

Photometric Stereo in a Scattering Medium

Zak Murez, Tali Treibitz, *Member, IEEE*, Ravi Ramamoorthi, *Senior Member, IEEE*,
and David Kriegman, *Fellow, IEEE*

Abstract—Photometric stereo is widely used for 3D reconstruction. However, its use in scattering media such as water, biological tissue and fog has been limited until now, because of forward scattered light from both the source and object, as well as light scattered back from the medium (backscatter). Here we make three contributions to address the key modes of light propagation, under the common single scattering assumption for dilute media. First, we show through extensive simulations that single-scattered light from a source can be approximated by a point light source with a single direction. This alleviates the need to handle light source blur explicitly. Next, we model the blur due to scattering of light from the object. We measure the object point-spread function and introduce a simple deconvolution method. Finally, we show how imaging fluorescence emission where available, eliminates the backscatter component and increases the signal-to-noise ratio. Experimental results in a water tank, with different concentrations of scattering media added, show that deconvolution produces higher-quality 3D reconstructions than previous techniques, and that when combined with fluorescence, can produce results similar to that in clear water even for highly turbid media.

Index Terms—Photometric Stereo, Scattering Medium, Fluorescence

1 INTRODUCTION

OBTAINING 3D information about an object submersed in fog, haze, water, or biological tissue is difficult because of scattering [1], [2], [3]. In this paper, we focus on photometric stereo, which estimates surface normals from intensity changes under varying illumination. In air, photometric stereo produces high-quality geometry, even in textureless regions with small details, and is a widely used 3D reconstruction method.

In a scattering medium, however, light propagation is affected by scattering which degrades the performance of photometric algorithms unless accounted for. Distance dependent attenuation caused by the medium has been dealt with in the past [4]. Here, our contributions lie in handling three scattering effects (Fig. 1), based on a single scatter model [5]: 1) light traveling *from the source to the object* is blurred due to forward scattering; 2) light traveling *from the object to the camera* is blurred due to forward scattering; 3) light traveling *from the source is scattered back towards the camera* without hitting the object. This is known as backscatter and is an additive component that veils the object. All these effects are distance dependent and thus depend on the object 3D surface: the property we aim to reconstruct. To handle this we introduce the *small surface variations* approximation for the object (Sec. 3), that assumes surface changes are small relative to the distance from the object (that is assumed to be known). This assumption removes the dependence on the unknown surface heights Z , but unlike the common distant light/camera approximations, it

still allows for dependencies on spatial locations X and Y . One important consequence of this is the ability to model anisotropic light sources, which is not possible for distant lights.

Forward scatter was previously compensated for iteratively for both pathways (light to object and object to camera) simultaneously [6]. We analyze the paths separately. The resulting algorithm is simpler, requires fewer images and yields better results. First, consider the blurring of light traveling *to the object* from the source (Fig. 1b). The photometric stereo formulation assumes a point light source, illuminating from a single direction. However, if the source is scattered by the medium, this no longer holds: the point light source is spread, and the direction of light rays incident on the object changes. Nonetheless, for a Lambertian surface, illuminated from a variety of directions, we still get a linear equation between the image intensities and the surface normals [7]. Here, we show through simulations in a large variety of single scattering media, that a forward-scattered light source illuminating a Lambertian surface can be well approximated by a non-blurred light source in an effective purely absorbing medium (Sec. 6). This allows for much easier calibration in practice.

Next, we observe that the blur caused by scattering *from the object to the camera* (Fig. 1c) significantly affects the shape of the surface reconstructed by photometric stereo. This important effect has been neglected in many previous works. In general, the point-spread function (PSF) for an object is spatially varying and dependent on the unknown scene depths. However, we demonstrate that a spatially invariant approximation can still achieve good results, when calibrated for the desired medium and approximate object distance. Although this means we must capture an additional calibration image for each medium and working distance, we do not believe this will be too cumbersome based on previous experience in the field. We estimate the PSF and use it to deconvolve the images after backscatter

• Z. Murez, R. Ramamoorthi, D. Kriegman are with the Department of Computer Science and Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0404.

E-mail: {zmurez,ravir,kriegman}@cs.ucsd.edu.

• T. Treibitz is with the Department for Marine Technologies, Charney School of Marine Sciences, University of Haifa, Haifa, Israel
E-mail: ttreibitz@univ.haifa.ac.il




Fig. 1. A perspective camera is imaging an object point at X , with a normal N , illuminated by a point light source at S . The object is in a scattering medium, and thus light may be scattered in the three ways shown, detailed in Sec. 3.

has been removed (Sec. 7). These corrected images are used as input to a linear photometric stereo algorithm to recover the surface normals, which are then integrated. This results in much higher quality 3D surfaces across varying turbidity levels.

Finally, consider the backscatter component (Fig. 1a). In a previous work (Tsiotsios et al. [8]) backscatter was calibrated and subtracted from the input images. This is similar in spirit to, and can in fact be used in conjunction with, ambient light subtraction, however underwater light is rapidly attenuated and thus ambient light is often minimal. However, when backscatter is strong relative to the object signal, subtracting it after image formation leads to lack of dynamic range and lower signal-to-noise ratios (SNR) [9] that significantly degrades deblurring and reconstruction. Here we show that if the object fluoresces, this can be leveraged to optically remove the backscatter prior to image formation (Sec. 8). Fluorescence is the re-emission of photons in wavelengths longer than the excitation light [10], and therefore the backscatter can be eliminated by optically blocking the excitation wavelengths and imaging only the fluorescence emission. This improves SNR, especially in high turbidity. This approach is feasible as many natural underwater objects such as corals and algae fluoresce naturally.

We demonstrate our method experimentally in a water tank (Sec. 10) with varying turbidity levels. Deblurring can be used separately or combined with fluorescence imaging to significantly improve the quality of photometric stereo reconstructions.

2 PREVIOUS WORK

The traditional setup for photometric stereo assumes a Lambertian surface, orthographic projection, distant light sources, and a non-participating medium [11]. However, underwater light is exponentially attenuated with distance, and thus the camera and lights must be placed close to the scene for proper illumination. This means that the orthographic camera model, and the distant light assumptions, are no longer valid. In addition, attenuation and scattering by the medium need to be accounted for. These effects were partially considered in previous works.

Near-Field Effects and Exponential Attenuation: Photometric stereo in air was solved with perspective cameras [12], [13], nearby light sources [14], or both [15], [16]. Kolagani et al. [4] uses a perspective camera, nearby light sources and includes exponential attenuation of light in a

medium. Their formulation leads to nonlinear solutions for the normals and heights. We handle these near field effects but linearize the problem (Sec. 3).

Photometric Stereo with Backscatter: Narasimhan et al. [17] handles backscatter and attenuation, with the assumption of distant light sources and an orthographic camera. Tsiotsios et al. [8] extends this to nearby point sources and assumes the backscatter saturates close to the camera and thus does not depend on the unknown surface height. Then, it can be calibrated and subtracted from the images. We use the method in [8] in one of our variants.

Backscatter Removal: Backscatter was previously removed for visibility enhancement, by structured light [18], range-gating [2], or using polarizers [19]. Nevertheless, these methods do not necessarily preserve photometric information. It is sometimes possible to reduce backscatter by increasing the camera light source separation [18], [20], but this often leads to more shadowed regions, creating problems for photometric stereo.

Fluorescence Imaging: Removing scatter using fluorescence is used in microscopy [21], where many objects of interest are artificially dyed to fluoresce. Hullin et al. [22] imaged objects immersed in a fluorescent liquid to reconstruct their 3D structure. It was recently shown that the fluorescence emission yields photometric stereo reconstructions [23], [24] in air that are superior to reflectance images as the fluorescence emission behaves like a Lambertian surface due to its isotropic emission.

Deblurring Forward Scatter: Zhang et al. [25] and Negahdaripour et al. [6] handle blur caused by forward scatter using the PSF derived in [20], [26]. Their PSF depends on the unknown distances, as well as three empirical parameters, and affects both the path from the light source to the object and from the object to the camera. They iteratively deconvolve and update the depths until a good result is achieved. Trucco et al. [27] simplify the PSF of [20], [26] to only depend on two parameters while assuming the depth is known. Our PSF is nonparametric, independent of the unknown depths and only affects the path to the camera, which allows for a direct solution without iteration. While we look at a Lambertian surface in a scattering medium, Inoshita et al. [28] and Dong et al. [29] consider the problem of photometric stereo in air on a surface that exhibits subsurface scattering, which blurs the radiance across the surface. They deconvolve the images to improve the quality of the normals recovered using linear photometric stereo. Tanaka et al. [30] also model forward scatter blur as a depth dependent PSF and combine it with multi (spatial) frequency illumination to recover the

Our Algorithm**Input:**

3 or more images: L^i
 source positions: \mathbf{S}^i
 mean depth: \bar{Z}
 backscatter images: L_b^i
 scattering parameters: PSF and $\tilde{\sigma}$ (Sec. 9.3)

Output:

normals: N
 surface heights: Z
 1: **if** reflectance images **then**
 2: subtract backscatter: $L^i - L_b^i$ (Sec. 5)
 3: **end if**
 4: deblur images: $L_o^i \leftarrow h^{-1} * (L^i - L_b^i)$ (Eq. 19)
 5: solve linear PS: $N_j \leftarrow [\tilde{L}_j^{eq}]^{-1} [L_{o_j}]$ (Eq. 10,12)
 6: integrate normals: $Z \leftarrow \int N$

Fig. 2. Our Algorithm. Steps 2 and 4 are applied to each image i independently. Step 5 is applied to each pixel j independently with the data from each image i stacked into a matrix.

appearance of a small number of inner slices of a translucent material.

An earlier version of this paper was originally presented in [31]. Here we have provided additional analysis of the effective point source approximation, a more complete description of our experimental procedure and more detailed results. We have also clarified details of our algorithm and added a summary in the form of pseudocode.

3 OVERVIEW AND ASSUMPTIONS

In this section we introduce the image formation model, considering each of the modes of light propagation in a single scattering medium, as shown in Fig. 1. We derive expressions for each component in the following sections.

