

Backscatter Compensated Photometric Stereo with 3 Sources

Chourmouzios Tsiotsios Maria E. Angelopoulou Tae-Kyun Kim Andrew J. Davison
Imperial College London, UK

{c.tsiotsios, m.angelopoulou, tk.kim, a.davison}@imperial.ac.uk

Abstract

Photometric stereo offers the possibility of object shape reconstruction via reasoning about the amount of light reflected from oriented surfaces. However, in murky media such as sea water, the illuminating light interacts with the medium and some of it is backscattered towards the camera. Due to this additive light component, the standard Photometric Stereo equations lead to poor quality shape estimation. Previous authors have attempted to reformulate the approach but have either neglected backscatter entirely or disregarded its non-uniformity on the sensor when camera and lights are close to each other.

We show that by compensating effectively for the backscatter component, a linear formulation of Photometric Stereo is allowed which recovers an accurate normal map using only 3 lights. Our backscatter compensation method for point-sources can be used for estimating the uneven backscatter directly from single images without any prior knowledge about the characteristics of the medium or the scene. We compare our method with previous approaches through extensive experimental results, where a variety of objects are imaged in a big water tank whose turbidity is systematically increased, and show reconstruction quality which degrades little relative to clean water results even with a very significant scattering level.

1. Scattering and Shape

Applying Photometric Stereo (PS) in murky media has had limited success so far, since the measured light carries information not only about the scene orientation but also about the medium itself. Specifically, light gets attenuated and scattered by the medium's particles, adding a strong unwanted signal to the measured brightness, the so-called backscatter component. Figure 1 illustrates the severity of this effect on image quality, dictating special treatment if any effective photometric task is to be considered.

In this paper we formulate PS within a scattering environment by estimating and taking account of the strong uneven backscatter created from point sources close to the



Figure 1: One of the three captured murky-water images, the restored image, and our final reconstruction result.

camera. We describe that the camera-lights baseline, which is irrelevant in pure air, is crucial in scattering media. Thus we relax the assumption of [12] that all lights are far away from the camera and investigate how this affects the varying backscatter from every source. As we describe, the additional ambiguity that is introduced by the backscatter makes a direct solution of the PS equations ineffective and hence we propose both a calibrated and an automatic way for estimating the unknown backscatter directly from the measured images. This gives further potential for single-image restoration in turbid media under directional lighting.

As we show, the backscatter compensation yields a linear PS solution that estimates the scene normals and a scaled version of the albedo. In our practical approach, only 3 sources are required as for PS in pure air, in a symmetric arrangement around the camera. We demonstrate our results through extensive experiments in a big water tank where the camera and light sources are both immersed into the turbid medium and separated by a small distance. We evaluate the performance of our method over a wide range of controlled scattering levels by adding milk of gradually increasing quantity and show that our method outperforms related approaches, compensating effectively for the backscatter effect and yielding shape recovery results similar to those in clean water even for high levels of water turbidity.

Related Work: A large amount of work has focused on modelling image formation within scattering environments with diffuse illumination in haze and fog [6, 11, 16], or sub-sea [3, 14]. There the lighting term is constant and equally illuminates all the particles of the medium between

the camera and the scene. As the depth of the imaged scene increases, the volume of illuminated particles increases as well. This implies that under diffuse illumination backscatter is directly associated with scene depth, and this can be used for estimating a rough depth map of the scene.

However, in Shape from Shading methods illumination is provided by artificial directional sources. This is also the case for deep sea underwater vehicles which carry a camera and light sources at a small offset [8, 13, 17]. Here, due to the limited beam angle of the sources, only some of the medium's particles are illuminated depending to their position from the source [5, 18]. Moreover, backscatter is saturated after a small depth from the source [17], after which it cannot serve the purpose of a depth prior as in diffuse lighting. These characteristics require careful backscatter modelling and compensation when directional sources are employed, as it happens with Photometric Stereo.

Instead of compensating for the backscatter effect, some previous approaches have aimed to minimize its impact by changing the hardware setup [8], using active methods [2, 5] or fusing multiple images [18]. Nevertheless, in typical imaging conditions where a camera and a conventional light are carried by the same body at a small offset, backscatter was still shown to be the main effect degrading image quality [10, 17]. In [17], it was measured using special polarizing hardware, while in [9, 10] it was estimated directly from a low-pass filtered version of the captured image.

