

Deep Polarization Cues for Single-shot Shape and Subsurface Scattering Estimation

Chenhao Li¹, Trung Thanh Ngo², and Hajime Nagahara¹

¹ Osaka University, 2-8 Yamadaoka, Suita, Osaka, Japan

² Hanoi University of Science and Technology, 1 Dai Co Viet, Hai Ba Trung, Hanoi, Vietnam

Abstract. In this work, we propose a novel learning-based method to jointly estimate the shape and subsurface scattering (SSS) parameters of translucent objects by utilizing polarization cues. Although polarization cues have been used in various applications, such as shape from polarization (SfP), BRDF estimation, and reflection removal, their application in SSS estimation has not yet been explored. Our observations indicate that the SSS affects not only the light intensity but also the polarization signal. Hence, the polarization signal can provide additional cues for SSS estimation. We also introduce the first large-scale synthetic dataset of polarized translucent objects for training our model. Our method outperforms several baselines from the SfP and inverse rendering realms on both synthetic and real data, as demonstrated by qualitative and quantitative results.

Keywords: Inverse Rendering · Subsurface scattering · Shape from Polarization

1 Introduction

Translucent objects, such as human skin, milk, wax, and crystals, are omnipresent in our daily lives. A physical phenomenon that substantially influences the appearance of translucent objects is subsurface scattering. It occurs when photons penetrate the surface of an object and scatter inside it. After multiple scattering events, the photons are either absorbed by the object or exit from another surface point. Accurate estimation of subsurface scattering is crucial for various applications, including virtual reality, material science, and computer graphics.

However, estimating parameters for translucent objects is an ill-posed problem, even with the homogeneous assumption. In contrast to works [7, 16, 48] that solely consider the interaction between light and the object on the surface, parameter estimation of translucent objects is more challenging since it requires accounting for both surface reflection and the multiple bounces and paths of SSS. Researchers often simplify scene representations to solve the parameter estimation problem of translucent objects. For instance, they may ignore surface




Fig. 1: Visual results on real-world translucent objects. Our method takes four linearly polarized images and a binary mask as input. We use Mitsuba [57] to render a sphere to visualize estimated SSS parameters. For a better understanding of the quality of the estimated shape (Normal and Depth), we provide the shading and relighting results.

reflection [9, 73] or use the simplified models such as Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF) to approximate optically thick materials in the form of complex surface reflection [14, 17, 29, 70], or dipole-diffusion approximation [76] that can only handle diffuse subsurface light transport. To the best of our knowledge, using polarization cues for SSS parameter estimation in the presence of surface reflections has not been explored, and we make the first attempt in this work.

Light is an electromagnetic wave in which the wavelength determines its color, the amplitude determines its intensity, and the geometrical direction of oscillation leads to its polarization. While the human visual system is sensitive to color and intensity, for most of the natural scenes, we barely perceive polarization. Nevertheless, optical elements such as polarizers or wave plates can be used to measure the polarization of light. The measurement of polarization offers a plethora of computer vision applications. When light interacts with an object, the polarization changes by the surface orientation and material, which can be used as an additional cue to estimate the intrinsic factors of the object.

Polarization is a suitable cue for SSS estimation. From the theoretical view, we can roughly classify polarization into diffuse and specular polarization. Diffuse polarization arises from the summation of SSS and multiple bounces between microsurfaces, while specular polarization is produced from single-bounce surface reflection. The polarization signal of reflected light (specular polarization) is distinct from that of refracted light after multiple scattering (diffuse polarization). The reflected light is strongly polarized, particularly near Brewster’s angle, whereas the light after multiple SSS is nearly unpolarized. SSS affects the ratio between diffuse and specular polarization, ultimately altering the captured polarization signal. This is the reason that polarization can be beneficial for SSS estimation. From the practical view, polarization cameras enable capturing four

linearly polarized images in a single shot, making the capturing feasible even in non-laboratory settings.

The challenges of polarization for SSS estimation derive from various ambiguities. One of the ambiguities worth mentioning is called diffuse/specular polarization ambiguity. Translucent objects exhibit both types of polarization simultaneously, making it difficult to determine whether the captured polarization signal originates from diffuse polarization, specular polarization, or both. Most existing works [11, 31, 61] exclusively consider either specular or diffuse polarization. A few researchers consider both types of polarization simultaneously; however, these methods are either applicable only to opaque objects [5] or use a simplified SSS model [4] in the form of Gaussian distribution.

This paper proposes a learning-based approach for the simultaneous estimation of shape and SSS parameters using polarization cues. We employ several strategies to mitigate the aforementioned ambiguity issues. Our method uses four linearly polarized images (captured at polarizer angles of 0 , $\frac{\pi}{4}$, $\frac{\pi}{2}$, and $\frac{3\pi}{4}$) as input for each translucent object. We propose novel *max* & *min* polarization representations and also use them as additional inputs for reducing the ambiguity of surface reflection and SSS since the representations are related to the ratio of specular and diffuse polarization. Inspired by the previous work [7] that step-by-step estimates shape and BRDF parameters to reduce ambiguities, we also train a stage-wise deep neural network that first estimates geometry and illumination, and then uses them as a guide for SSS estimation. We propose a novel reconstruction network to disambiguate SSS estimation. Using a differentiable renderer to reconstruct the scene is commonly used in previous inverse rendering works to improve the accuracy of estimation. However, there is no general-purpose differentiable renderer for the polarized inverse scattering, we bypass this issue by using a deep neural network to reconstruct the scene and compute a reconstruction network. To train our network, we constructed a large-scale synthetic dataset comprising 117K scenes. In each scene, we rendered a human-created 3D model with a specular BSDF and homogeneous SSS under environmental illumination using a polarized renderer. Our contributions can be summarized as follows:

- We introduce polarization cues to the deep learning domain for the joint estimation of shape and SSS parameters.
- We propose *max* & *min* polarization representations using an additional input as explicit cues for specular reflection and SSS separation.
- We propose a reconstruction network for supervising SSS parameters from images.
- We built a large-scale synthetic dataset of polarized translucent objects.