Consider a perspective camera placed at the origin, with the image (x, y) coordinates parallel to the world's (X, Y) axes, and the Z -axis aligned with the camera's optical axis. Let the point $\mathbf{X} = (X, Y, Z)$ be the point on the object's surface along the line of sight of pixel $\mathbf{x} = (x, y)$. Let \mathbf{S} be the world coordinates of a point light source, and define $\mathbf{D}(\mathbf{X}) = \mathbf{S} - \mathbf{X}$ as the vector from the object to the source.

We assume a single scattering medium which allows us to express the radiance L_o reflected by a surface point as the sum of two terms:

$$L_o(\mathbf{x}) = L_d(\mathbf{x}) + L_s(\mathbf{x}) \quad (1)$$

where L_d is the direct radiance from the source (Sec. 4), and L_s is the radiance from the source which is scattered from other directions onto \mathbf{X} (Sec. 6 and Fig. 1b).

Next, we express the radiance arriving at the camera as the sum of three terms:

$$L(\mathbf{x}) = L_o(\mathbf{x})e^{-\sigma\|\mathbf{X}\|} + L_b(\mathbf{x}) + L_c(\mathbf{x}) \quad (2)$$

where L_o is the light reflected by the surface point \mathbf{X} which arrives at the camera without undergoing scattering. Note that it is attenuated by $e^{-\sigma\|\mathbf{X}\|}$ where σ is the extinction coefficient. L_b is composed of rays of light emitted by the source that are scattered into \mathbf{x} 's line of sight before hitting

the surface (Sec. 5 and Fig. 1a). This term is known as backscatter. Finally, L_c is composed of rays of light reflected by other points on the surface that are scattered into pixel \mathbf{x} 's line of sight (Sec. 7 and Fig. 1c).

In order to write analytic expressions for these terms and derive a simple solution we make **two assumptions**. First, the surface is Lambertian with a spatially varying albedo $\rho(\mathbf{X})$. Second, we assume that surface variations in height are small compared to object distance from the camera. We call this the *small surface variations approximation* and note that it is weaker than the common distant light sources and orthographic projection approximations. Let \bar{Z} be the average Z coordinate of the surface (assumed to be known). Then, the approximation claims that for every point on the surface: $|Z(\mathbf{X}) - \bar{Z}| \ll \bar{Z}, \forall \mathbf{X}$.

The approximation results in a weak perspective such that the projection \mathbf{x} of \mathbf{X} in the image plane is given by

$$\mathbf{x} = \left(f \frac{X}{\bar{Z}}, f \frac{Y}{\bar{Z}} \right)^t ; \quad \mathbf{X} = \left(\frac{\bar{Z}}{f} x, \frac{\bar{Z}}{f} y, \bar{Z} \right)^t, \quad (3)$$

where f is the known focal length. Note that for a given pixel, since we know its (x, y) coordinates and the average object distance \bar{Z} , the world coordinates \mathbf{X} are known. Specifically, $\mathbf{D}(\mathbf{X})$ is independent of the unknown object height Z but still depends on X and Y , whereas in the distant light sources approximation $\mathbf{D}(\mathbf{X})$ is a constant.

Outline of Our Method

Given an input image L we eliminate the backscatter L_b by one of two methods. The first follows [8]: backscatter from each light source is measured by imaging it with no objects in the scene, and then the measured backscatter is subtracted from the input images. In the second, backscatter is optically eliminated using fluorescence as we explain in Sec. 8. Once backscatter is removed, the resulting images are deblurred, using a calibrated PSF, to recover L_o (Sec. 7, Eq. 19). Next we write L_o as a linear equation between the unknown surface normals, albedo and an equivalent light source (Sec. 6, Eq. 10), which we approximate as an effective point source in a purely absorbing medium with effective extinction coefficient (Sec. 6, Eq. 12). With a minimum of 3 images under distinct light locations the normals can be solved for, as in conventional photometric stereo. The normals are then integrated to recover a smooth surface. This algorithm is summarized in Fig. 2.

4 DIRECT RADIANCE

First, consider the direct reflected radiance from a Lambertian surface [11]:

$$L_d(\mathbf{x}) = I(\mathbf{X}) \frac{\rho(\mathbf{X})}{\pi} \hat{\mathbf{D}}(\mathbf{X}) \cdot \hat{\mathbf{N}}, \quad (4)$$

where $\hat{\mathbf{N}}$ is the unit surface normal and $\hat{\mathbf{D}}$ is the normalized source-to-object vector. The radiance on the object surface $I(\mathbf{X})$ depends on the radiant intensity I_0 of the source in direction¹ $(-\hat{\mathbf{D}})$:

$$I(\mathbf{X}) = \left(I_0(-\hat{\mathbf{D}}(\mathbf{X})) e^{-\sigma\|\mathbf{D}(\mathbf{X})\|} \right) / \|\mathbf{D}(\mathbf{X})\|^2. \quad (5)$$

1. the direction is negative as we consider outgoing rays from the source.

Eq. 5 accounts for nearby angularly-varying sources, exponential attenuation along the optical path length with extinction coefficient σ , and inverse-square distance falloff.

Can Distance-Dependent Falloff Be Neglected?

We have introduced a near-field source, but often, in photometric stereo, the distant light source assumption is used, as it simplifies the mathematical development, computation of shape, and calibration of implemented systems because the light source direction can be treated as a constant, and the incident irradiance does not depend upon depth allowing it to also be treated as a constant. In this section, we explore whether similar simplifications are possible for photometric stereo in a medium.

In a medium, the incident irradiance falls off as function of distance due to the product of two factors; free space falloff and medium attenuation. Consider the situation depicted in Fig. 3a for which two points P and P_Δ are illuminated by a light source that is respectively at distances d and $d + \Delta$ from the points. The ratio Ψ of the irradiance at P and P_Δ can be expressed as $\Psi = E_{P_\Delta}/E_P = \Psi_{\text{freespace}}\Psi_{\text{medium}}$. The irradiance of a point light source propagating in a medium at distance d , falls off by $1/d^2$, and is attenuated by the medium by $e^{-\beta d}$. Thus,

$$\Psi_{\text{freespace}} = d^2/(d + \Delta)^2, \quad \Psi_{\text{medium}} = e^{-\beta \Delta}. \quad (6)$$

Interestingly, $\Psi_{\text{freespace}}$ depends on both the absolute distance d to the light source as well as Δ , while Ψ_{medium} is independent of d , and depends only on the path difference Δ and on the attenuation coefficient β . To get an idea of object dimensions where the variation in incident irradiance is small and might be treated as constant, let us consider an example where the path difference Δ yields $\Psi = 0.9$ (i.e., a 10% difference in the incident irradiance at P and P_Δ). Figure 3b shows a plot of Δ vs d where $\Psi_{\text{freespace}} = 0.9$, and we see that the object's size can increase linearly with distance. Figure 3c shows a plot of Δ vs β the attenuation parameter where $\Psi_{\text{medium}} = 0.9$, and we see that as the medium becomes murkier the size decreases and is independent of distance. In other words, the opportunity to neglect distance-dependent falloff of lighting in a medium depends upon the clarity of the medium, even when the distance is large enough to allow the freespace falloff to be ignored. Due to this we chose to use the *small surface variations approximation* instead of the distant source approximation.

5 BACKSCATTER

Light is scattered as it travels through a medium. The fraction of light scattered to each direction is determined by the phase function $P(\alpha)$, where $\alpha \in [0, 2\pi]$ is the angle between the original ray direction and scattered ray, and β is the scattering coefficient.

Light which is scattered directly into the camera by the medium without reaching the object is termed *backscatter* and is given by [5], [32] (Fig. 1a):

$$L_b(\mathbf{x}) = \beta \int_0^{\|\mathbf{x}\|} I(r\hat{\mathbf{X}})P(\alpha)e^{-\sigma r} dr \quad (7)$$

The integration variable r is the distance from the camera to the imaged object point \mathbf{X} along the line of sight (LOS),

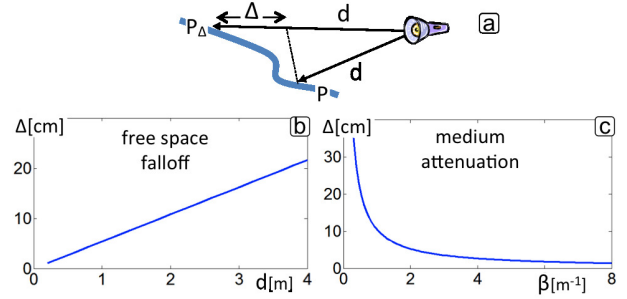


Fig. 3. (a) Light source can be considered distant if the irradiant light intensity across it is uniform. (b) Path length differences yield 10% intensity difference of point light source intensity due to free space falloff, as a function of d . (c) Path length differences along an object that yield 10% difference in medium attenuation across it, as a function of β . For $\beta > 1m^{-1}$, which represents fairly clear water, path lengths greater than 10cm already result in noticeable intensity changes, ruling out the distant light source assumption.

that is a unit direction $\hat{\mathbf{X}}$. The scattering angle α is given by $\cos(\alpha) = \hat{\mathbf{D}} \cdot \hat{\mathbf{X}}$ (recall that $\hat{\mathbf{D}}$ is the direction to the light) and $I(r\hat{\mathbf{X}})$ is the direct radiance of the source at point $r\hat{\mathbf{X}}$ as defined in Eq. 5.

Note that for the *small surface variations approximation*, \mathbf{X} and hence the limits of the integral are known for a given pixel. Therefore, the backscatter does not depend on the unknown height of the object and is a (different) constant for each pixel, similar to Tsitsios et al. [8].

Instead of analytically computing L_b , we found that it was easier and more accurate to directly measure it using the calibration method of [8]: for each light an image is captured with no object in the field-of-view.