As we show in this paper, the effect of backscatter on photometric-based shape recovery is highly significant. Related works [13, 19] neglected backscatter under the assumption that the water murkiness level is low or a special setup has been employed for its minimization. In [1], the decreasing effectiveness of PS due to the presence of an additive light component was also described, yet no further effort was made to model or remove its effect. PS in a scattering medium was studied in [12], suggesting that a non-linear system of at least 5 equations should be solved. The problem can be reduced to a 4-source linear solution if the scattering is isotropic and the sources are symmetric. Although effective shape reconstruction was achieved, the sources in the setup of [12] were outside water and distant from the camera. By adopting this type of setup, the backscatter saturation that occurs as depth increases and its non-uniformity on the image sensor were neglected.

2. Light Propagation Model

When illumination comes from an artificial source, due to Inverse Square Law (ISL) its intensity I_k is inversely proportional to the distance d travelled. In turbid media, light is further attenuated exponentially with distance:

$$I_d = \frac{I_k}{d^2} e^{-cd}, \quad (1)$$




Figure 2: Geometry of light propagation. The measured brightness is the sum of the scene originated direct component D , and the additive backscatter component B .

where c is the total attenuation coefficient of the medium which describes the light loss per unit distance [8], modelled as $c = a + b$. a and b denote the medium's absorption and scattering coefficients. b reflects the superposition of all scattering incidents around a unit volume particle and can be calculated by integrating the angular scattering function of the medium $\beta(\theta)$ over all the directions around the particle: $b = 2\pi \int_0^\pi \beta(\theta) \sin(\theta) d\theta$. We employ the low-order representation of [12] defining $\beta(\theta)$ as $\beta(\theta) = \frac{b}{4\pi} (1 + g \cos \theta)$, where $g \in (-1, 1)$.

Consider a camera and an artificial point-light source both located on the same plane in a scattering medium and separated by a small distance h as in Figure 2. We define a coordinate system whose xy plane coincides with the camera's image sensor, and z -axis coincides with the camera's optical axis. As in related works [12, 13, 17], we assume that the size of the imaged object is small with respect to its distance (depth) z_O from the camera, such that we can use an orthographic camera model. The total irradiance due to light source k that is measured by a pixel with coordinates (x, y) on the image sensor is the sum of the direct and backscatter components, denoted as $D_k(x, y)$ and $B_k(x, y)$, respectively¹. For simplicity we omit the pixel (x, y) coordinates notation in the following derivations.

Direct Component: The direct component is the light beam travelling distance $\sqrt{h^2 + z_O^2}$ from the source to the scene that then gets reflected and travels a distance z_O to the camera sensor. Considering the attenuation of (1), the incident component on the scene is $I_k \frac{e^{-c\sqrt{h^2 + z_O^2}}}{h^2 + z_O^2}$. We assume here that the surface is Lambertian with a constant albedo ρ and unit normal vector \mathbf{n} , while light direction is denoted by unit vector \mathbf{s}_k . The scene-reflected light will be then attenuated again by e^{-cz_O} before it reaches the sensor. Hence, the total direct component is:

¹As in the prior art, we adopt the single scattering model. We ignore forward scattering phenomena [19] and focus on the backscatter which was shown to have the dominating impact on image quality [17, 10].

$$D_k = \frac{I_k}{h^2 + z_O^2} e^{-c(z_O + \sqrt{h^2 + z_O^2})} \varrho \mathbf{n} \cdot \mathbf{s}_k. \quad (2)$$

Backscatter Component: Consider now a differential scattering volume at distance z across the Line-Of-Sight (LOS) of the sensor pixel. As before, the scattering volume is illuminated by $I_v = \frac{I_k}{z^2 + h^2} e^{-c\sqrt{h^2 + z^2}}$. This backscatters through angle θ a light component equal to $\beta(\theta)I_v$, which is attenuated by e^{-cz} until it reaches the sensor. Thus, replacing $\beta(\theta)$ by $\frac{b}{4\pi}(1 + g \cos \theta)$, and $\cos \theta$ by $\frac{z}{\sqrt{h^2 + z^2}}$, the differential backscatter component that reaches the sensor pixel is $dB_k(z) = \frac{bI_k}{4\pi} (1 + g \frac{z}{\sqrt{z^2 + h^2}}) \frac{e^{-c(z + \sqrt{h^2 + z^2})}}{h^2 + z^2}$.