2 Related Work

2.1 Inverse rendering

Inverse rendering is a longstanding problem in computer vision that seeks to estimate material, shape, and illumination from single or multiple views. In this

section, we will focus exclusively on single-view methods. We roughly divide these methods into three categories: pure surface reflection models, pure SSS models, and both.

So far, **pure surface reflection** models have dominated inverse rendering. These approaches assume that light interacts with objects only at the surface. Numerous studies have been published for estimating shape [18, 44, 60], BRDF [1, 2, 15, 45, 47, 54], and illumination [21, 22], either separately or jointly [62, 69]. Li *et al.* [48] achieved a milestone in this field by proposing a method that simultaneously estimates shape, illumination, and BRDF using deep neural networks. Later, some studies have shown that the performance could be further improved by using a flash and no-flash image pair [7], a recursive network [50], polarization cues [16], or an RGB-D camera [35]. Meanwhile, some researchers have demonstrated that Li *et al.* [48] could be extended to scene-level inverse rendering by introducing a Residual Appearance [63] or spatially-varying illumination [46, 68, 75].

Another group of researchers has focused on **pure SSS** models. Such scene representations are often used to model participating media without clear boundaries, such as smoke, dust, or fog. Traditional methods [24, 25, 33, 40, 72] were mainly based on analysis by synthesis. However, they suffered from many issues like long optimization time and local minimum. A recent breakthrough has been made by Che *et al.* [9]. They combined a neural network and a differentiable renderer [57] and trained them end-to-end to estimate SSS parameters.

Only a few works have considered **both surface reflection and SSS**, and it is still an open research topic. Our method belongs to this group. Baek *et al.* [4] used a sophisticated temporal-polarimetric capturing system to simultaneously optimize the surface and subsurface parameters. However, the rendering model they used for SSS is a simple depolarization model in the form of Gaussian distribution. Additionally, they suffered from a long optimization time. Compared to their work, we use polarization to estimate SSS parameters. Besides, our method realizes the parameter estimation at an acceptable cost. The inference process is just a single forward propagation. At the same time, our capturing requirement is more flexible, allowing users to capture images without a laboratory environment. Li *et al.* [41] introduced a novel method that uses flash and non-flash image pairs to disambiguate the inverse scattering. Compared to their method, we allow the data capture within just a single shot and do not require additional lighting information.

2.2 Polarization

As noted, polarization is an active research area in computer vision and graphics. Numerous works have been reported for various applications such as SfP [3, 10, 19, 28, 31, 38, 64, 74, 77], BRDF estimation [5, 16, 27, 34, 71], depth estimation [67], image segmentation [32, 49, 53, 59], reflection removal [37, 43, 51], white balance [58], pose estimation [12, 20, 78], compass [66], and sensor design [36]. This section only discusses some main topics highly relevant to our work.

SfP has been extensively studied and can be considered a sub-task of our work. The theoretical basis of SfP comes from the Fresnel equations. When a beam of light interacts with the object’s surface, the polarization changes according to the surface orientation and material of the object. Therefore, the measured polarization signal can be used for shape estimation. The challenges of SfP arise from various ambiguities. For example, the current capturing system cannot disambiguate the polarization signal for the Angle of Polarization θ and $\theta + \pi$. Also, the ambiguity between specular and diffuse polarization is another critical issue. To reduce the difficulty of the problem, researchers usually use a simplified polarization model. For example, most existing works [11, 31, 61] exclusively consider specular and diffuse polarization. Another research direction addresses ambiguity by combining polarization with cues such as shading [52, 56, 65], depth [31], or time-of-light [4]. Recent breakthroughs have been dominated by deep learning, using convolutional networks to compensate for ambiguity. Ba *et al.* [3] proposed the first learning-based method for normal estimation. Lei *et al.* [38] extended their work to the scene-level normal estimation using a viewing encoding to handle the non-orthographic projection problem.

A recent trend in polarization research is the **joint estimation of normal and reflectance**. Baek *et al.* [5, 6] proposed a model that captures the polarimetric BRDF and surface normal simultaneously. The proposed method was later improved by utilizing the time-of-flight cues [4]. Deschaintre *et al.* [16] were the first to introduce the joint estimation of BRDF and surface normal to the realm of deep learning. Dave *et al.* [13] successfully addressed the joint normal and BRDF estimation problem by using an implicit radiance field of polarization. Li *et al.* [42] solved the multiple bounced polarimetric light problem using a hybrid light path representation.

Finally, we discuss some polarization works on **challenging materials** such as transparent objects. The difficulty of dealing with transparent objects is that they inevitably exhibit reflection and refraction at the same time, and it is not easy to separate the reflection/refraction portion. Our target objects (translucent objects) face similar challenges to these works (transparent objects). Polarization is a suitable cue for solving such problems because the polarization signal for reflected and refracted light is different. Many existing works used this property for applications such as reflection removal [37], glass segmentation [32], and normal estimation for transparent objects [64].