6 SINGLE SCATTERED SOURCE RADIANCE

Because of the medium, light rays that are not originally pointed at an object point may be scattered and reach it from the entire hemisphere of directions Ω (Fig. 1b), termed forward scattered radiance

$$L_s(\mathbf{x}) = \frac{\rho(\mathbf{X})}{\pi} \int_{\omega \in \Omega} L_i(\omega)(\omega \cdot \hat{\mathbf{N}}) d\omega. \quad (8)$$

where $L_i(\omega)$ is the total radiance scattered into the direction ω and is given by

$$L_i(\omega) = \beta \int_{t=0}^{\infty} I(\mathbf{X} + t\omega)P(\alpha)e^{-\sigma t} dt, \quad (9)$$

where t is the distance from the object, and the angle α is given by $\cos(\alpha) = \hat{\mathbf{D}}(\mathbf{X} + t\omega) \cdot \omega$. Note that $\hat{\mathbf{D}}$ is the direction of the integration point to the light.

Substituting Eqs. 4,8 into Eq. 1 and rearranging yields:

$$L_o(\mathbf{X}) = L_d(\mathbf{X}) + L_s(\mathbf{X}) = \frac{\rho(\mathbf{X})}{\pi} L^{\text{eq}}(\mathbf{X}) \cdot \hat{\mathbf{N}}, \quad (10)$$

where

$$L^{\text{eq}}(\mathbf{X}) = I(\mathbf{X})\hat{\mathbf{D}}(\mathbf{X}) + \int_{\omega \in \Omega} L_i(\omega)\omega d\omega. \quad (11)$$

Here, the direct light as well as the integrated scattered contributions can be thought of as an equivalent distant source. However, this equivalent source may be different (in

direction and magnitude) for each surface point, and thus is not a true distant source. Furthermore, the integration domain Ω in Eq. 11 depends on $\hat{\mathbf{N}}$, preventing Eq. 10 from giving us a simple linear equation for the unknown normals and albedo.

We next show through simulations, that for a wide variety of media, $L^{\text{eq}}(\mathbf{X})$ can be approximated as an effective point source in a purely absorbing medium with effective extinction coefficient. This eliminates the non-linearity in Eq. 10 allowing for a linear solution for the normals given 3 or more images.

Effective Point Source Simulations

We approximate $L^{\text{eq}}(\mathbf{X})$ as:

$$\tilde{L}^{\text{eq}}(\mathbf{X}) \approx \frac{\kappa I_0 (-\hat{\mathbf{D}}(\mathbf{X})) e^{-\tilde{\sigma} \|\mathbf{D}(\mathbf{X})\|}}{\|\mathbf{D}(\mathbf{X})\|^2} \hat{\mathbf{D}}(\mathbf{X}), \quad (12)$$

which has the same form as Eq. 5, but the extinction coefficient is replaced by the effective extinction coefficient $\tilde{\sigma}$, and the intensity is scaled by κ . The effective source has the same position \mathbf{S} and intensity distribution I_0 as the real source. Note that κ is a global brightness scale, which is the same for all the lights, and thus does not need to be explicitly calibrated, as it cancels out in the normal estimation.

The intuition for why the source direction is unchanged is visualized in Fig. 4a. Although light is arriving from the entire hemisphere of directions, the vector sum of most of these directions lies in the original direction due to symmetry. Only the area of asymmetric scattering does not have symmetrical rays since the symmetrical rays lie below the visible hemisphere (in attached shadow) for the surface point. Although these asymmetric rays could potentially shift the equivalent direction, their contribution is often small for two reasons. First, these rays correspond to larger scattering angles, which are often much weaker than for rays with smaller scattering angles. Second, these paths are on average longer than for paths with smaller scattering angles and thus are more attenuated. Guided by this intuition we formulated the effective source approximation, and verified it's accuracy through extensive simulations in a wide variety of media.

For our simulations we used an isotropic point source at a distance d from a Lambertian surface patch with an angle ϕ between the surface normal and light direction. Note that the parametrization of a surface patch by d and ϕ fully parametrize the space of possible surface patches. For the scattering function, we used the common Henyey-Greenstein phase function [33], which can represent a large space of scattering functions by tuning a single parameter $g \in [-1, 1]$. In water, g is usually between 0.7 – 0.9 [32].

We compute $L_o(d, \phi)$ for $d \in [200, 600]\text{mm}$, $\phi \in [0, \pi]$, for a variety of media given by $\beta \in [0, 0.005]\text{mm}^{-1}$ and $g \in [0, 0.9]$. Note that we choose I_0 such that $L_o(200, 0)$ is normalized to 1. To reduce the number of parameters we set $\sigma = \beta$, which does not influence the analysis². For

2. In general $\beta \leq \sigma$, but since β purely scales L_s , a smaller value of beta would make L_o closer to L_d and thus \tilde{L}_0 would be an even better fit than we calculated.

each parameter pair g, β we compute the approximation parameters κ and $\tilde{\sigma}$ by minimizing:

$$\min_{\kappa, \tilde{\sigma}} \sum_{d, \phi} |L_o(d, \phi) - \tilde{L}_o(d, \phi)|^2. \quad (13)$$

The error in the approximation is then given by the residuals $\text{RE}(d, \phi) = |L_o(d, \phi) - \tilde{L}_o(d, \phi)|$.

The residuals for $g = 0.8, \beta = 0.0026$ (common in our setup) are plotted in Fig. 4b. We can see that the difference between the approximation and true values are small for all values of d and ϕ . Note that the error is largest near $\phi = 90^\circ$, where the area of asymmetric scattering is largest.

For each $\beta \in [0, .005]\text{mm}^{-1}$ and $g \in [0, .9]$ we compute the mean residual (MRE) over d, ϕ and plot it in Figure 4c. We see that the MRE is less than 2% across all medium conditions tested justifying the approximation. Further, we see that the MRE increases slowly with scattering coefficient β , and more rapidly with phase parameter g . As such, in water, which is mostly forward scattering (g between 0.7 – 0.9), our approximation is very accurate, even for highly turbid media. On the other hand, if scattering is more isotropic (g close to zero), then the approximation might not be valid for large β .

7 SINGLE SCATTER OBJECT BLUR

Similar to the light source blur, radiance from the object is also blurred while it propagates to the camera (Fig. 1c). As we demonstrate, this effect deteriorates the performance of photometric stereo, although it has been neglected in previous works [5].

The contribution of object blur to the pixel intensity is computed by integrating light scattered into the LOS of \mathbf{X} from all other points on the surface:

$$L_c(\mathbf{x}) = \beta \int_{r=0}^{\|\mathbf{X}\|} \int_{\omega \in \Omega} L_o(\mathbf{X}') P(\alpha) e^{-\sigma(t+r)} d\omega dr. \quad (14)$$

Here r is the distance along the LOS, \mathbf{X}' is the object surface point intersected by the ray starting at point $r\mathbf{X}$ in direction ω . Its radiance is $L_o(\mathbf{X}')$ and its distance to the scatter point in the LOS is given by $t = \|r\hat{\mathbf{X}} - \mathbf{X}'\|$ with scattering angle $\cos \alpha = \omega \cdot (-\hat{\mathbf{X}})$.

We now show that L_o can be recovered from $L_o e^{-\sigma \|\mathbf{X}\|} + L_c$ by deconvolution with a constant PSF.

Deblurring Object Scatter

First we rewrite Eq. 14 to integrate over the area of the object surface $dA = d\omega \cdot t^2 / \cos \theta$ instead of solid angle $d\omega$, where t is the distance from \mathbf{X}' to the scattering event, and θ is the angle between the normal at \mathbf{X}' and the ray of light before scattering. Eq. 14 now becomes

$$L_c(\mathbf{x}) = \beta \int_{\mathbf{X}'} L_o(\mathbf{X}') \int_0^{\|\mathbf{X}\|} P(\alpha) e^{-\sigma(t+r)} \frac{\cos \theta}{t^2} dr dA. \quad (15)$$

Now we define the scattering kernel

$$K(\mathbf{X}, \mathbf{X}') = \delta(\mathbf{X} - \mathbf{X}') e^{-\sigma \|\mathbf{X}\|} + \beta \int_0^{\|\mathbf{X}\|} P(\alpha) e^{-\sigma(t+r)} \frac{\cos \theta}{t^2} dr, \quad (16)$$

where $\delta(\mathbf{X} - \mathbf{X}')$ is the Dirac delta function. Now,

$$L_o(\mathbf{x})e^{-\sigma\|\mathbf{x}\|} + L_c(\mathbf{x}) = \int_{\mathbf{X}'} K(\mathbf{X}, \mathbf{X}')L_o(\mathbf{X}') dA(\mathbf{X}') . \quad (17)$$

In general, the kernel K , depends on \mathbf{X} , \mathbf{X}' and the unknown normals $\hat{\mathbf{N}}'$. For an orthographic camera viewing a plane at constant depth, K is shift invariant and Eq. 17 can be written as a convolution with a PSF. Motivated by this, we found empirically that for a given \bar{Z} it is approximately shift invariant (and rotationally symmetric).

Denoting the PSF as h , we get

$$L_o(\mathbf{x})e^{-\sigma\|\mathbf{x}\|} + L_c(\mathbf{x}) \approx h * L_o . \quad (18)$$

We emphasize here that we have shown that under a single scattering model, the forward scatter from the object can be written as an integral transform with kernel K . This justifies approximating the forward scattering as a PSF which is not obvious in the form of Eq. 14.

We solve Eq. 18 for L_o by writing the image as a column vector and representing the convolution as a matrix operation

$$(L - L_b) = HL_o \quad (19)$$

where we have substituted the known backscatter compensated image $(L - L_b)$ for $L_o e^{-\sigma\|\mathbf{x}\|} + L_c$, and H is the matrix representation of h . Here H is a large nonsparse matrix and thus storing it in memory and directly inverting it is infeasible. Instead we solve the linear system of Eq. 19 using conjugate gradient descent. This requires only the matrix vector operation which can be computed as a convolution and implemented using a Fast Fourier Transform (FFT).