In order to estimate the total backscattered light, we have to integrate over all depth positions z along the pixel's LOS where differential volumes of particles are illuminated from the source and backscatter part of this light towards the sensor. Thus, it should be emphasized here that the lower limit of the integral does not equal 0 as in [12], but the minimum depth position on the LOS where the particles are illuminated, denoted by z_k in Figure 2. z_k corresponds to the intersection point between the pixel's LOS and the limited beam angle of the illumination source. The maximum limit of the integral equals the scene depth z_O which is the final point on the LOS where particles receive and backscatter light towards the sensor. Finally, the total backscatter component on pixel (x, y) due to the source k is given as

$$B_k = \int_{z=z_k}^{z=z_O} \frac{bI_k}{4\pi} (1 + g \frac{z}{\sqrt{z^2 + h^2}}) \frac{e^{-c(z + \sqrt{h^2 + z^2})}}{h^2 + z^2} dz. \quad (3)$$

There exists no closed-form solution for the integral. However, it was shown that it is a smooth function that is straightforward to evaluate numerically [15].

3. Photometric Stereo in a Scattering Medium

In pure-air PS, the product $\varrho \mathbf{n} \cdot \mathbf{s}_k$ yields a system of linear equations that can recover the normals and albedo for every pixel using at least 3 sources. Based on the image formation model of Section 2, within a scattering medium the total measured intensity E_k for every pixel/source equals the sum $D_k + B_k$. The equation is non-linear and has a complicated form. Increasing the number of sources is a common practice to determine the PS problem when further unknowns are added. Narasimhan *et al.* [12] showed that within turbid media where sources are far away from the camera at least 5 sources are required for determining the non-linear system of equations. Otherwise, for a 4-source symmetric setup and isotropic medium ($g = 0$), the backscatter for each pixel should be equal for all 4 sources: $B_{k_i} = B_{k_j} \forall i, j \in \{1, 2, 3, 4\}$, and thus can be eliminated by subtracting pairs of measured intensities: $E_{k_i} - E_{k_j} = D_{k_i} - D_{k_j}$. Let us examine the problem for

our case, where, in addition to the considerations of [12], the limited beam angle of the k_{th} source next to the camera that leads to $z_k \neq 0$ and ISL are also considered.

The medium coefficients b , c and g are constants-global unknowns within a uniform scattering medium. The unknown values for every pixel are: the albedo ϱ , the normal vector \mathbf{n} (ϱ and \mathbf{n} are considered as 3 unknowns together since \mathbf{n} is a unit vector), and the depth variables z_O, z_k . At a first glance the total number of unknowns for each pixel are 5 plus 3 global unknowns, and thus 6 sources would be enough to determine all of the unknowns for every pixel.




Figure 3: The LOS of a sensor pixel intersects the beam angle of each light source k_i at a different depth z_{k_i} . Thus, each source creates a distinct backscatter component on the sensor (illustrated by backscatter images B_1 and B_2).

3.1. Backscatter Variation for Each Source - z_k

In fact, contrary to the rest of the parameters, z_k does not have a single value for every pixel (x, y) since it differs for each source (Figure 3). When a pixel is closer to a light source k_i , its LOS will intersect the beam angle of k_i at a smaller depth point z_{k_i} than it will intersect the other sources. Thus, for this source the integration path between z_{k_i} and z_O will be larger than for other sources, adding a higher backscatter value B_i to the pixel. This holds for every pixel, and hence every source finally creates an uneven backscatter component on the sensor according to its position with respect to each pixel. The synthetical backscatter images (using (3)) of Figure 3 illustrate this non-uniformity due to varying z_k for every pixel/source. In order to estimate z_k we would have to calculate the exact intersection point of each pixel's LOS with the beam angle of every source. Such a task would be very difficult since the precise 3D position, rotation and beam angle of the source should be known. Note here that z_k , as opposed to scene depth z_O which might be far away from the camera-lights setup (as




Figure 4: Due to ISL backscatter is saturated after a small depth z_{sat} away from the camera and hence it captures no information about scene's depth z_O .

in orthographic projection), is only a few cm away from the sensor² where the backscatter function (3) is highly variant and hence small errors in calibration of z_k are expected to have a strong impact on backscatter estimation [17].

Due to the varying $z_k(x, y)$ for every pixel-source combination, increasing the number of sources also increases the number of the unknowns. This leaves the problem underdetermined regardless of the number of sources. Furthermore, due to this variation the linear solution of [12] that assumes that backscatter for every pixel would be equal for every source $B_{k_i} = B_{k_j}$ does not hold, as $z_{k_i} \neq z_{k_j}$.

3.2. Linear Solution with 3 Sources

Since the extra ambiguity due to z_k lies solely in the additive backscatter term, we suggest that the whole backscatter component is estimated for each pixel-source and subtracted from the measured brightness without having to determine the exact value of z_k (Section 4). This would leave us with the direct component for every source, given by (2). The albedo and the attenuating factor can be combined into a scene-depth dependent function defined as $\varrho_{sc}(z_O) = \varrho \frac{1}{h^2 + z_O^2} e^{-c(z_O + \sqrt{h^2 + z_O^2})}$ for every pixel.

