are more accurate than the three-pair HS, and less sensitive to noise.

5.2. Real Experiments

Prototype. We implement a physical system for capturing the polar-HS image pairs. We mount a polarization camera and a pico projector on an automated rotating wheel so that their positions can be precisely exchanged. Both the camera and the project are calibrated geometrically [32] and radiometrically [25]. More details on the system construction can be found in the supplementary material. A reciprocal pair is captured by rotating the wheel at 180° . At each position in the pair, we capture two polarization images under 0° and 90° polarized light.

Quantitative Evaluation. We quantitatively evaluate our method on two real scenes: a billiard ball (59mm diameter) and a cardboard corner (corner angle of 110°). We compare the normal and surface reconstruction of our method with one pair and two pairs, and the standard HS with three pairs. The scene settings and the reconstruction results are shown in Fig. 6. We report the MAE of recovered normals, and compare the recovered shapes in cross-sections. Our one-pair reconstruction results have large errors at oblique polarization angles (*i.e.*, $\phi = 90^\circ$) where the polarimetric constraint becomes unstable. However, our two-pair results are highly accurate in both scenes.

Qualitative Evaluation. We test our method on a variety of scenes with different types of reflectance, ranging from purely diffuse to highly specular. Some are composite scenes that contain multiple surfaces types. Fig 7 shows our recovered surfaces. Here we demonstrate the two-pair results as they are more accurate than one-pair ones. Results on more scenes and one reciprocal pair can be found in the supplementary material. We can see that our recovered surfaces preserve fine geometric details (for example, bandage on the statue and beard of the gnome). Our method also works for concave scenes that do not exhibit very strong



Figure 7. Surface reconstruction results on real scenes. Here the results are computed with two reciprocal pairs.




Figure 8. Comparisons of our method with classical 3D reconstruction methods. We show the recovered normal maps (row one), relit surfaces (row two), and zoom-in views of the surfaces (row three).

interreflection (for example, the soap dish and bunny ears).

We also compare our method against classical 3D reconstruction methods: shape-from-polarization (SfP) [44], structured light (SL) [30], photometric stereo (PS) [53], and standard HS [55] with three pairs. We show visual comparison results on an orange scene in Fig. 8. The overall shapes of our one-pair and two-pair results are closer to the structured light scanned result (SL). Our method is also able to recover more surface details than SL.

6. Conclusion and Discussion

In summary, we extend the classical Helmholtz stereopsis to the polarimetric case by deriving a new transpositional reciprocity relationship. We exploit the polarimetric cues and reduce the minimal number of reciprocal pairs to perform HS to one. Our proposed polar-HS can recover various types of surfaces with high accuracy.

Limitations. Although we have demonstrated successful 3D reconstruction on a variety of scenes, our method has

limitations on handling strong interreflection and transparent scenes (see failure examples in the supplementary material). If the interreflection is too strong and results in caustic effects on the surface, our method fails at the regions where the caustics occurs. As for the transparent objects, the captured images are transmission dominant. However, our method relies on analyzing the reflected light for surface reconstruction. One possible solution is to separate the weak reflected image from the transmitted one.

Future Directions. As our one-pair solution suffers from large errors at oblique polarization angles, we plan to exploit additional physical constraints (*e.g.*, shading cues or multi-view constraint) for improvement. As our polar-HS only needs one reciprocal pair in theory, it is possible to build a compact acquisition system without using a rotating wheel (*e.g.*, in a binocular setting with polarization-multiplexing), which makes our solution more practical.

Acknowledgements. This project is supported by NSF awards CRII-1948524 and NRI-2024795.