3 Background

Light and optical element representation. Following most SfP works [3, 16, 38], we only consider linear polarization. The linear polarization state of a beam of light is usually described by Stokes vectors $\mathbf{s} = [s_0, s_1, s_2]$. Where $s_0 = I$ represents the unpolarized image intensity, s_1 represents the horizontal vs. vertical polarization, and s_2 represents the 45-degree polarization vs. the 135-degree polarization. Then the Degree of Polarization (DoP) ρ and the Angle




Fig. 2: (a) SSS influences the polarization of objects. Images of a bottle (BSDF) were captured under the same illumination with three different liquids (SSS). For black coffee, most photons are absorbed upon entering the object, and the surface reflection dominates. For the milk, photons are back-scattered, and the SSS contributes significantly to the overall appearance. The latté falls somewhere in between these two extremes. (b) Our scene representation. We assume unpolarized light sources. The captured light intensity consists of two components. One comes from the single-bounce surface reflection, which exhibits **specular polarization** (orange path). Another one is the refracted light (blue path): An unpolarized light refracts into the object and becomes partially polarized light. Then, after undergoing the multi-bounce SSS, it becomes unpolarized light. Finally, the light leaves the object, undergoes Fresnel refraction again, and exhibits **diffuse polarization**.

of Polarization(AoP) ϕ can be computed as follows:

$$\rho = \frac{\sqrt{s_1^2 + s_2^2}}{s_0}, \quad (1)$$

$$\phi = \frac{1}{2} \arctan \frac{s_2}{s_1}. \quad (2)$$

Mueller matrices $\mathbf{M} \in \mathbb{R}^{3 \times 3}$ are usually used to represent optical elements that change the polarization state. The final captured Stokes vector $\mathbf{s}_{\text{out}}$ is an integral over all possible light paths:

$$\mathbf{s}_{\text{out}} = \int_{\mathbb{P}} f(\bar{x}) d\bar{x}, \quad (3)$$

where $\mathbb{P}$ is the space of all light paths, $\bar{x}$ is a single path, and f is the path contribution function.

Our scene assumption. Following existing inverse rendering works [7, 48], we assume the presence of two unpolarized light sources - natural lighting and a camera flashlight. The camera flashlight dominates the illumination. For the translucent object, we assume a smooth surface and a homogeneous SSS (SSS

parameters are uniform within the object). A specular BSDF describes the object boundary, and similar to most SfP works [3], we assume a constant Index of Refraction (IoR) of 1.5046 for the BSDF. To model the homogeneous SSS, we use the Radiative Transport Equation (RTE). The Henyey-Greenstein phase function [26] is used to represent the probability distribution of the outgoing direction when a photon interacts with particles inside the object. In theory, SSS depolarizes incident light based on the distance the light travels within the object [23]. However, we assume that the light travels a sufficiently long distance and is already unpolarized before leaving the object's surface.

Motivation of Polarization for SSS estimation. We analyze two typical path contribution functions that majorly affect the captured Stokes Vectors $\mathbf{s}_{\text{out}}$. The first one is:

$$f_r = \mathbf{M}_r \mathbf{s}_{\text{in}}, \quad (4)$$

where $\mathbf{s}_{\text{in}}$ is the Stokes vector of the light source, and $\mathbf{M}_r$ is the Mueller matrices of dielectric reflection. For convenience, we omit the rotation matrices here. This path describes the single bounce surface reflection (Orange path in Figure 2 (b)) and exhibits **specular polarization**. The second one is:

$$f_t = \mathbf{M}_t^o \left[\prod_i \sigma_t f_p(\theta_i, g) G(i, i+1) T(i, i+1) \right] \mathbf{M}_d \mathbf{M}_t^i \mathbf{s}_{\text{in}}, \quad (5)$$

where $\mathbf{M}_t^o$ and $\mathbf{M}_t^i$ are the Mueller matrices of dielectric transmission, $\mathbf{M}_d$ is a depolarizer, σ_t is the extinction coefficient, f_p is the phase function, G is the geometry term, and T is the transmission term. Refer to the blue path in Figure 2 (b). $\mathbf{M}_t^o$ creates **diffuse polarization** when the light leaves the object's surface. SSS can affect the light intensity before $\mathbf{M}_t^o$, thus affecting f_t . The final polarization $\mathbf{s}_{\text{out}}$ is affected by the ratio between f_t and f_r . Therefore, SSS can contribute to $\mathbf{s}_{\text{out}}$. In other words, $\mathbf{s}_{\text{out}}$ can be considered as a cue for SSS estimation. From Figure 2 (a), we can observe that by switching the SSS (liquid in the bottle), the captured polarization signal changes significantly. Specifically, when the liquid is dark (the 1st row), most photons are absorbed during SSS. In this case, the reflected portion f_r dominates, and the object exhibits specular polarization. When the liquid turns white (the 2nd and 3rd rows), the SSS portion f_t increases and the object exhibits diffuse polarization. For most incident angles, the DoP of specular polarization is stronger than diffuse polarization. Therefore, we can observe that the DoP gets smaller as the liquid gets brighter. The AoP between the diffuse and specular polarization is $\frac{\pi}{2}$. Thus, the AoP also changes according to the SSS. This is why we bring polarization into the realm of SSS estimation.

Polarization capturing. In our case, we use a polarization camera to capture four linearly polarized images (I_0 , I_{45} , I_{90} , and I_{135}) in a single shot. The Stokes vectors can be calculated by the following equations:

$$s_0 = \frac{1}{2}(I_0 + I_{45} + I_{90} + I_{135}), \quad (6)$$

$$s_1 = I_{90} - I_0, \quad (7)$$




Fig. 3: Overview of the proposed model. It takes four linearly polarized images ($I_0, I_{45}, I_{90}, I_{135}$), a *max & min* physical prior ($I_{\max}, I_{\min}$), and a binary mask (M) as input. We use a stage-wise network architecture, and the model estimates shape and illumination first, then uses the estimated shape and illumination to guide the SSS parameter estimation. A novel reconstruction network whose inputs are four pure BSDF images ($I_{b0}, I_{b45}, I_{b90}, I_{b135}$) and the estimated SSS parameters is proposed to further reduce the ambiguity of SSS estimation. Note that the reconstruction network is only used during the training.

$$s_2 = 2 * I_{45} - s_0. \quad (8)$$

Then, we can obtain images with the maximum intensity $I_{\max}$ and minimum intensity $I_{\min}$:

$$I_{\max} = \frac{1}{2}(s_0 + \sqrt{s_1^2 + s_2^2}), \quad (9)$$

$$I_{\min} = \frac{1}{2}(s_0 - \sqrt{s_1^2 + s_2^2}). \quad (10)$$

We use these *max & min* polarization representations because, for most incident angles, the specular polarization is stronger than the diffuse polarization. Therefore, when specular polarization dominates, the difference between $I_{\max}$ and $I_{\min}$ is significant, and we can observe a slight difference when diffuse polarization dominates. The network can learn the relationship between the specular and diffuse polarization and the *max & min* images work for the separation cue between surface reflection and SSS components.