8 BACKSCATTER REMOVAL USING FLUORESCENCE

While we are able to subtract the backscatter component, it is an additive component that effectively reduces the dynamic range of the signal from the object, degrades the image quality and reduces SNR [9]. As such it is beneficial to optically remove it when imaging. Here, we use the observation that for fluorescence images taken with non-overlapping excitation and emission filters, there is no backscatter in the image (Fig. 5a). In fluorescence imaging, the signal of interest is composed of wavelengths that are longer than that of the illumination, and a barrier filter on the camera is used to block the reflected light. The backscatter is composed of light scattered by the medium *before* it reaches the object. Thus, the backscatter has the same spectral distribution as the light source, which is blocked by the barrier filter on the camera. This insight enables imaging without loss of dynamic range even in highly turbid media. Compared to a backscatter subtracted reflectance image, a fluorescence image has less noise (Fig. 5b,c). This difference becomes even more apparent after deconvolution (Fig. 5d,e).

In addition, in [23], [24] it was shown that the fluorescence emission acts as a Lambertian surface in photometric reconstructions. Thus, imaging fluorescence has an additional advantage as it relaxes the need for a Lambertian surface.

In the development of our algorithm we assumed a single set of medium parameters β , σ and $P(\alpha)$. However

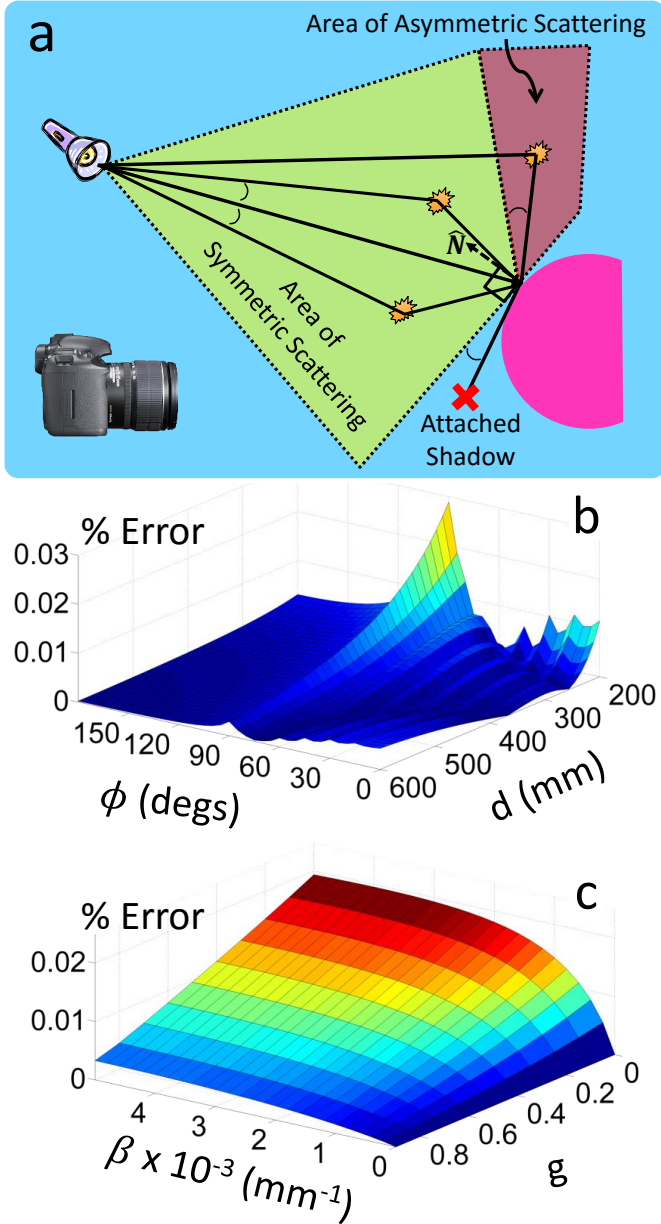


Fig. 4. a) Diagram of the intuition for the effective source approximation. Although light is arriving from the entire hemisphere of directions, the vector sum of most of these directions lies in the original direction due to symmetry. Only the area of asymmetric scattering does not have symmetrical rays since the symmetrical rays lie below the visible hemisphere (in attached shadow) for the surface point. The contribution from these rays are often small. b) The relative error between $L_o(d, \phi)$ and $\tilde{L}_o(d, \phi)$ for $g = 0.8$ and $\beta = 0.0026$. Note that the spike only reaches 3% and is located at $\phi = 90^\circ$ where $\tilde{L}_o = 0$ due to shadowing. ϕ near 90° and above is not usually relevant for photometric stereo. c) The mean relative error between L_o and \tilde{L}_o for $\beta \in [0, 0.005]\text{mm}^{-1}$ and $g \in [0, 0.9]$. The approximation errors are small over a wide variety of media.

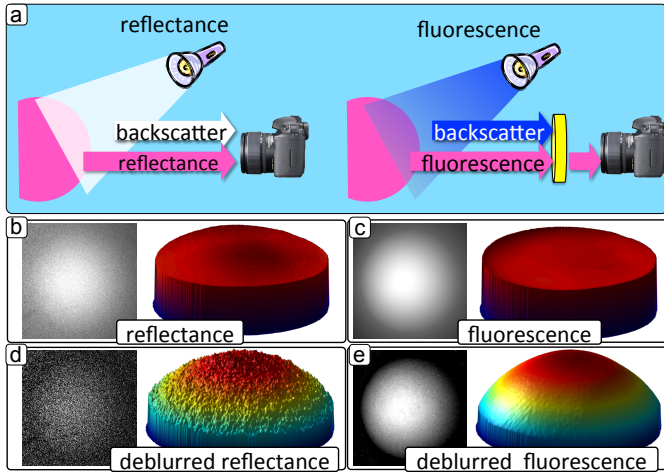


Fig. 5. a) Backscatter is caused by light that is scattered into the camera by a medium, before it reaches the object and has the same color as the illumination. Thus, the barrier filter used to block fluorescence excitation also blocks backscatter, while imaging the signal from the object. We use this property to remove backscatter in input images. b) Photometric stereo reconstruction of a fluorescent sphere using backscatter subtracted reflectance images. One of the input images is shown on the left, with visible noise and blur. Blur in the input images flattens the reconstruction. c) Looking at fluorescence images as an input, the backscatter is eliminated while maintaining a higher SNR. However the blur still flattens the reconstruction. d) Deblurring the backscatter subtracted images recovers the general shape but suffers from noise as seen by the spiky surface. e) Deblurring the fluorescence results in the correct shape with much less noise.

these quantities are in general wavelength dependent. In reflectance imaging, the wavelength of the light is the same on both pathways: light to object, and object to camera. However, in fluorescence imaging they are different. Nevertheless, the only parameters that require calibration in our solution are the effective extinction coefficient $\tilde{\sigma}$ and the PSF. The parameter $\tilde{\sigma}$ is estimated for the excitation wavelength and the PSF is estimated for the emission wavelength, and as such we do not need to calibrate any extra parameters in the case of fluorescence imaging.

9 IMPLEMENTATION

9.1 Experimental Setup

Our setup is shown in Fig. 6. We used a Canon 1D camera with a 28mm lens placed 2cm away from a 10 gallon glass aquarium. All sides except the front (where the camera looks in) were painted black to reduce reflection. In addition, a black panel was suspended just below the surface of the water to remove reflection from the air-water interface. The objects were placed at an average distance of 40cm from the front of the tank. For point illumination we used Cree XML - RGBW Star LEDs. The LEDs were water proofed by coating the electrical terminals with epoxy. Reflection images were taken under white illumination while fluorescent images were taken under blue illumination with a Tiffen #12 emission filter on the camera. We used tap water, and the turbidity was increased using a mixture of whole milk and grape juice (milk is nearly purely scattering, while grape juice is nearly purely absorbing and thus by mixing them




Fig. 6. [Left] Our experimental setup consists of a camera looking through a glass port into a tank. [Right] 8 LEDs are mounted inside the tank around the camera port illuminating the object placed at the back of the tank.

	Level 1	Level 2	Level 3	Level 4
Objects	Spherical Cap & Lobster			
Milk (ml)	1.25	2.50	3.75	5.00
juice (ml)	15.0	30.0	45.0	60.0
$\beta (\times 10^{-3} mm^{-1})$.602	1.20	1.81	2.41
$\sigma (\times 10^{-3} mm^{-1})$.642	1.28	1.93	2.57
Objects	Toy Gun & Mask			
Milk (ml)	1.25	2.50	3.75	5.00
juice (ml)	0	0	0	0
$\beta (\times 10^{-3} mm^{-1})$.602	1.20	1.81	2.41
$\sigma (\times 10^{-3} mm^{-1})$.602	1.20	1.81	2.41

Fig. 7. Tabulated values for the amount of milk and grape juice added in our experiments, and the associated scattering and extinction coefficients. The coefficients were computed using the data provided in [32].




Fig. 8. Cross-sections of the spherical cap reconstruction in turbid medium using various methods compared to ground truth. The clear water reconstruction resembles the ground truth. Only correcting for the backscatter (by subtraction or fluorescence) yields flattened results. Deblurring the backscatter subtracted images recovers the shape but is degraded by noise (the surface is jagged). Deblurring the fluorescence images produces the best results.

we can achieve a variety of scattering conditions [32]). The LEDs were mounted inside the tank on a square around the camera, four on the corners and four on the edges. Their positions were measured. Images were acquired in Raw mode which is linear and the normal integration was done using the method of [34].