References

- [1] Gilad Adiv. Inherent ambiguities in recovering 3-D motion and structure from a noisy flow field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(5):477–489, 1989.
- [2] Gary A. Atkinson. Polarisation photometric stereo. *Computer Vision and Image Understanding*, 160:158–167, 2017.
- [3] Gary A. Atkinson and Edwin R. Hancock. Multi-view surface reconstruction using polarization. In *IEEE International Conference Computer Vision (ICCV)*, 2005.
- [4] Gary A. Atkinson and Edwin R. Hancock. Recovery of surface orientation from diffuse polarization. *IEEE Transaction on Image Processing*, 15(6):1653–1664, 2006.
- [5] Gary A. Atkinson and Edwin R. Hancock. Shape estimation using polarization and shading from two views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):2001–2017, 2007.
- [6] Yunhao Ba, Alex Gilbert, Franklin Wang, Jinfa Yang, Rui Chen, Yiqin Wang, Lei Yan, Boxin Shi, and Achuta Kadambi. Deep shape from polarization. In *European Conference on Computer Vision (ECCV)*, 2020.
- [7] Seung-Hwan Baek, Daniel S. Jeon, Xin Tong, and Min H. Kim. Simultaneous acquisition of polarimetric svbrdf and normals. *ACM Transaction on Graphics*, 37(6), Dec. 2018.
- [8] Seung-Hwan Baek, Tizian Zeltner, Hyun Jin Ku, Inseung Hwang, Xin Tong, Wenzel Jakob, and Min H. Kim. Image-based acquisition and modeling of polarimetric reflectance. *ACM Transactions on Graphics (Proc SIGGRAPH 2020)*, 39(4), 2020.
- [9] Ronen Basri, David Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72:239–257, 2007.
- [10] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001.
- [11] Subrahmanyan Chandrasekhar. *Radiative transfer*. Courier Corporation, 2013.
- [12] Lixiong Chen, Yinqiang Zheng, Art Subpa-Asa, and Imari Sato. Polarimetric three-view geometry. In *European Conference on Computer Vision (ECCV)*, 2018.
- [13] Alessandro Chiuso, Roger Brockett, and Stefano Soatto. Optimal structure from motion: Local ambiguities and global estimates. *International Journal of Computer Vision*, 39:195–228, 2000.
- [14] F. Clarke and D. Parry. Helmholtz reciprocity: its validity and application to reflectometry. *Lighting Research & Technology*, 17:1–11, 1985.
- [15] Yan Cui, Sebastian Schuon, Derek Chan, Sebastian Thrun, and Christian Theobalt. 3D shape scanning with a Time-of-Flight camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [16] Zhaopeng Cui, Jinwei Gu, Boxin Shi, Ping Tan, and Jan Kautz. Polarimetric multi-view stereo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [17] Amaël Delaunoy, Emmanuel Prados, and Peter N. Belhumeur. Towards full 3d Helmholtz stereovision algorithms. In *Asian Conference on Computer Vision (ACCV)*, 2010.
- [18] Ondřej Drbohlav and Radim Šára. Unambiguous determination of shape from photometric stereo with unknown light sources. In *IEEE International Conference Computer Vision (ICCV)*, 2001.
- [19] Jinwei Gu, Toshihiro Kobayashi, Mohit Gupta, and Shree K. Nayar. Multiplexed illumination for scene recovery in the presence of global illumination. In *IEEE International Conference Computer Vision (ICCV)*, 2011.
- [20] Mohit Gupta, Amit Agrawal, Ashok Veeraraghavan, and Srinivasa G. Narasimhan. Structured light 3D scanning in the presence of global illumination. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [21] Mohit Gupta and Shree K. Nayar. Micro Phase Shifting. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [22] Felix Heide, Lei Xiao, Wolfgang Heidrich, and Matthias B. Hullin. Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive Time-of-Flight sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [23] Tomoaki Higo, Yasuyuki Matsushita, and Katsushi Ikeuchi. Consensus photometric stereo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [24] Cong Phuoc Huynh, Antonio Robles-Kelly, and Edwin R. Hancock. Shape and refractive index recovery from single-view polarisation images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [25] Zsolt Jankó, Ondřej Drbohlav, and Radim Šára. Radiometric calibration of a Helmholtz stereo rig. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [26] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Polarized 3D: High-quality depth sensing with polarization cues. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [27] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Depth sensing using geometrically constrained polarization normals. *International Journal of Computer Vision*, 125(1-3):34–51, 2017.
- [28] Daisuke Miyazaki, Takuya Shigetomi, Masashi Baba, Ryo Furukawa, Shinsaku Hiura, and Naoki Asada. Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In *International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, 2012.
- [29] Daisuke Miyazaki, Robby T. Tan, Kenji Hara, and Katsushi Ikeuchi. Polarization-based inverse rendering from a single view. In *IEEE International Conference Computer Vision (ICCV)*, 2003.
- [30] Raymond A. Morano, Cengizhan Ozturk, Robert Conn, Stephen Dubin, Stanley Zietz, and Jonathan Nissanov. Structured light using pseudorandom codes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):322–327, 1998.