4 Our Model

Given a translucent object with unknown illumination, shape, and material, our target is to estimate the shape and SSS parameters simultaneously. Similar to existing works [7], we use a depth map D to represent coarse geometry and a normal map N to provide local details. The illumination is also estimated as

a side prediction to assist in the estimation. Following previous works [48], we use the first three-order spherical harmonics sh to model the environment light and a flash intensity i to determine the brightness of the camera flashlight. We model our SSS with three terms: a volumetric albedo α , which determines the probability of whether photons are scattered or absorbed during a volume event; an extinction coefficient σ_t , which defines the optical density; and a Henyey-Greenstein phase function [26] parameter g which controls whether the scattering is forward ($g > 0$), backward ($g < 0$) or isotropic ($g = 0$). In summary, our **input** images are four linearly polarized images I_0, I_{45}, I_{90} , and $I_{135} \in \mathbb{R}^{H \times W \times 3}$, two *max* & *min* images $I_{\max}$ and $I_{\min} \in \mathbb{R}^{H \times W \times 3}$, and a binary mask $M \in \mathbb{R}^{H \times W}$ that locates the object. Our **output** parameters are:

- Shape: A normal map $N \in \mathbb{R}^{H \times W \times 3}$ and a depth map $D \in \mathbb{R}^{H \times W}$.
- Illumination: The spherical harmonics $sh \in \mathbb{R}^{3 \times 9}$ and flashlight intensity $i \in \mathbb{R}$.
- SSS: The extinction coefficient $\sigma_t \in \mathbb{R}^3$, volumetric albedo $\alpha \in \mathbb{R}^3$, and Henyey-Greenstein phase function parameter $g \in \mathbb{R}$.

The existing works on BRDF estimation [7] have used a cascaded network to address the shape/reflection ambiguity. Inspired by them, we also adopted this architecture. To be specific, we first estimate the shape and illumination, and then use the estimated parameters for SSS estimation. The model architecture is illustrated in Figure 3, and further details will be given in the following sections.

4.1 Shape and illumination estimation

For depth and normal estimation, we use a UNet-like network from a SOTA SfP work [38]. We concatenate four polarization images, *max* & *min* images, and a binary mask in the channel dimension. The depth and normal are estimated via two separate networks f_D and f_N :

$$\tilde{N} = f_N(P), \quad (11)$$

$$\tilde{D} = f_D(P), \quad (12)$$

where $P = (I_0, I_{45}, I_{90}, I_{135}, I_{\max}, I_{\min}, M)$ represents the input. The training losses of the estimated depth $\tilde{D}$ and normal $\tilde{N}$ are based on the L_1 distance between the estimated parameters and their GTs.

We also estimate the illumination as a side prediction to assist with the SSS estimation. For the illumination network, we use a small CNN followed by a fully connected layer, with the same input as the depth and normal estimation. For the output, we use the first three-order spherical harmonics, which results in nine coefficients for each RGB channel. Combining these coefficients with the one-dimensional flash-light intensity, we have a total of 28 parameters for illumination:

$$\tilde{sh}, \tilde{i} = f_{\text{illum}}(P). \quad (13)$$

The illumination net f_{illum} is trained using the L_2 loss between the estimated coefficients $\tilde{sh}$ and flash light intensity $\tilde{i}$ and their GTs.

4.2 Shape and illumination guided SSS estimation

The final appearance of an image is affected by the interplay of shape, illumination, and material. With the known shape and illumination, the ambiguity of SSS parameters estimation can be significantly reduced. Therefore, in addition to the original inputs, we also input the estimated shape and illumination information into the SSS net f_{SSS} :

$$\tilde{\sigma}_t, \tilde{\alpha}, \tilde{g} = f_{\text{SSS}}(P, \tilde{N}, \tilde{D}, \tilde{s}\tilde{h}, \tilde{i}), \quad (14)$$

where $\tilde{\sigma}_t, \tilde{\alpha}, \tilde{g}$ are estimated extinction coefficients, volumetric albedo, and phase function parameters, respectively. We use a ResNet-like architecture for our SSS net. The SSS net is also trained by L_2 loss. Note that, similar to the idea of shape and illumination guided SSS estimation. Conversely, knowing the SSS can also disambiguate the shape or illumination estimation. Existing works [39, 55] used an iterative manner or cascaded networks [7, 48] to achieve this. For simplicity and computational efficiency, we only use one module, which is the shape and illumination-guided SSS estimation. However, the proposed model can be extended by cascading more X-guided X estimation modules.

4.3 Reconstruction network

In this section, we propose a novel reconstruction network to further disambiguate SSS estimation. As mentioned before, similar appearances can be achieved by different combinations of illumination, shape, and material, which makes inverse rendering extraordinarily challenging and ill-posed. To alleviate this problem, a typical approach is using a differentiable renderer [9] to reconstruct the scene and compute the reconstruction network. However, there are several issues. First, to the best of our knowledge, there is no existing differentiable renderer that supports both surface reflection and SSS at the same time in polarized mode. While the latest work mitsuba3 [30] claims to support differentiable polarization, the differentiable polarization mode is incompatible with our scene representation (SSS+BSDF). Second, a well-known problem with differentiable renderers is the large memory usage and computational cost, and this problem is exacerbated in polarized mode.