Thus, after subtracting the estimated backscatter we end up with the direct component that carries the shape information of the scene and a scaled version of the albedo $\varrho_{sc}(z_O)$:

$$E_k(x, y) - B_k(x, y) = I_k \varrho_{sc}(z_O(x, y)) \mathbf{n} \cdot \mathbf{s}_k. \quad (4)$$

Given that the light intensity and direction are known for each source, we can estimate the unit normal vector \mathbf{n} and scaled albedo $\varrho_{sc}(z_O)$ of the surface patch, given at least 3 light sources. Since $E_k(x, y)$ is the measured intensity, the problem now becomes estimation of the varying backscatter component $B_k(x, y)$ for every pixel (x, y) and source k .

3.3. Backscatter Variation with Scene Depth z_O

As (3) indicates, backscatter is a function of both the minimum lighted depth z_k and the scene depth z_O . These two values determine the integration path (pointing arrows in Figure 3), whose length determines the intensity of the

²Typical values of $z_k \in [10 - 30cm]$ were approximated in [17].

backscatter. According to Section 3.1, pixels that are positioned closer to a source will have a bigger integration path due to smaller z_k . Let us now examine the upper limit z_O .

Contrarily with cases of diffuse or distant from camera illumination [6, 12], backscatter for point-sources is saturated after a small depth value due to ISL [17]. Figure 4a shows the backscatter function for increasing scene-depth when ISL is considered. It equals 0 below the minimum lighted depth z_k and then exhibits a rapid increase until it reaches saturation, while it is smoothly increased with scene depth when ISL is omitted. The saturation indicates that backscatter dependence on z_O can be safely omitted after z_{sat} , where the scattered light by the particles becomes negligible: $B_k(z_O) = B_k(\infty), \forall z_O \in [z_{sat}, \infty]$ (Figure 4b). z_{sat} is within the small range of 0.5 – 1.5m from the camera [17]³, which indicates that in typical orthographic PS setup (Section 5) backscatter will be saturated, and z_O can be replaced by ∞ in the backscatter function term.

Interestingly, according to numerical evaluations using (2) and (3), even for $z_O < z_{sat}$ backscatter variation with scene depth can be omitted when the total measured brightness E_k is considered. For the small depths below z_{sat} we expect the intensity of the illuminating light, and the direct component D_k to be very high due to ISL. Figure 4c shows the respective direct component over the varying depth z , along with the absolute values of the backscatter component of Figure 4a. For small depths where backscatter isn't still totally saturated, the measured brightness seems to be dominated by the direct component intensity. In order to estimate the impact of the backscatter variation with scene depth with respect to the actual measured brightness E , we calculate the ratio $\frac{B(z)}{E(z)}$ within a varying scene depth range z , and the ratio $\frac{B(\infty)}{E(z)}$ which approximates the backscatter by its saturation value even for scene depths below z_{sat} where it is still unsaturated. Figure 4d indicates that these two differ with a small error value $\epsilon(z)$ at every depth.

This is negligible at a depth point smaller than z_{sat} ⁴.

³This range limit coincides with the one we found through numerical evaluations using (3).

⁴We consider as negligible any value smaller than 1/255 that is the minimum sensing capability of the majority of imaging sensors (8-bit).

Using (3), we have run extensive numerical simulations for a wide range of its unknown values ($b, c \in (0 - 2m^{-1})$, $z_O \in (0.1 - 10m)$, $z_k \in (0.01 - 0.5m)$) and the error ϵ for any $z_O > 0.5m$ was measured to be as low as 0.02 brightness values. The respective error for any $z_O > 1m$ which is the imaging range in our PS setup was below 0.005 brightness levels which is similar to sensor or quantization noise level. Considering also that in PS applications we expect a much smaller scene depth range than the one used in the numerical simulations, due to ISL we can safely neglect the backscatter dependence on any scene depth $z_O > 0.5m$ and replace $z_O = \infty$ in (3).

4. Backscatter Estimation

Previous Work: The task of estimating the backscatter component when directional sources are employed has drawn limited attention compared with the respective cases of diffuse lighting [6, 14]. The work of Mortazavi and Oakley [9, 10] was the only work we found estimating this directly from the image brightness. The dependence of backscatter on z_k and its saturation with scene depth were omitted in this work and as it was assumed, the measured backscatter was proportional to a low-pass filtered version of the image $B(x, y) \simeq \gamma \bar{E}(x, y)$. $\bar{E}(x, y)$ was used for extracting the illumination variation of the image by suppressing any high-frequency details and thus a recursive Gaussian filter with a large parameter σ was used. As we show next, the assumption that backscatter follows a low-pass filtered version of the image is unrealistic in many cases, overestimating the backscatter and introducing high errors in PS.