- [31] Olivier Morel, Fabrice Meriaudeau, Christophe Stolz, and Patrick Gorria. Polarization imaging applied to 3D reconstruction of specular metallic surfaces. In *SPIE Machine Vision Applications in Industrial Inspection XIII*, volume 5679, 2005.
- [32] Daniel Moreno and Gabriel Taubin. Simple, accurate, and robust projector-camera calibration. In *International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, 2012.
- [33] Hironori Mori, Roderick Köhle, and Markus Kamm. Scene depth profiling using Helmholtz stereopsis. In *European Conference on Computer Vision (ECCV)*, 2016.
- [34] Volker Müller. Elimination of specular surface-reflectance using polarized and unpolarized light. In *European Conference on Computer Vision (ECCV)*, 1996.
- [35] Nikhil Naik, Achuta Kadambi, Christoph Rhemann, Shahram Izadi, Ramesh Raskar, and Sing Bing Kang. A light transport model for mitigating multipath interference in Time-of-Flight sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [36] Shree K. Nayar, Xi-Sheng Fang, and Terrance E. Boult. Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3):163–186, 1997.
- [37] Trung Thanh Ngo, Hajime Nagahara, and Rin-ichiro Taniguchi. Shape and light directions from shading and polarization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [38] Stefan Rahmann and Nikos Canterakis. Reconstruction of specular surfaces using polarization imaging. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.
- [39] Nadejda Roubtsova and Jean-Yves Guillemaut. Bayesian Helmholtz stereopsis with integrability prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(9):2265–2272, 2017.
- [40] Nadejda Roubtsova and Jean-Yves Guillemaut. Colour Helmholtz stereopsis for reconstruction of dynamic scenes with arbitrary unknown reflectance. *International Journal of Computer Vision*, 124(1):18–48, 2017.
- [41] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47:7–42, 2004.
- [42] Steven M. Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [43] Zdeněk Sekera. Scattering matrices and reciprocity relationships for various representations of the state of polarization. *Journal of the Optical Society of America*, 56(12):1732–1740, 1966.
- [44] William A. P. Smith, Ravi Ramamoorthi, and Silvia Tozza. Linear depth estimation from an uncalibrated, monocular polarisation image. In *European Conference on Computer Vision (ECCV)*, 2016.
- [45] William A. P. Smith, Ravi Ramamoorthi, and Silvia Tozza. Height-from-polarisation with unknown lighting or albedo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(12):2875–2888, 2019.
- [46] Christoph Strecha, Wolfgang von Hansen, Luc Van Gool, Pascal Fua, and Ulrich Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [47] Silvia Tozza, William A. P. Smith, Dizhong Zhu, Ravi Ramamoorthi, and Edwin R. Hancock. Linear differential constraints for photo-polarimetric height estimation. In *IEEE International Conference Computer Vision (ICCV)*, 2017.
- [48] Peter Tu and Paulo R. S. Mendonça. Surface reconstruction via Helmholtz reciprocity with a single image pair. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.
- [49] Shimon Ullman. *The Interpretation of Visual Motion*. MIT Press, 1st edition, 1979.
- [50] Hermann von Helmholtz. *Physiological Optics*. Optical Society of America, 1924.
- [51] Michael Weinmann, Roland Ruiters, Aljosa Osep, Christopher Schwartz, and Reinhard Klein. Fusing structured light consistency and Helmholtz normals for 3D reconstruction. In *British Machine Vision Conference (BMVC)*, 2012.
- [52] Lawrence B. Wolff and Terrance E. Boult. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):635–657, 1991.
- [53] Robert J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [54] Dizhong Zhu and William A. P. Smith. Depth from a polarisation + RGB stereo pair. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [55] Todd E. Zickler, Peter N. Belhumeur, and David J. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision*, 49(2-3):215–227, 2002.
- [56] Todd E. Zickler, Jeffrey Ho, David J. Kriegman, Jean Ponce, and Peter N. Belhumeur. Binocular Helmholtz stereopsis. In *IEEE International Conference Computer Vision (ICCV)*, 2003.