An advantage of using synthetic data is that we can create specific variations to the original dataset so that we can compute additional losses to train our model. Specifically, for each scene, we additionally render four pure BSDF polarized images (I_{b0}, I_{b45}, I_{b90} , and I_{b135}) by removing the SSS (see Figure 3 for reference). Except for the SSS, the pure BSDF images and the original four polarized images have exactly the same illumination, shape, and BSDF. After Equation 14, we input the estimated SSS parameters and pure BSDF images into a reconstruction network f_{re} to reconstruct the original input images:

$$\tilde{I}_0, \tilde{I}_{45}, \tilde{I}_{90}, \tilde{I}_{135} = f_{\text{re}}(I_{b0}, I_{b45}, I_{b90}, I_{b135}, \tilde{\sigma}_t, \tilde{\alpha}, \tilde{g}). \quad (15)$$

The reconstructed images are also penalized by the L_1 loss between their GT images I_0, I_{45}, I_{90} , and I_{135} . Our reconstruction network has the following merits: During reconstruction, the shape and illumination information is implicitly

provided by the pure BSDF images, so there is less ambiguity of SSS. In addition, compared to the loss that is directly computed from a low-dimensional parameter space, a high-dimensional image space loss can provide more detail for gradient computing. Note that pure BSDF images are only required during training.

5 Experiment

5.1 Training details

The training was divided into two stages. We first trained the f_N , f_D , and f_{illum} for 20 epochs using an Adam optimizer with a constant learning rate equal to 0.0002 for the first 10 epochs and a linearly decayed learning rate for the remaining 10 epochs. The batch size was set to 64. After that, we used the same approach to train the f_{SSS} and f_{re} . Note that in the second stage, we did not freeze the weights of f_N , f_D , and f_{illum} . Instead, they continued to be trained.

5.2 Datasets

In this section, we introduce the proposed synthetic dataset. A key issue in deep learning for inverse rendering is the lack of training datasets. Measuring the physical parameters of real-world objects, such as BRDF, SSS, and surface normals, can be time-consuming, especially with a large number of objects. With the rise of the computer graphics community, the current physically-based renderer can produce photorealistic images even for polarization. Therefore, many existing works [16, 48] use synthetic datasets for network training.

In this work, we used Mitsuba 2 [57] for rendering the **synthetic dataset**. For 3D objects, we selected 5,847 human-created 3D objects from ShapeNet [8]. To create more variation of the surface normal of 3D objects, we collected 2,745 bump maps from several public sources. For the environment illumination, we used the Laval Indoor HDR dataset [21], which consists of 2,357 high-resolution indoor panoramas. Additionally, we also used a small sphere area light to mimic the camera flashlight. Each scene contains a randomly sampled 3D object, bump map, environment map, and SSS parameters. In total, we obtained 117K scenes, and each scene contains four linearly polarized images, four pure BSDF polarized images, and their GT parameters. All images' resolutions are 256×256 with 4,096 samples per pixel. We used 100K for training and 17K for testing. To verify that the model trained on the synthetic dataset can be applied to the real-world translucent objects, we also constructed a **real-world dataset**. We calibrated our capturing system to mimic the synthetic data scene representation. For each object, we used a 3D scanner to obtain the shape and used Mitsuba to optimize the position to calculate the GT surface normals. We demonstrate the results on our real-world dataset in Figure 1.




Fig. 4: Visual comparison of estimated normal with Deep SfP [3] and SfPW [38]. Mean angular errors are provided on the top-left corners.

Table 1: MAE results on our 17K synthetic test dataset.

	N	D	σ_t	α	g
Che <i>et al.</i> [9]	-	-	.1714	.1026	.2015
Deep SfP [3]	.0831	-	-	-	-
SfPW [38]	.0652	-	-	-	-
<i>w/o</i> Pol	.0750	.0512	.1398	.0957	.1717
4 Pol	.0657	.0487	.1369	.0928	.1602
4 Pol + DoPAoP	.0641	.0486	.1358	.0969	.1603
4 Pol + MaxMin	.0649	.0483	.1345	.0931	.1558
Ours (4 Pol + MaxMin + Re)	.0651	.0473	.1340	.0918	.1554

5.3 Comparison

We evaluate our approach by comparing it with the methods from the SfP and inverse rendering realm. During the following experiments, all methods are trained on our dataset with the same hyper-parameters.

Comparison to SfP baselines. Since the recent progress of SfP was dominated by deep learning, we only compare our methods with learning-based SfP. Finally, We selected two methods which are Deep SfP [3] and SfPW [38]. **Deep SfP** is the first learning-based SfP method. They used four polarization images and pseudo-normal as input to train a UNet-like network. **SfPW** extended Deep SfP to the scene-level normal estimation by introducing a novel view encoding and an encoded AoP. We illustrate a visual comparison in Figure 4 and report the Mean Absolute Error (MAE) in Table 1. It is observed that SfPW tends to estimate plane normals, while the results of Deep SfP are noisy. Our method can estimate high-quality normals for both synthetic and real data.

Comparison to inverse rendering baselines. It is not easy to find competitors for our method in the field of inverse rendering. Compared with **pure surface reflection models** is impossible. Because pure surface reflection models [7, 16, 48] are usually modeled by BRDF, for our scene representation, we use




Fig. 5: Visual comparison with an inverse scattering method. The 1st row is the GT intensity images. Images in the 2nd row are rendered by the GT shape, illumination, and estimated SSS by Che *et al.* [9]. The 3rd row is the images rendered by SSS parameters estimated by our method. Error maps are provided in the upper right corner.

BSDF to model the boundary and RTE to model the SSS. Thus, the parameters to be estimated are entirely different. Although Baek *et al.* [4] also considered **both surface and SSS**, they only used a simplified SSS model for depolarization in the form of a Gaussian distribution. Although Li *et al.* [41] has a similar task as our method, comparing with them is unfair as they require additional lighting conditions and images as inputs. Therefore, it is not suitable to compare with their method. Finally, we selected a **pure SSS** method proposed by Che *et al.* [9] to compare with our model. They trained an encoder-renderer model for SSS estimation, which can be considered a subtask of our work. Because it is not easy to compare the quality of estimated SSS without GT references of real-world objects, we only compare their method on the proposed synthetic dataset in Figure 5 and report MAE results in Table 1. Our method can estimate more accurate SSS parameters, and the images rendered by our estimated SSS parameters are close to the GT intensity images.