Proposed Calibrated Method: As it was described in Section 3, for point-light sources next to the camera backscatter is saturated, and thus the varying integration path that results in an uneven backscatter for every pixel-source combination is attributed to z_k . Thus, the backscatter component can be estimated by capturing images when the camera looks at ∞ , directly measuring the saturation value $B_k^\infty(x, y)$ of every pixel⁵. In a finite tank, this can be done using a flat matte black canvas to produce $D = 0$.

Proposed Automatic Method: As $z_k(x, y)$ varies smoothly for every pixel due to its clear geometrical meaning (Section 3.1), the respective backscatter function is also smooth. Specifically, $B_k(x, y)$ would have its maximum at the pixel position which is closest to the source k and then smoothly decrease for the rest of sensor pixels (Figure 3). This smoothness gives us insight that knowing the backscatter intensity of only a few pixels, we can approximate the whole smooth backscatter function over the sensor:

$$B_k \simeq f_k(\mathbf{X}, \boldsymbol{\alpha}), \quad (5)$$

⁵The calibration should be done separately for every source creating a backscatter lookup table for each pixel-source combination.

where \mathbf{X} is the set of all pixel coordinates (x, y) , and $\boldsymbol{\alpha}$ are the unknown parameters of the model that approximates B_k . Due to the smoothness of the function and its unique maximum on image border, we found that a 2D quadratic function $f_k(x, y) = \alpha_0 + \alpha_1 x^2 + \alpha_2 y^2 + \alpha_3 xy + \alpha_4 x + \alpha_5 y$ can estimate with negligible errors the true B_k function of (3)⁶. A set of at least 6 points with known backscatter component are needed in order to define the 6 unknown coefficients of $\boldsymbol{\alpha}$, although more pixels would be necessary for robustness, as described in Section 5. Figure 5a shows the true backscatter function $B_k(x, y)$ simulated using the model of (3), and the resulting fitted quadratic function f_k using the backscatter values of only 6 points.




Figure 5: (a) Simulated ground-truth backscatter and fitted quadratic function using 6 random points. (b) Automatic backscatter estimation for a murky-water object image.

The problem now comes to the selection of at least 6 backscatter pixels that are the input for our regression. Potential candidates are pixels that correspond either to dark scene points, *i.e.* $q(x_B, y_B) = 0$, or to infinite depth points $z_O(x_B, y_B) = \infty$. From (2), the respective direct component for these pixels will be 0 and hence the measured brightness corresponds to the actual backscatter intensity:

$$E_k(x_B, y_B) = \frac{D_k(x_B, y_B)}{z_k(x_B, y_B)} + B_k(x_B, y_B).$$

In order to select a potential set of backscatter pixels for which $D_k(x_B, y_B) = 0$, we divide the image into a number of $N \times N$ blocks and choose the pixel with the lowest intensity in each block (Figure 5b). In reality, not all of the selected points have zero direct component, which introduces a number of outliers. For this purpose, we exploit a RANSAC approach which iteratively evaluates a 6-point set out of the selected pixels. We also take advantage of the physical characteristics of our model in order to facilitate the outlier rejection. Specifically, given that backscatter for each source has its maximum on a border pixel that is closer to the source, we reject solutions that estimate the maximum of f_k on non-border pixels. Furthermore, in our case the outliers should be always additive to our model f_k since they correspond to a positive direct component: $f_k + D_k$, $D_k > 0$. Thus, we penalize solutions that have outliers below the fitted function, by adding the absolute

⁶The selection of the quadratic function for the regression is supported by ground-truth comparison of real data in Section 5.

number of these outliers to the RANSAC score count. Figure 5b shows the resulting estimated function f_k , together with the inliers and outliers of our RANSAC approach. This procedure yields an automatic backscatter estimation for each light source, which requires no prior knowledge about the characteristics of the source, the medium or the scene.

5. Experiments

Our experimental setup (Figure 6) consists of a rectangular-frame pool with a water volume of 5000L. Both the underwater lights and the camera were placed in the water, imitating the setup of an underwater robotic vehicle. Specifically, 4 lights were on the corners of a square baseline with side length 0.4m around the camera. The camera is a Nikon D60 with a AF-S Nikkor 35mm f/1.8G lens. The imaged objects are matte, their size (each dimension) is within 10 – 15cm, and they were all captured at approximately 1.2m depth, enforcing the orthographic assumption. To simulate the scattering effect, we made a linear scale of 15 turbidity steps ranging from totally clean up to heavily murky, by adding milk to the water (Figure 9).