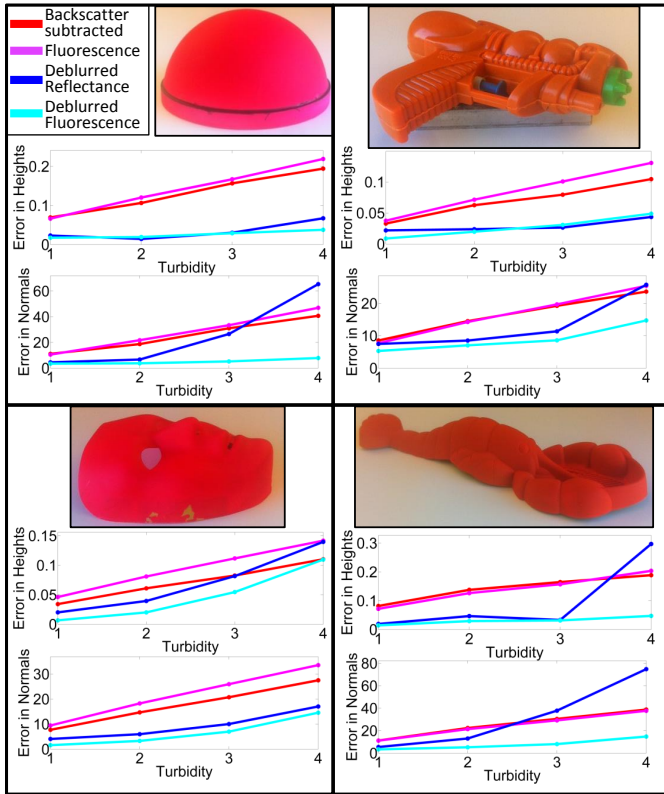


Fig. 9. Errors in the reconstructions of four objects as a function of turbidity, compared to clear water reconstruction. Top rows are average percent errors in heights and bottom rows are average angular errors in normals. Removing backscatter by either subtraction or using fluorescence performs similarly. Deblurring the backscatter compensated images significantly improves the reconstructions. In high turbidity where the backscatter is strong compared to the object signal deblurring the backscatter subtracted images degrades due to noise, while deblurring the fluorescence suffers less, as the fluorescence images have a higher SNR.

9.2 Geometric and Radiometric Calibration

Images of a checkerboard (in clear water) were used to calibrate the intrinsic camera parameters (implicitly accounting for refraction) [35]. The location of each light was measured using a ruler, and transformed to the camera reference frame. To calibrate each light's angular intensity distribution we imaged a matte painted (assumed to be Lambertian) plane at a known position under illumination from each light in clear water. Using Eq. 4, the known geometry, and $\sigma = 0$, we compute $I_0(-\mathbf{D})$, the angular dependence of the light source.

9.3 Calibration of Medium Parameters

The backscatter component is measured using the calibration method of [8]. For each light an image is captured with no object in the field-of-view and subsequently subtracted from future reflectance images. This is not used when imaging fluorescence.

Our method works independent of how the PSF is calibrated and thus a variety of methods could be used including that of Narasimhan et. al. [36]. Here we chose a procedure using a calibration target similar to [37] due to its ease of implementation. We use a matte painted

checkerboard which is imaged with its axis aligned to the image plane at the approximate depth of the objects we plan to reconstruct. As the PSF is rotationally symmetric its parameters are the values along a radius $[h_0, \dots, h_s]$, where h_0 is the center value and h_s is the value on the support radius s . The PSF and the effective extinction coefficient $\tilde{\sigma}$ are estimated by optimizing

$$\min_{\tilde{\sigma}} \min_{h_0 \dots h_s} \sum_{\mathbf{x}} \|h * L_o(\mathbf{x}, \tilde{\sigma}) - (L - L_b)\| \quad (20)$$

where L is the image of the checkerboard in the medium and L_b is the image of the backscatter. L_o is computed from Eq. 10 using the calibrated lights, known geometry, and registering the checkerboard albedo, measured in clear water to the image in turbid water. The inner optimization is an overdetermined linear system holding $\tilde{\sigma}$ fixed. We sweep over the values of $\tilde{\sigma}$ and choose the one with the minimum error. Note that the PSF is not normalized due to loss of energy (attenuation) from the object to the camera.

10 RESULTS

We imaged four objects: a spherical cap (Fig. 10), a plastic toy lobster (Fig. 11), a plastic toy squirt gun (Fig. 12), and a fluorescent painted mask (Fig. 13) in clear water as well as four increasing turbidities. Each turbidity level corresponded to adding 1.25ml of milk to the 10 gallon tank. For the spherical cap and the lobster, we also added 15ml of grape juice per turbidity level to increase absorption. In this case, since attenuation is exponential with distance, the signal from the object, which travels further, is relatively weaker than the backscatter which comes mostly from shorter paths. This exacerbates the loss of signal-to-noise ratio in backscatter subtracted images. To get an idea of the true scattering parameters of our various media we use the data provided in [32]. Tabulated values are shown in Fig. 7.

We employ two error metrics to evaluate the quality of our reconstructions: The mean absolute difference in heights ($\text{Err } Z = \text{mean}(Z - Z_{gt})$) and the mean angular error in the normals ($\text{Err } N = \text{mean}(\text{acos}(N \cdot N_{gt}))$), where Z_{gt} and N_{gt} are the ground truth heights and normals. Note that during integration, random noise in the normals cancels out locally, resulting in reconstructions with the correct overall shape, but with rough surfaces. As such $\text{Err } Z$ captures systematic errors that affect the overall shape, but is less sensitive to noise in the normals.

We see that the reconstructed spherical cap in clear water nearly perfectly matches the ground truth (Fig. 8) with an $\text{Err } Z$ of 1.4% and $\text{Err } N$ of 3°. This justifies our use of clear water reconstructions as ground truth for the other objects where true ground truth is not available.

The quality of results as a function of turbidity level is demonstrated in Fig. 9. The plots show how $\text{Err } Z$ and $\text{Err } N$ increase for each method as the turbidity increases, where the lowest error is achieved using the deblurred fluorescence images. In the highest turbidity level, the deblurred reflectance image often performs worse than all other methods, as the deblurring degrades with noise.

Figures 10,11,12,13 depict an input image and the resulting reconstruction for the spherical cap, the toy lobster, the toy squirt gun, and the mask respectively. In each figure, the

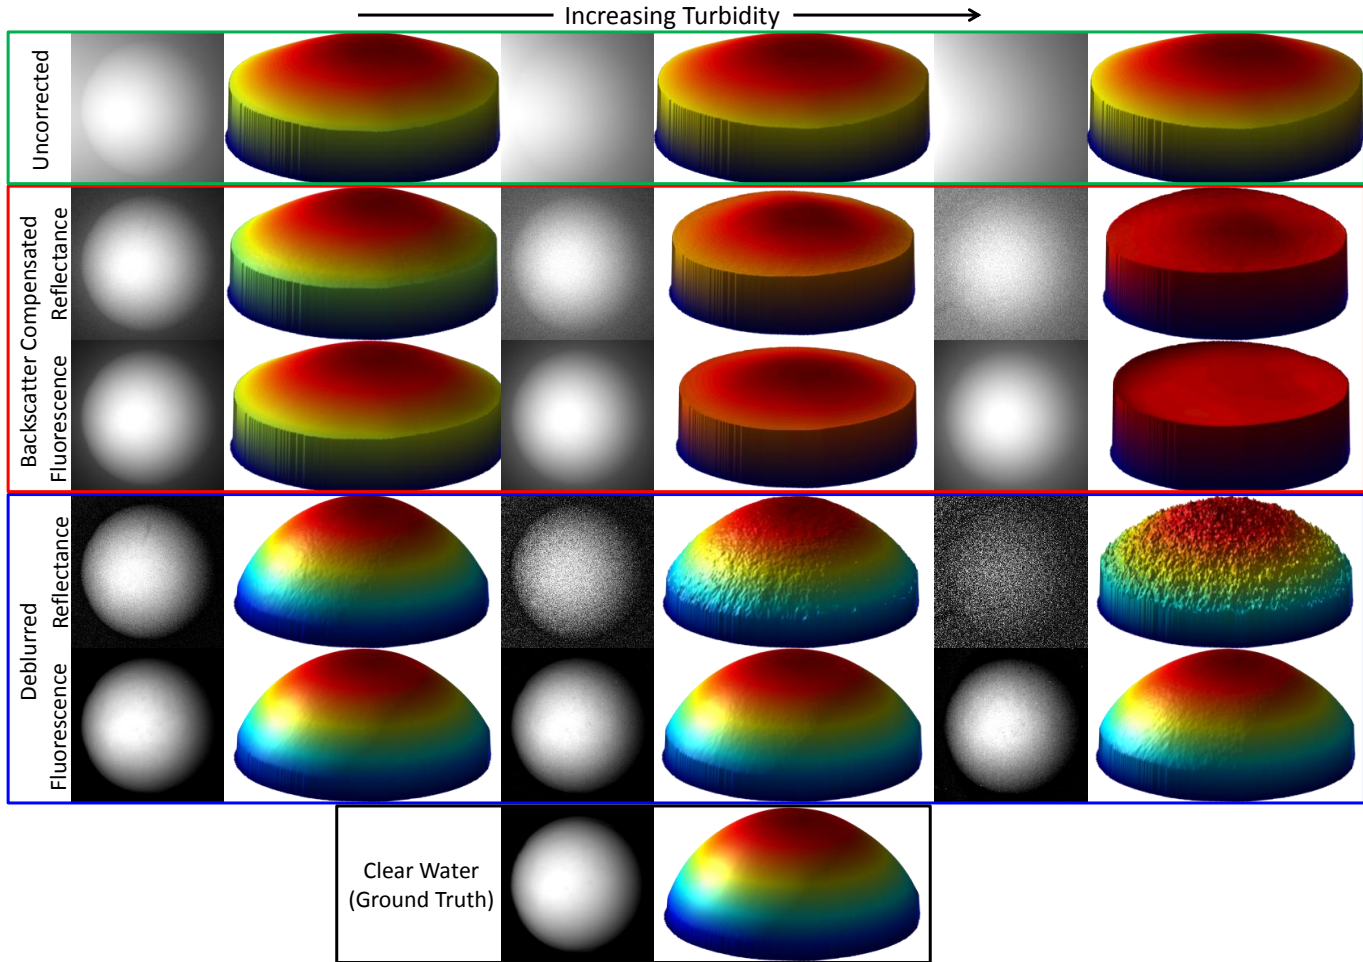


Fig. 10. Input images and resulting surface reconstructions of the spherical cap. The columns depict three levels of increasing turbidity from left to right. [1st row] result of standard photometric stereo (scattering is ignored). The shape is not reconstructed correctly. [2nd row] Result of removing the backscatter as in [8]. The reconstruction is improved but still unsatisfactory. [3rd row] Using fluorescence to remove backscatter. The result is basically the same as backscatter subtraction. [4th row] result of deblurring the backscatter subtracted images. This recovers the shape quite well when the SNR is not too low. However this is not the case in high turbidity. [5th row] result of deblurring the fluorescence images. Here the SNR remains high even in high turbidity and thus we continue to get excellent quality reconstructions. Note the roughness on the fourth row, second column due to noise. [Bottom row] Clear water reconstruction (ground truth).

rows show various reconstruction methods and the columns show the results for turbidity levels 2-4. The bottom row shows an input image and reconstruction in clear water which are treated as ground truth.