5.4 Ablation study

In this section, we analyze the proposed method via several controlled experiments. We started from a naive baseline by removing the reconstruction network and using unpolarized images as input: $P = (I, M)$, where I is the unpolarized light intensity. We call this experiment “*w/o Polar*”. Next, instead of unpolarized images, we used four linearly polarized images: $P = (I_0, I_{45}, I_{90}, I_{135}, M)$, and denote this experiment as “4 Pol”. Then, we added AoP and DoP to our input: $P = (I_0, I_{45}, I_{90}, I_{135}, \rho, \phi, M)$, and used “4 Pol + DoPAoP” to represent this experiment. After that, we replaced DoP and AoP images with the proposed *max* & *min* images: $P = (I_0, I_{45}, I_{90}, I_{135}, I_{\min}, I_{\max}, M)$, and denoted this experiment as “4 Pol + MaxMin”. Finally, we added the reconstruction network (Equation 15) back to obtain our full model. We compared these controlled experiments and reported MAE results in Table 1. From the table, we can observe that polarization cues significantly improved both shape and SSS estimation (See




Fig. 6: Visual comparison of SSS estimation results between our method and “*w/o Polar*”.

the comparison between “*w/o Polar*” and “4 Pol”). Although “4 Pol + DoPAoP” further improved the accuracy of shape estimation, the performance of scattering parameters was similar to “4 Pol”. The comparison between “4Pol + MaxMin” and “4 Pol + DoPAoP” demonstrated the efficiency of the proposed polarization representation in SSS estimation. Finally, the contribution of the proposed reconstruction network was validated by comparing our full model with the “4 Pol + MaxMin” experiment. In addition, we provide a qualitative result of SSS estimation between our method and “*w/o Polar*” in Figure 6 to demonstrate the contribution of polarization cues.

6 Conclusion and Limitation

In this paper, we proposed the first method for joint estimation of shape and SSS parameters using polarization cues. We also constructed the first large-scale synthetic dataset of polarized translucent objects. The efficiency of the proposed *max* & *min* polarization representation and reconstruction network were validated by several controlled experiments. In addition, we demonstrated that the proposed method outperforms the existing works from both SfP and inverse rendering realm, qualitatively and quantitatively.

However, we made several assumptions to reduce ambiguity. The first was a smooth surface to ensure the diffuse polarization only comes from SSS. However, if the object has a rough surface, the multiple reflections between microfacies also contribute to diffuse polarization. This can escalate ambiguity. Second, we assumed that photons travel far enough inside the object, which is inaccurate when the object has high transparency. Therefore, our method cannot be applied to nearly transparent objects. Solving these problems can be promising for future work.

Acknowledgement This paper is partially supported by JSPS KAKENHI 23H05490.

References

1. Aittala, M., Aila, T., Lehtinen, J.: Reflectance modeling by neural texture synthesis. *ACM Transactions on Graphics (ToG)* **35**(4), 1–13 (2016)
2. Aittala, M., Weyrich, T., Lehtinen, J., et al.: Two-shot svbrdf capture for stationary materials. *ACM Trans. Graph.* **34**(4), 110–1 (2015)
3. Ba, Y., Gilbert, A., Wang, F., Yang, J., Chen, R., Wang, Y., Yan, L., Shi, B., Kadambi, A.: Deep shape from polarization. In: *European Conference on Computer Vision*. pp. 554–571. Springer (2020)
4. Baek, S.H., Heide, F.: All-photon polarimetric time-of-flight imaging. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17876–17885 (2022)
5. Baek, S.H., Jeon, D.S., Tong, X., Kim, M.H.: Simultaneous acquisition of polarimetric svbrdf and normals. *ACM Trans. Graph.* **37**(6), 268–1 (2018)
6. Baek, S.H., Zeltner, T., Ku, H., Hwang, I., Tong, X., Jakob, W., Kim, M.H.: Image-based acquisition and modeling of polarimetric reflectance. *ACM Trans. Graph.* **39**(4), 139 (2020)
7. Boss, M., Jampani, V., Kim, K., Lensch, H., Kautz, J.: Two-shot spatially-varying brdf and shape estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3982–3991 (2020)
8. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: ShapeNet: An Information-Rich 3D Model Repository. Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago (2015)
9. Che, C., Luan, F., Zhao, S., Bala, K., Gkioulekas, I.: Towards learning-based inverse subsurface scattering. In: *2020 IEEE International Conference on Computational Photography (ICCP)*. pp. 1–12. IEEE (2020)
10. Chen, G., He, L., Guan, Y., Zhang, H.: Perspective phase angle model for polarimetric 3d reconstruction. In: *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*. pp. 398–414. Springer (2022)
11. Cui, Z., Gu, J., Shi, B., Tan, P., Kautz, J.: Polarimetric multi-view stereo. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1558–1567 (2017)
12. Cui, Z., Larsson, V., Pollefeys, M.: Polarimetric relative pose estimation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2671–2680 (2019)
13. Dave, A., Zhao, Y., Veeraraghavan, A.: Pandora: Polarization-aided neural decomposition of radiance. In: *European Conference on Computer Vision*. pp. 538–556. Springer (2022)
14. Deng, X., Luan, F., Walter, B., Bala, K., Marschner, S.: Reconstructing translucent objects using differentiable rendering. In: *ACM SIGGRAPH 2022 Conference Proceedings*. pp. 1–10 (2022)
15. Deschaintre, V., Aittala, M., Durand, F., Drettakis, G., Bousseau, A.: Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)* **37**(4), 1–15 (2018)
16. Deschaintre, V., Lin, Y., Ghosh, A.: Deep polarization imaging for 3d shape and svbrdf acquisition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15567–15576 (2021)