Figure 6: Top row: Calibrating the camera in clean water. Bottom row: Imaged man-made objects.




Figure 7: (a) Backscatter estimation error according to the number of backscatter pixels (x-axis). (b) Backscatter estimation error for the Head, Turtle, Deer, GNOME, Ladybug and Sphere objects (x-axis).

Backscatter Estimation: The backscatter component for all 4 sources and 15 scattering levels was first measured through the calibration step by placing a matte black canvas at 2m depth in our setup. In order to examine how well the quadratic function can fit to the ground-truth backscatter given only a small number of its points, we estimated the

RMSE between the real and the estimated function (Figure 7a) after selecting a different random combination of its points (x-axis). Regardless of the number of blocks, the error was as low as noise variation when at least 8 pixels were used, supporting the validity of the quadratic function.

Then the performance of our automatic estimation method under the presence of objects, where a high number of outliers existed, was compared with the calibrated ground-truth backscatter. Figure 7b shows the RMSE for our method, and the one of Mortazavi and Oakley [9] (Section 4). Our method outperformed [9] for all the imaged objects. Regarding the Sphere and Head objects, whose images include a significantly large number of white pixels, backscatter was still estimated effectively, while the error for [9] increased significantly, overestimating the backscatter due to the unrealistic assumption that this is proportional to a low-pass filtered version of the image. For these objects our method rejected all the outliers on the white object, approximating the backscatter from a small number of scene points on the background. As we described earlier, an advantage of our method is that due to backscatter saturation with scene depth, the selected backscatter pixels do not have to be dark patches on the object; they can also be dark or infinity points on the background⁷. In the case where white objects covered the whole image would lead to erroneous backscatter estimation. However, such a case would be rare in deep-sea scenarios where infinite depth usually surrounds the imaged objects [17], and even then additional frames could be employed by moving the camera to target surrounding dark or infinity pixels.

Shape Recovery: The recovered normal vectors were quantitatively assessed using the Sphere object whose normals are a priori known. Figure 8a shows the *RMSE* value between the estimated and the ground-truth maps for each scattering level considering: our 3-source proposed method for both calibrated and uncalibrated backscatter estimation, traditional PS method where backscatter is neglected, the 4-source linear approach of Narasimhan et al. [12] described in Section 3, and PS method after estimating and subtracting backscatter using [9]. Our approach yields effective normals estimation similar to those in totally clean water, for a very wide range of scattering levels. The performance is decreased beyond 1.5L of milk, in a similar manner that other methods are decreased within the lowest murkiness levels. This effect is reasonable, since beyond 1.5L the degrading effects are so severe that the backscatter component takes up almost all of the dynamic range of the sensor (Figure 9).

The shape reconstruction results using the man-made objects of Figure 6 were next evaluated. The outputs of the PS method are the normals and the scaled albedo ϱ_{sc} of

⁷Contrary to diffuse lighting method of [6] which assumed that a dark point exists in a small neighbourhood around every scene patch, our method requires only a very small number of either dark or infinite points.

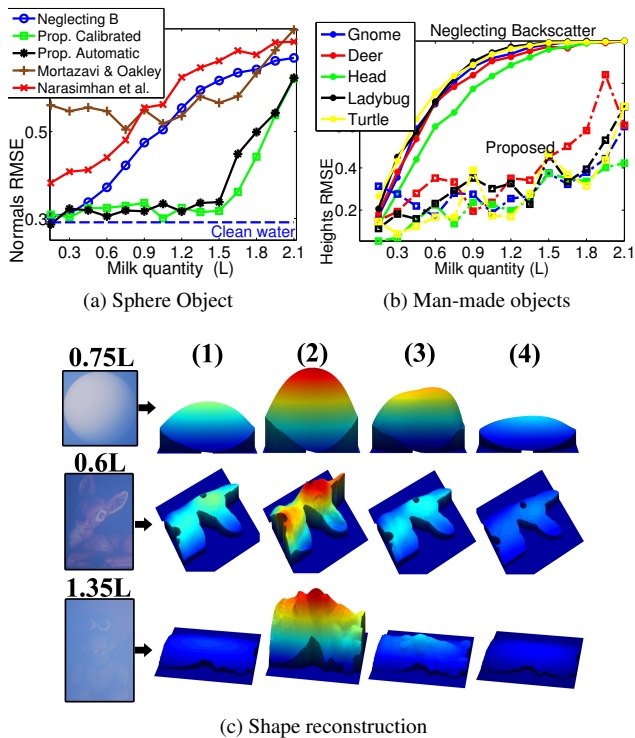


Figure 8: Shape reconstruction results using different methods: (1) neglecting backscatter, (2) our proposed method, (3) Mortazavi and Oakley [9], and (4) Narasimhan et al [12].