In all results, reconstructions from uncorrected images are flattened. Removing backscatter, either by backscatter subtraction (current state-of-the-art [8]), or using fluorescence, but without handling blur, also produces flattened results. For lower turbidities deblurring backscatter subtracted images produces excellent results, but in the highest turbidity, where the backscatter dominates the signal, using fluorescence reduces the noise and results in a smoother surface.

11 CONCLUSION

In this paper, we have developed a comprehensive and novel solution for photometric stereo in a scattering medium. We address each of the three key modes of single scattering, showing how a scattered light source can be

modeled as an unscattered point light source, accounting for blur due to scattering from the object through a novel deconvolution framework, and demonstrating how fluorescence imaging can optically eliminate backscatter, increasing SNR in high turbidity. With the simple *small surface variations approximation*, we reduce the problem to a linear system for the surface normals, almost identical to conventional photometric stereo. Our practical methods for deconvolution and fluorescence can be combined to produce reconstructions almost as accurate as those obtained in air, and significantly better than previous methods.

Future work includes removing the need to know the average object distance, removing the *small surface variations approximation*, and an automated PSF calibration procedure for varying turbidities and depths. Although our theory only applies to a single scattering medium, in practice, our calibrated PSF may be taking multiple scattering effects into account. Extending our theory to multiple scattering would provide further insight.

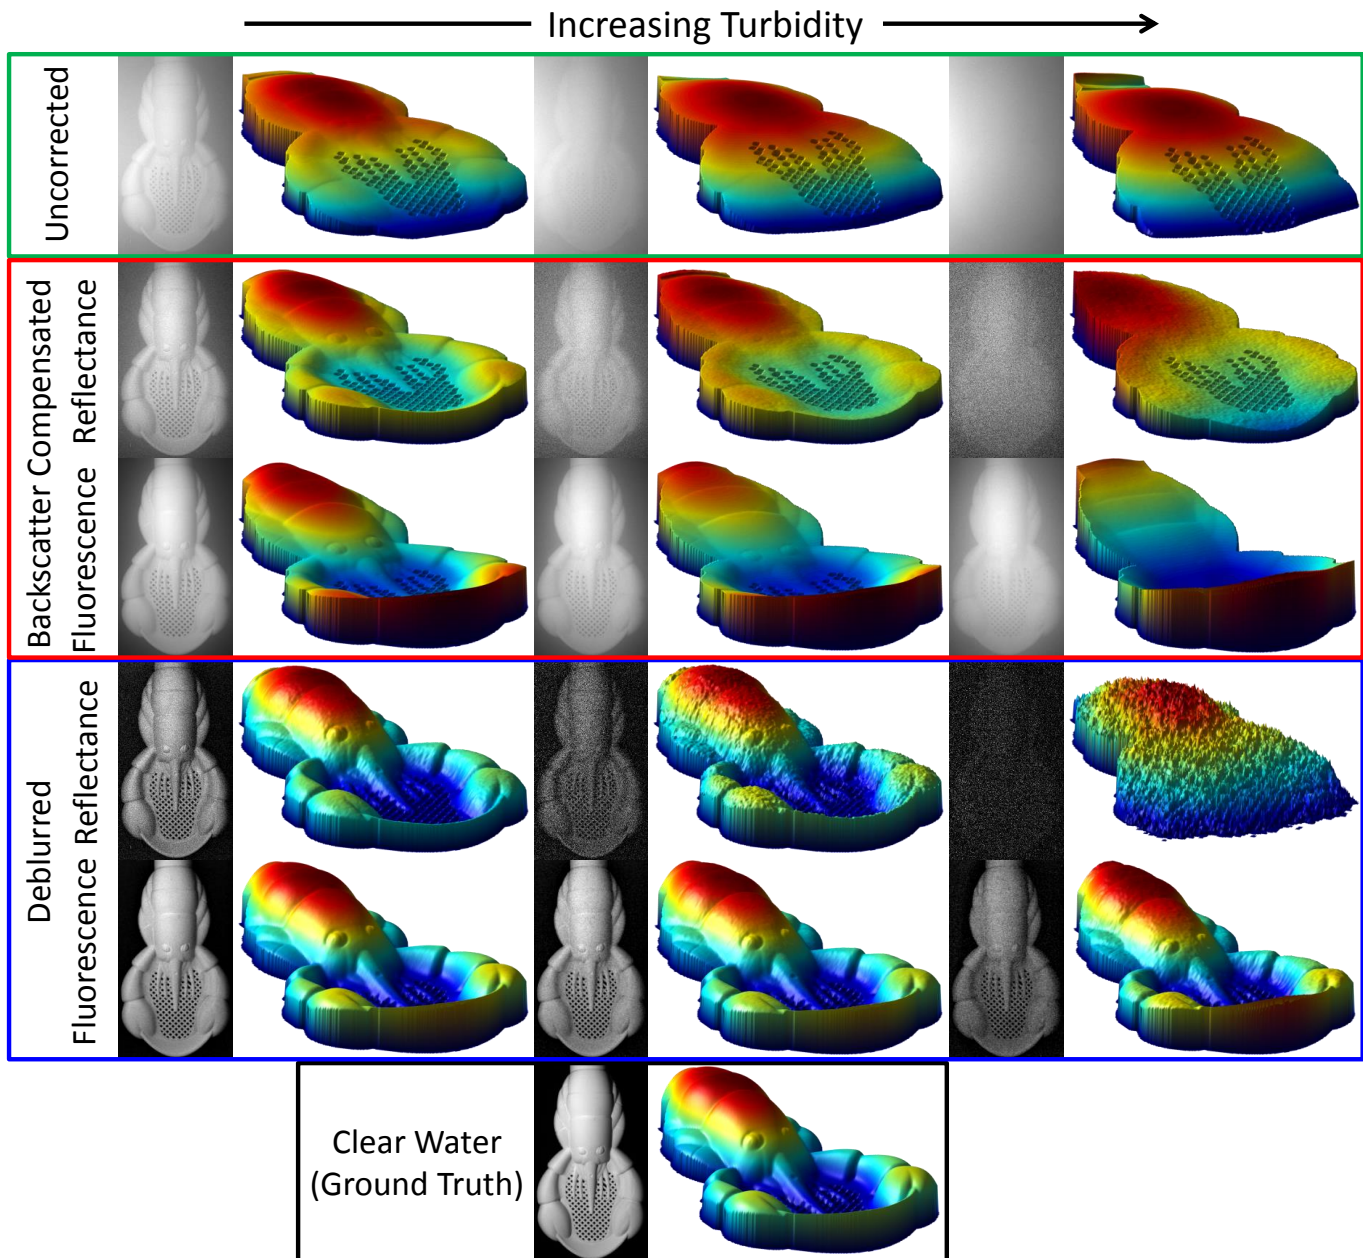


Fig. 11. Input images and resulting surface reconstructions of the toy lobster. The columns depict three levels of increasing turbidity from left to right. [1st row] result of standard photometric stereo (scattering is ignored). The shape is not reconstructed correctly. [2nd row] Result of removing the backscatter as in [8]. The reconstruction is improved but still unsatisfactory. [3rd row] Using fluorescence to remove backscatter. The result is basically the same as backscatter subtraction. [4th row] result of deblurring the backscatter subtracted images. This recovers the shape quite well when the SNR is not too low. However this is not the case in high turbidity. [5th row] result of deblurring the fluorescence images. Here the SNR remains high even in high turbidity and thus we continue to get excellent quality reconstructions. Note the roughness on the fourth row, second column due to noise. [Bottom row] Clear water reconstruction (ground truth).