17. Dong, B., Moore, K.D., Zhang, W., Peers, P.: Scattering parameters and surface normals from homogeneous translucent materials using photometric stereo. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2291–2298 (2014)
18. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2002–2011 (2018)
19. Fukao, Y., Kawahara, R., Nobuhara, S., Nishino, K.: Polarimetric normal stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 682–690 (2021)
20. Gao, D., Li, Y., Ruhkamp, P., Skobleva, I., Wysocki, M., Jung, H., Wang, P., Guridi, A., Busam, B.: Polarimetric pose prediction. In: Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX. pp. 735–752. Springer (2022)
21. Gardner, M.A., Sunkavalli, K., Yumer, E., Shen, X., Gambaretto, E., Gagné, C., Lalonde, J.F.: Learning to predict indoor illumination from a single image. *ACM Transactions on Graphics (TOG)* **36**(6), 1–14 (2017)
22. Georgoulis, S., Rematas, K., Ritschel, T., Fritz, M., Tuytelaars, T., Van Gool, L.: What is around the camera? In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5170–5178 (2017)
23. Germer, T.A.: Evolution of transmitted depolarization in diffusely scattering media. *JOSA A* **37**(6), 980–987 (2020)
24. Gkioulekas, I., Levin, A., Zickler, T.: An evaluation of computational imaging techniques for heterogeneous inverse scattering. In: European Conference on Computer Vision. pp. 685–701. Springer (2016)
25. Gkioulekas, I., Zhao, S., Bala, K., Zickler, T., Levin, A.: Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics (TOG)* **32**(6), 1–13 (2013)
26. Henyey, L.G., Greenstein, J.L.: Diffuse radiation in the Galaxy. *The Astrophysical Journal* **93**, 70–83 (Jan 1941). <https://doi.org/10.1086/144246>
27. Hwang, I., Jeon, D.S., Muñoz, A., Gutierrez, D., Tong, X., Kim, M.H.: Sparse ellipsometry: portable acquisition of polarimetric svbrdf and shape with unstructured flash photography. *ACM Transactions on Graphics (TOG)* **41**(4), 1–14 (2022)
28. Ichikawa, T., Purri, M., Kawahara, R., Nobuhara, S., Dana, K., Nishino, K.: Shape from sky: Polarimetric normal recovery under the sky. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14832–14841 (2021)
29. Inoshita, C., Mukaigawa, Y., Matsushita, Y., Yagi, Y.: Surface normal deconvolution: Photometric stereo for optically thick translucent objects. In: European Conference on Computer Vision. pp. 346–359. Springer (2014)
30. Jakob, W., Speierer, S., Roussel, N., Vicini, D.: Dr.jit: A just-in-time compiler for differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* **41**(4) (Jul 2022). <https://doi.org/10.1145/3528223.3530099>
31. Kadambi, A., Taamazyan, V., Shi, B., Raskar, R.: Polarized 3d: High-quality depth sensing with polarization cues. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3370–3378 (2015)
32. Kalra, A., Taamazyan, V., Rao, S.K., Venkataraman, K., Raskar, R., Kadambi, A.: Deep polarization cues for transparent object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8602–8611 (2020)

33. Khungurn, P., Schroeder, D., Zhao, S., Bala, K., Marschner, S.: Matching real fabrics with micro-appearance models. *ACM Trans. Graph.* **35**(1), 1–1 (2015)
34. Kondo, Y., Ono, T., Sun, L., Hirasawa, Y., Murayama, J.: Accurate polarimetric brdf for real polarization scene rendering. In: *European Conference on Computer Vision*. pp. 220–236. Springer (2020)
35. Ku, H.J., Hat, H., Lee, J.H., Kang, D., Tompkin, J., Kim, M.H.: Differentiable appearance acquisition from a flash/no-flash rgb-d pair. In: *2022 IEEE International Conference on Computational Photography (ICCP)*. pp. 1–12. IEEE (2022)
36. Kurita, T., Kondo, Y., Sun, L., Moriuchi, Y.: Simultaneous acquisition of high quality rgb image and polarization information using a sparse polarization sensor. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 178–188 (2023)
37. Lei, C., Huang, X., Zhang, M., Yan, Q., Sun, W., Chen, Q.: Polarized reflection removal with perfect alignment in the wild. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1750–1758 (2020)
38. Lei, C., Qi, C., Xie, J., Fan, N., Koltun, V., Chen, Q.: Shape from polarization for complex scenes in the wild. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12632–12641 (2022)
39. Lensch, H., Kautz, J., Goesele, M., Heidrich, W., Seidel, H.P.: Image-based reconstruction of spatially varying materials. In: *Eurographics Workshop on Rendering Techniques*. pp. 103–114. Springer (2001)
40. Levis, A., Schechner, Y.Y., Aides, A., Davis, A.B.: Airborne three-dimensional cloud tomography. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3379–3387 (2015)
41. Li, C., Ngo, T.T., Nagahara, H.: Inverse rendering of translucent objects using physical and neural renderers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12510–12520 (2023)
42. Li, C., Ono, T., Uemori, T., Mihara, H., Gatto, A., Nagahara, H., Moriuchi, Y.: Neisf: Neural incident stokes field for geometry and material estimation. *arXiv preprint arXiv:2311.13187* (2023)
43. Li, R., Qiu, S., Zang, G., Heidrich, W.: Reflection separation via multi-bounce polarization state tracing. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII* 16. pp. 781–796. Springer (2020)
44. Li, R., Xian, K., Shen, C., Cao, Z., Lu, H., Hang, L.: Deep attention-based classification network for robust depth prediction. In: *Asian Conference on Computer Vision*. pp. 663–678. Springer (2018)
45. Li, X., Dong, Y., Peers, P., Tong, X.: Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (ToG)* **36**(4), 1–11 (2017)
46. Li, Z., Shafiei, M., Ramamoorthi, R., Sunkavalli, K., Chandraker, M.: Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2475–2484 (2020)
47. Li, Z., Sunkavalli, K., Chandraker, M.: Materials for masses: Svbrdf acquisition with a single mobile phone image. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 72–87 (2018)
48. Li, Z., Xu, Z., Ramamoorthi, R., Sunkavalli, K., Chandraker, M.: Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Transactions on Graphics (TOG)* **37**(6), 1–11 (2018)