each pixel. In order to reconstruct the height map from the respective normals we employ the integration method of [4]. Figure 8b shows the $RMSE$ between the reconstructed height of each object in clean water and that estimated at each scattering level, using our proposed method and that of neglecting backscatter which had the best performance amongst the other methods, while Figure 8c compares the recovered shape of various objects using all methods. As can be observed, our method successfully preserves the reconstructed shape, while the rest of the methods tend to smoothly flatten the result over increased turbidity levels.

In order to have photometric representation of the recovered objects we approximated the albedo of each pixel, under the assumption that the depth variation δz_O of the object surface is small compared to the camera-scene depth. Then, the scaling factor of ρ_{sc} (Section 3.2) is a constant that differs for each channel due to the wavelength-dependent coefficient c . Thus, we rescale each channel by its maximum value in order to recover full contrast. Figure 9 demonstrates our results for various objects and murkiness levels.

6. Discussion and Future Work

When Photometric Stereo is attempted in turbid media, the light is supplied by directional sources close to the camera. This causes significant effects in the backscatter com-

ponent according to the exact pixel-source baseline which have not previously been accounted for. We demonstrate that the scene depth saturation of the backscatter due to ISL makes its direct approximation experimentally appropriate. In order to estimate the smoothly varying backscatter we propose both a calibrated and an automatic way that exploits a very small number of dark points on the surface or infinite points in the surrounding of the imaged object. Our automatic method gives further potential for restoring the poor visibility from a single image. Figure 10 shows the result of backscatter compensation from a single image, captured in murky deep sea port waters using a directional source ⁸.

This backscatter compensation leads to effective PS shape estimation over a very wide range of scattering levels. Our model and setup are consistent with a real world imaging scenario such as robotic inspection underwater, using only 3 light sources. The small source number requirement facilitates rapid PS imaging in a robotic scenario. Further effects might degrade the performance of PS and deserve further research. Additional to the small camera-lights baseline which has a severe effect on the backscatter component, a small camera-scene baseline might require further compensation regarding illumination or perspective projection. Such effects could motivate the combination of our method with depth information from additional cues such as multi-view stereo [7].




Figure 10: Murky and restored port-water images.

Acknowledgements: This work was supported by the contract #270180 of the European Communities FP7 (NOPTILUS). We are very grateful to Prof. Maria Petrou for guiding and inspiring the project work in its early stage, and to Eirini Takoulidou for helping with the experiments.

References

- [1] M. E. Angelopoulou and M. Petrou. Evaluating the effect of diffuse light on photometric stereo reconstruction. *Machine Vision and Applications*, 25(1):199–210, Jan. 2014.
- [2] F. Bruno, G. Bianco, M. Muzzupappa, S. Barone, and A. Rationale. Experimentation of structured light and stereo vision for underwater 3d reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(4), 2011.
- [3] J. Chiang and Y.-C. Chen. Underwater image enhancement by wavelength compensation and dehazing. *Image Processing, IEEE Transactions on*, 21(4):1756–1769, April 2012.

⁸Full contrast was rescaled after backscatter compensation.