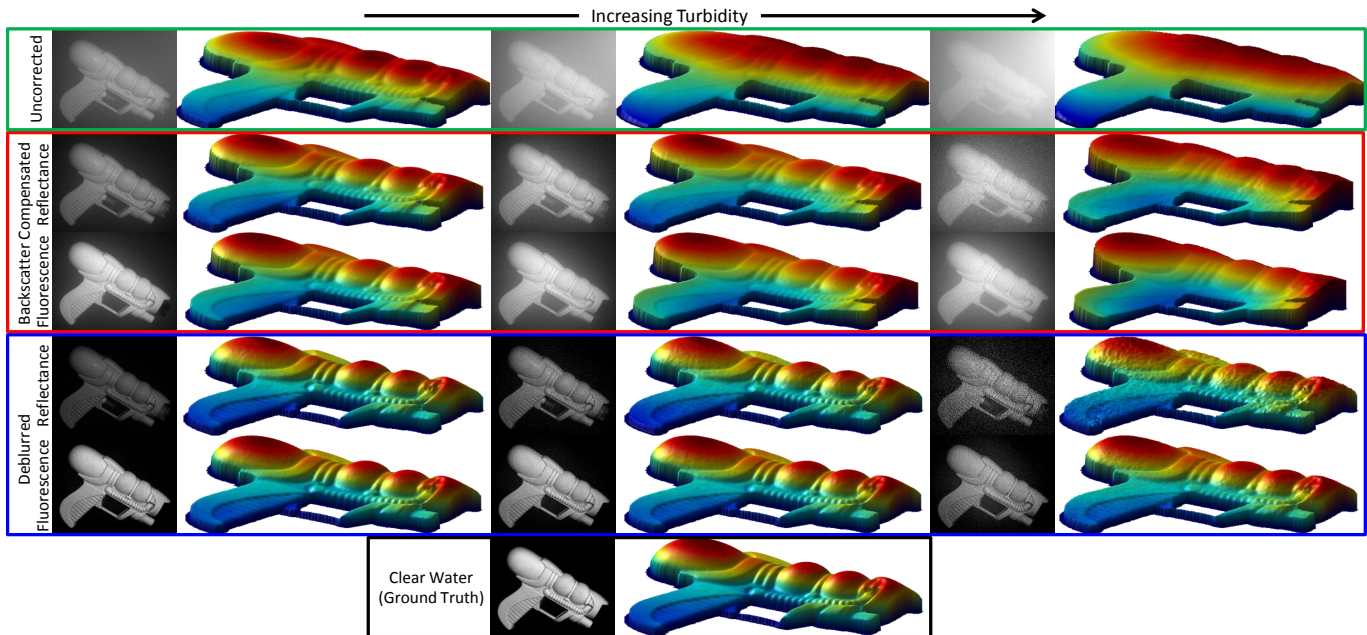


Fig. 12. Input images and resulting surface reconstructions of the toy squirt gun. The columns depict three levels of increasing turbidity from left to right. [1st row] result of standard photometric stereo (scattering is ignored). The shape is not reconstructed correctly. [2nd row] Result of removing the backscatter as in [8]. The reconstruction is improved but still unsatisfactory. [3rd row] Using fluorescence to remove backscatter. The result is basically the same as backscatter subtraction. [4th row] result of deblurring the backscatter subtracted images. This recovers the shape quite well when the SNR is not too low. However this is not the case in high turbidity. [5th row] result of deblurring the fluorescence images. Here the SNR remains high even in high turbidity and thus we continue to get excellent quality reconstructions. Note the roughness on the fourth row, third column due to noise. The object signal is stronger in this case than the lobster and sphere since the medium doesn't contain juice which increases attenuation. [Bottom row] Clear water reconstruction (ground truth).

ACKNOWLEDGMENTS

This work was supported in part by NSF grant ATM-0941760, ONR grant N00014-15-1-2013, W.M. Keck Foundation, and by the UC San Diego Center for Visual Computing. Tali Treibitz was supported by the Ministry of Science, Technology and Space grant 3-12487, and the Technion Ollendorff Minerva Center for Vision and Image Sciences.

REFERENCES

- [1] N. Gracias, P. Ridao, R. Garcia, J. Escartin, M. L'Hour, F. Cibecchini, R. Campos, M. Carreras, D. Ribas, N. Palomeras *et al.*, "Mapping the moon: Using a lightweight auv to survey the site of the 17th century ship la lune," in *Proc. MTS/IEEE OCEANS*. IEEE, 2013.
- [2] D. Kocak, F. Dalgleish, F. Caimi, and Y. Schechner, "A focus on recent developments and trends in underwater imaging," *Marine Technology Society Journal*, vol. 42, no. 1, pp. 52–67, 2008.
- [3] R. Pintus, S. Podda, and M. Vanzì, "An automatic alignment procedure for a four-source photometric stereo technique applied to scanning electron microscopy," *IEEE Trans. Instrumentation and Measurement*, vol. 57, no. 5, pp. 989–996, 2008.
- [4] N. Kolagani, J. Fox, and D. Blidberg, "Photometric stereo using point light sources," in *Proc. IEEE Int. Conf. Robotics and Automation*, 1992, pp. 1759–1764.
- [5] B. Sun, R. Ramamoorthi, S. G. Narasimhan, and S. K. Nayar, "A practical analytic single scattering model for real time rendering," *ACM Transactions on Graphics (TOG)*, vol. 24, pp. 1040–1049, 2005.
- [6] S. Negahdaripour, H. Zhang, and X. Han, "Investigation of photometric stereo method for 3-D shape recovery from underwater imagery," in *Proc. MTS/IEEE OCEANS*, vol. 2, 2002, pp. 1010–1017.
- [7] A. Shashua, "Geometry and photometry in 3d visual recognition," Ph.D. dissertation, Massachusetts Institute of Technology, 1992.
- [8] C. Tsiotsios, M. E. Angelopoulou, T.-K. Kim, and A. J. Davison, "Backscatter compensated photometric stereo with 3 sources," in *Proc. IEEE CVPR*, 2014, pp. 2259–2266.
- [9] T. Treibitz and Y. Y. Schechner, "Resolution loss without imaging blur," *JOSA A*, vol. 29, no. 8, pp. 1516–1528, 2012.
- [10] G. Guilbault, *Practical fluorescence*. CRC, 1990, vol. 3.
- [11] R. Woodham, "Photometric method for determining surface orientation from multiple images," *Opt. Eng.*, vol. 19, no. 1, pp. 139–144, January 1980.
- [12] K. M. Lee, C. J. Kuo *et al.*, "Shape from photometric ratio and stereo," *J. Vis. Comm. & Image Rep.*, vol. 7, no. 2, pp. 155–162, 1996.
- [13] A. Tankus and N. Kiryati, "Photometric stereo under perspective projection," in *Proc. IEEE ICCV*, 2005.
- [14] B. Kim and P. Burger, "Depth and shape from shading using the photometric stereo method," *CVGIP: Image Understanding*, vol. 54, no. 3, pp. 416–427, 1991.
- [15] M. Galo and C. L. Tozzi, "Surface reconstruction using multiple light sources and perspective projection," in *Proc. IEEE ICIP*, vol. 1, 1996, pp. 309–312.
- [16] Y. Iwahori, H. Sugie, and N. Ishii, "Reconstructing shape from shading images under point light source illumination," in *Proc. IEEE ICPR*, vol. 1, 1990, pp. 83–87.
- [17] S. G. Narasimhan, S. K. Nayar, B. Sun, and S. J. Koppal, "Structured light in scattering media," in *Proc. IEEE ICCV*, 2005, pp. 420–427.
- [18] M. Gupta, S. Narasimhan, and Y. Schechner, "On controlling light transport in poor visibility environments," in *Proc. IEEE CVPR*, 2008.
- [19] T. Treibitz and Y. Y. Schechner, "Active polarization descattering," *IEEE Trans. PAMI*, vol. 31, pp. 385–399, 2009.
- [20] J. S. Jaffe, "Computer modeling and the design of optimal underwater imaging systems," *IEEE J. Oceanic Engineering*, vol. 15, no. 2, pp. 101–111, 1990.
- [21] V. V. Tuchin and V. Tuchin, *Tissue optics: light scattering methods and instruments for medical diagnosis*. SPIE press Bellingham, 2007, vol. 13.
- [22] M. Hullin, M. Fuchs, I. Ihrke, H. Seidel, and H. Lensch, "Fluorescent immersion range scanning," *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3, pp. 87–87, 2008.
- [23] I. Sato, T. Okabe, and Y. Sato, "Bispectral photometric stereo based on fluorescence," in *Proc. IEEE CVPR*, 2012.