49. Liang, Y., Wakaki, R., Nobuhara, S., Nishino, K.: Multimodal material segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19800–19808 (2022)
50. Lichy, D., Wu, J., Sengupta, S., Jacobs, D.W.: Shape and material capture at home. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6123–6133 (2021)
51. Lyu, Y., Cui, Z., Li, S., Pollefeys, M., Shi, B.: Reflection separation using a pair of unpolarized and polarized images. *Advances in neural information processing systems* **32** (2019)
52. Mahmoud, A.H., El-Melegy, M.T., Farag, A.A.: Direct method for shape recovery from polarization and shading. In: 2012 19th IEEE International Conference on Image Processing. pp. 1769–1772. IEEE (2012)
53. Mei, H., Dong, B., Dong, W., Yang, J., Baek, S.H., Heide, F., Peers, P., Wei, X., Yang, X.: Glass segmentation using intensity and spectral polarization cues. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12622–12631 (2022)
54. Meka, A., Maximov, M., Zollhoefer, M., Chatterjee, A., Seidel, H.P., Richardt, C., Theobalt, C.: Lime: Live intrinsic material estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6315–6324 (2018)
55. Nam, G., Lee, J.H., Gutierrez, D., Kim, M.H.: Practical svbrdf acquisition of 3d objects with unstructured flash photography. *ACM Transactions on Graphics (TOG)* **37**(6), 1–12 (2018)
56. Ngo Thanh, T., Nagahara, H., Taniguchi, R.i.: Shape and light directions from shading and polarization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2310–2318 (2015)
57. Nimier-David, M., Vicini, D., Zeltner, T., Jakob, W.: Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)* **38**(6), 1–17 (2019)
58. Ono, T., Kondo, Y., Sun, L., Kurita, T., Moriuchi, Y.: Degree-of-linear-polarization-based color constancy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19740–19749 (2022)
59. Qiao, Y., Dong, B., Jin, A., Fu, Y., Baek, S.H., Heide, F., Peers, P., Wei, X., Yang, X.: Multi-view spectral polarization propagation for video glass segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23218–23228 (2023)
60. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* (2020)
61. Riviere, J., Reshetouski, I., Filipi, L., Ghosh, A.: Polarization imaging reflectometry in the wild. *ACM Transactions on Graphics (TOG)* **36**(6), 1–14 (2017)
62. Sang, S., Chandraker, M.: Single-shot neural relighting and svbrdf estimation. In: European Conference on Computer Vision. pp. 85–101. Springer (2020)
63. Sengupta, S., Gu, J., Kim, K., Liu, G., Jacobs, D.W., Kautz, J.: Neural inverse rendering of an indoor scene from a single image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8598–8607 (2019)
64. Shao, M., Xia, C., Yang, Z., Huang, J., Wang, X.: Transparent shape from a single view polarization image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9277–9286 (2023)
65. Smith, W.A., Ramamoorthi, R., Tozza, S.: Height-from-polarisation with unknown lighting or albedo. *IEEE transactions on pattern analysis and machine intelligence* **41**(12), 2875–2888 (2018)

66. Sturzl, W.: A lightweight single-camera polarization compass with covariance estimation. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5353–5361 (2017)
67. Tian, C., Pan, W., Wang, Z., Mao, M., Zhang, G., Bao, H., Tan, P., Cui, Z.: Dps-net: Deep polarimetric stereo depth estimation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3569–3579 (2023)
68. Wang, Z., Phillion, J., Fidler, S., Kautz, J.: Learning indoor inverse rendering with 3d spatially-varying lighting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 12538–12547 (October 2021)
69. Wu, S., Makadia, A., Wu, J., Snavely, N., Tucker, R., Kanazawa, A.: De-rendering the world’s revolutionary artefacts. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6338–6347 (2021)
70. Yang, J., Xiao, S.: An inverse rendering approach for heterogeneous translucent materials. In: *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry-Volume 1*. pp. 79–88 (2016)
71. Zhao, J., Monno, Y., Okutomi, M.: Polarimetric multi-view inverse rendering. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)
72. Zhao, S., Wu, L., Durand, F., Ramamoorthi, R.: Downsampling scattering parameters for rendering anisotropic media. *ACM Transactions on Graphics (TOG)* **35**(6), 1–11 (2016)
73. Zheng, Q., Singh, G., Seidel, H.P.: Neural relightable participating media rendering. *Advances in Neural Information Processing Systems* **34**, 15203–15215 (2021)
74. Zhu, D., Smith, W.A.: Depth from a polarisation+ rgb stereo pair. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7586–7595 (2019)
75. Zhu, R., Li, Z., Matai, J., Porikli, F., Chandraker, M.: Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2822–2831 (2022)
76. Zhu, Y., Garigipati, P., Peers, P., Debevec, P., Ghosh, A.: Estimating diffusion parameters from polarized spherical-gradient illumination. *IEEE Computer Graphics and Applications* **33**(3), 34–43 (2013)
77. Zou, S., Zuo, X., Qian, Y., Wang, S., Xu, C., Gong, M., Cheng, L.: 3d human shape reconstruction from a polarization image. In: *European Conference on Computer Vision*. pp. 351–368. Springer (2020)
78. Zou, S., Zuo, X., Wang, S., Qian, Y., Guo, C., Cheng, L.: Human pose and shape estimation from single polarization images. *IEEE Transactions on Multimedia* (2022)