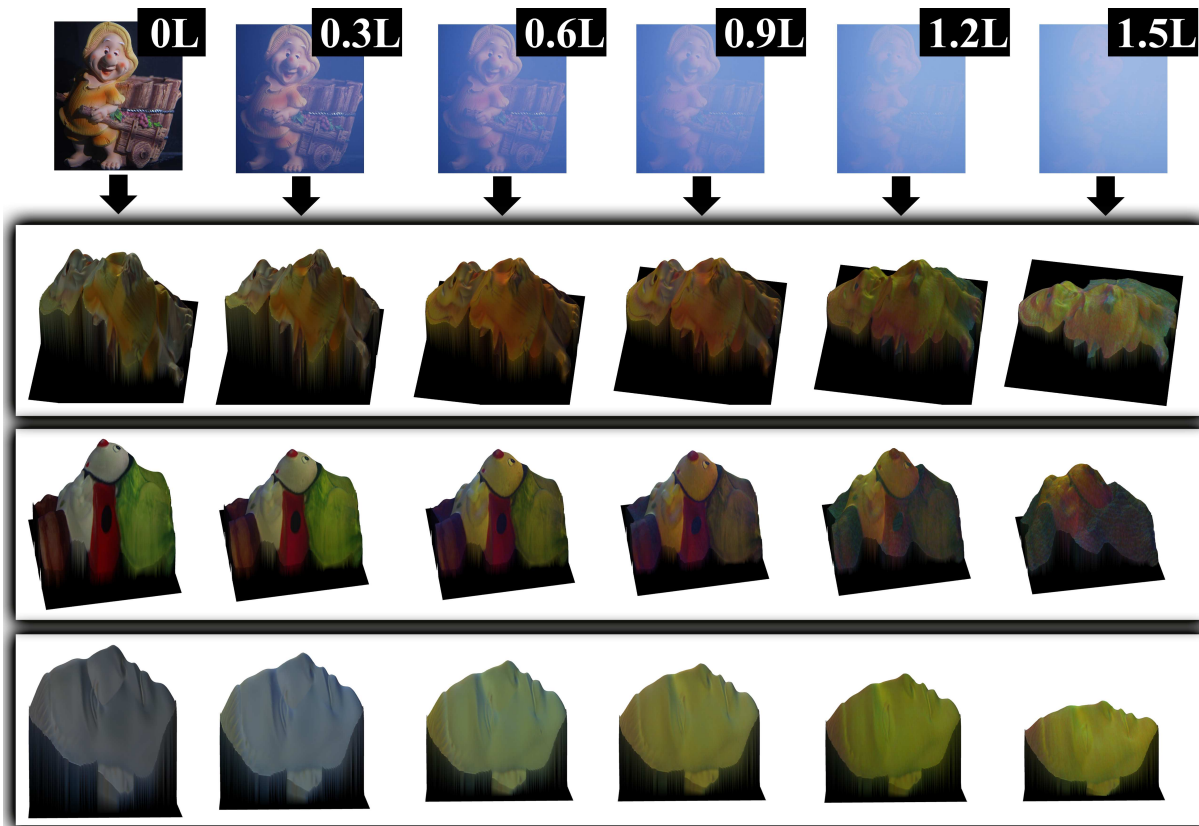


Figure 9: Top row: captured images indicating the respective level of water murkiness. Rows 2 – 4: our reconstruction result.

- [4] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 10(4):439–451, 1988.
- [5] M. Gupta, S. Narasimhan, and Y. Schechner. On controlling light transport in poor visibility environments. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [6] K. He, J. Sun, and X. Tang. Single image haze removal using dark channel prior. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 99(1), 2010.
- [7] T. Higo, Y. Matsushita, N. Joshi, and K. Ikeuchi. A handheld photometric stereo camera for 3-d modeling. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 1234–1241. IEEE, 2009.
- [8] J. S. Jaffe. Computer modeling and the design of optimal underwater imaging systems. *Oceanic Engineering, IEEE Journal of*, 15(2):101–111, 1990.
- [9] H. Mortazavi and J. Oakley. Underwater image enhancement by backscatter compensation. In *International Conference on Modeling, Simulation and Applied Optimization. ICMSAO 2007*. IEEE, 2007.
- [10] H. Mortazavi and J. Oakley. *Mitigation of Contrast Loss in Underwater Images*. University of Manchester, 2010.
- [11] S. Narasimhan and S. Nayar. Contrast restoration of weather degraded images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(6):713–724, 2003.
- [12] S. Narasimhan, S. Nayar, B. Sun, and S. Koppal. Structured light in scattering media. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, 2005.
- [13] S. Negahdaripour, H. Zhang, and X. Han. Investigation of photometric stereo method for 3-d shape recovery from underwater imagery. In *OCEANS’02 MTS/IEEE*, volume 2, pages 1010–1017, 2002.
- [14] Y. Y. Schechner and N. Karpel. Recovery of underwater visibility and structure by polarization analysis. *Oceanic Engineering, IEEE Journal of*, 30(3):570–587, 2005.
- [15] B. Sun, R. Ramamoorthi, S. G. Narasimhan, and S. K. Nayar. A practical analytic single scattering model for real time rendering. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 1040–1049. ACM, 2005.
- [16] J. P. Tarel and N. Hautiere. Fast visibility restoration from a single color or gray level image. In *Computer Vision, 2009 IEEE 12th International Conference on*, 2009.
- [17] T. Treibitz and Y. Schechner. Active polarization descattering. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(3):385–399, 2009.
- [18] T. Treibitz and Y. Schechner. Turbid scene enhancement using multi-directional illumination fusion. *Image Processing, IEEE Transactions on*, 21(11):4662–4667, 2012.
- [19] S. Zhang and S. Negahdaripour. 3-d shape recovery of planar and curved surfaces from shading cues in underwater images. *IEEE J. Ocean. Eng.*, 27(1):100–116, 2002.