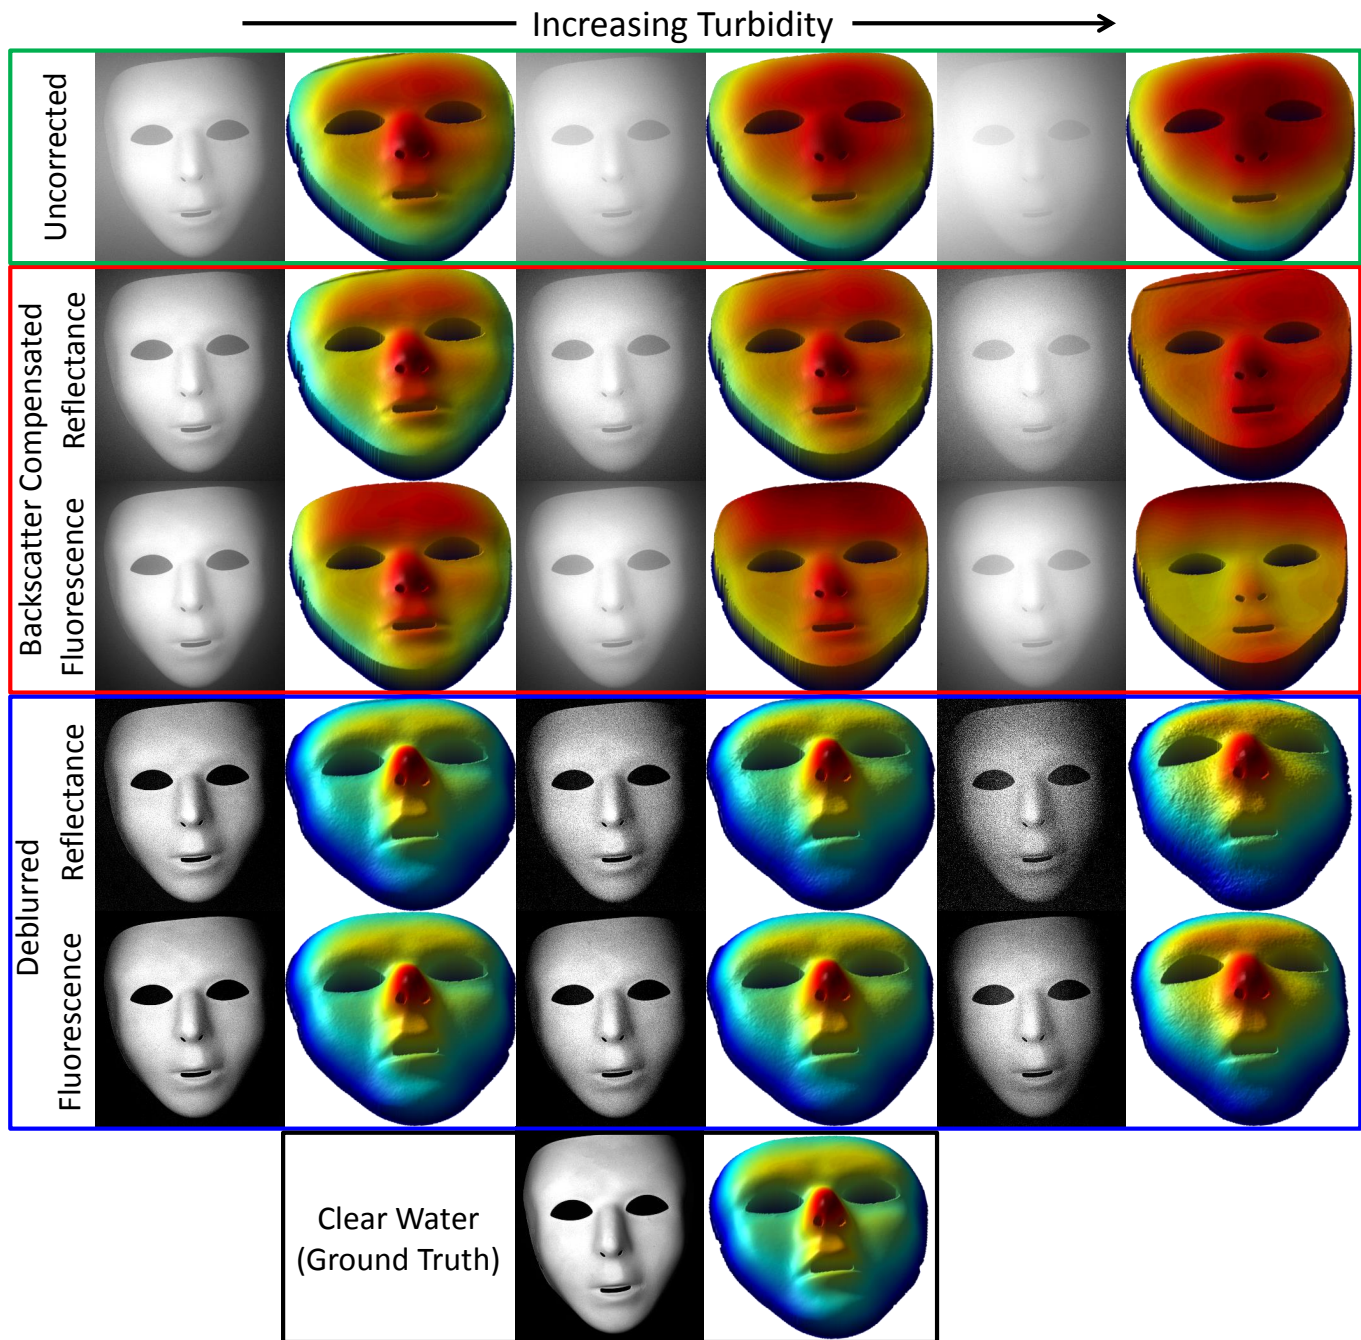


Fig. 13. Input images and resulting surface reconstructions of the mask. The columns depict three levels of increasing turbidity from left to right. [1st row] result of standard photometric stereo (scattering is ignored). The shape is not reconstructed correctly. [2nd row] Result of removing the backscatter as in [8]. The reconstruction is improved but still unsatisfactory. [3rd row] Using fluorescence to remove backscatter. The result is basically the same as backscatter subtraction. [4th row] result of deblurring the backscatter subtracted images. This recovers the shape quite well when the SNR is not too low. However this is not the case in high turbidity. [5th row] result of deblurring the fluorescence images. Here the SNR remains high even in high turbidity and thus we continue to get excellent quality reconstructions. Note the roughness on the fourth row, third column due to noise. Similar to the toy gun, the object signal is stronger in this case than the lobster and sphere since the medium doesn't contain juice which increases attenuation. [Bottom row] Clear water reconstruction (ground truth).

- [24] T. Treibitz, Z. Murez, B. G. Mitchell, and D. Kriegman, "Shape from fluorescence," in *Proc. ECCV*, 2012.
- [25] S. Zhang and S. Negahdaripour, "3-D shape recovery of planar and curved surfaces from shading cues in underwater images," *IEEE J. Oceanic Engineering*, vol. 27, no. 1, pp. 100–116, 2002.
- [26] B. McGlamery, "A computer model for underwater camera systems," in *Ocean Optics VI*. International Society for Optics and Photonics, 1980, pp. 221–231.
- [27] E. Trucco and A. T. Olmos-Antillon, "Self-tuning underwater image restoration," *IEEE J. Oceanic Engineering*, vol. 31, no. 2, pp. 511–519, 2006.
- [28] C. Inoshita, Y. Mukaigawa, Y. Matsushita, and Y. Yagi, "Surface normal deconvolution: Photometric stereo for optically thick translucent objects," in *Proc. ACCV*. Springer, 2014, pp. 346–359.
- [29] B. Dong, K. D. Moore, W. Zhang, and P. Peers, "Scattering parameters and surface normals from homogeneous translucent materials using photometric stereo," in *Proc. IEEE CVPR*, 2014, pp. 2299–2306.
- [30] K. Tanaka, Y. Mukaigawa, H. Kubo, Y. Matsushita, and Y. Yagi, "Recovering inner slices of translucent objects by multi-frequency illumination," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 5464–5472.
- [31] Z. Murez, T. Treibitz, R. Ramamoorthi, and D. Kriegman, "Photometric stereo in a scattering medium," in *Proc. IEEE ICCV*, 2015.
- [32] S. G. Narasimhan, M. Gupta, C. Donner, R. Ramamoorthi, S. K. Nayar, and H. W. Jensen, "Acquiring scattering properties of participating media by dilution," *ACM Transactions on Graphics (TOG)*, vol. 25, no. 3, pp. 1003–1012, 2006.
- [33] L. G. Henyey and J. L. Greenstein, "Diffuse radiation in the galaxy," *The Astrophysical Journal*, vol. 93, pp. 70–83, 1941.
- [34] A. Agrawal, R. Raskar, and R. Chellappa, "What is the range of surface reconstructions from a gradient field?" in *Proc. ECCV*. Springer, 2006, pp. 578–591.
- [35] J. Y. Bouguet, "Camera calibration toolbox for matlab," vision.caltech.edu/bouguetj/calib/_doc.
- [36] S. G. Narasimhan and S. K. Nayar, "Shedding light on the weather," in *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, vol. 1. IEEE, 2003, pp. 1–665.
- [37] N. Joshi, R. Szeliski, and D. Kriegman, "PSF estimation using sharp edge prediction," in *Proc. IEEE CVPR*. IEEE, 2008, pp. 1–8.



Zak Murez received his BS degree in mathematics and computer science from Yale University in 2011 and his MS degree in computer science from the University of California, San Diego in 2015. He is currently a PhD student at the University of California, San Diego, advised by David J. Kriegman and Ravi Ramamoorthi.



Tali Treibitz received her BA degree in computer science and her PhD degree in electrical engineering from the Technion-Israel Institute of Technology in 2001 and 2010, respectively. Between 2010–2013 she has been a post-doctoral researcher in the department of computer science and engineering, in the University of California, San Diego and in the Marine Physical Lab in Scripps Institution of Oceanography. Since 2014 she is heading the marine imaging lab in the school of Marine Sciences in the University

of Haifa. She was the recipient of the Google Anita Borg Scholarship in 2009 and the Weizmann Institute of Science National Postdoctoral Award for Advancing Women in Science in 2010.



Ravi Ramamoorthi received his BS degree in engineering and applied science and MS degrees in computer science and physics from the California Institute of Technology in 1998. He received his PhD degree in computer science from the Stanford University Computer Graphics Laboratory in 2002, upon which he joined the Columbia University Computer Science Department. He was on the UC Berkeley EECS faculty from 2009–2014. Since July 2014, he is a Professor of Computer Science and Engineering at

the University of California, San Diego, where he holds the Ronald L. Graham Chair of Computer Science. He is also the founding Director of the UC San Diego Center for Visual Computing. His research interests cover many areas of computer vision and graphics, with more than 100 publications. His research has been recognized with a number of awards, including the 2007 ACM SIGGRAPH Significant New Researcher Award in computer graphics, and by the white house with a Presidential Early Career Award for Scientists and Engineers in 2008 for his work on physics-based computer vision. He has advised more than 20 Postdoctoral, PhD and MS students, many of whom have gone on to leading positions in industry and academia; and he has taught the first open online course in computer graphics on the EdX platform in fall 2012, with more than 100,000 students enrolled in that and subsequent iterations. He is a finalist for the inaugural edX Prize for exceptional contributions in online teaching and learning.



David J. Kriegman received the BSE degree in electrical engineering and computer science from Princeton University in 1983. He received the MS degree in 1984 and the PhD degree in electrical engineering in 1989 from Stanford University. Since 2002, he has been a Professor of Computer Science and Engineering in the Jacobs School of Engineering, University of California, San Diego (UCSD). Prior to joining UCSD, he was an Assistant and Associate Professor of Electrical Engineering and Computer

Science at Yale University (1990–1998) and an Associate Professor with the Computer Science Department and Beckman Institute at the University Illinois at Urbana-Champaign (1998–2002). He was founding CEO of Taaz Inc., the leader in photorealistic virtual try on, and a founding partner of KBVT. Presently, he is on leave from UCSD and is part of the Machine Learning Team at Dropbox.

His research is in computer vision with particular application to face recognition, robotics, computer graphics, microscopy, and coral reef ecology. He was chosen for the National Science Foundation Young Investigator Award, and has received Best Paper Awards at the 1996 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), the 1998 European Conference on Computer Vision, and the 2007 International Conference on Computer Vision (Marr Prize, runner up) as well as the 2003 Paper of the Year Award from the Journal of Structural Biology. He has served as program co-chair of CVPR 2000 and general co-chair of CVPR 2005. He was the editor-in-chief of the IEEE Transactions on Pattern Analysis and Machine Intelligence from 2005 to 2008. He is a Fellow of the IEEE.