Image and Vision Computing 57 (2017) 44–57

Contents lists available at ScienceDirect

Image and Vision Computing

journal homepage: www.elsevier.com/locate/imavis

Near-lighting Photometric Stereo for unknown scene distance and medium attenuation☆

Chourmouzios Tsiotsios [a,*], Andrew J. Davison [a], Tae-Kyun Kim [b]

[a] Department of Computing, Imperial College London, SW72AZ, UK
[b] Electrical and Electronic Engineering Department, Imperial College London, SW72AZ, UK

ARTICLE INFO

Article history:
Received 16 October 2015
Received in revised form 10 July 2016
Accepted 6 October 2016
Available online 15 November 2016

Keywords:
Underwater vision
Photometric Stereo
Near-lighting
Uncalibrated

ABSTRACT

Photometric Stereo in murky water is subject to light attenuation and near-field illumination, and the resulting image formation model is complex. Apart from the scene normals and albedo, the incident illumination varies per-pixel and it depends on the scene depth and the attenuation coefficient of the medium. When these are unknown, e.g. in a realistic scenario where a robotic platform explores an underwater scene (unknown shape and distance) within the dynamic subsea environment (unknown scattering level), Photometric Stereo becomes ambiguous. Previous approaches have tackled the problem by assuming distant-lighting and resorting to external hardware for estimating the unknown model variables. In our work, we show that the Photometric Stereo problem can be determined as soon as some additional constraints regarding the scene albedo and the presence of pixels with local intensity maxima within the image are incorporated into the optimization framework. Our proposed solution leads to effective Photometric Stereo and yields detailed 3*D* reconstruction of objects in murky water when the scene distance and the medium attenuation are unknown. We evaluate our work using both numerical simulations and real experiments in the controlled environment of a water tank and real port water using a remotely operated vehicle.

1. Introduction

In murky maritime environments such as ports and lakes, the human activity is significant. Man-made structures need to be monitored by underwater vehicles in order to prevent hazardous situations [1]. At the same time, scientific fields such as marine biology [2] and archeology [3] are also benefited from the use of autonomous visual systems that are able to monitor and evaluate the condition of important targets underwater.

Scene reconstruction in such environments is a very demanding task. Attenuation and scattering lead to dark and noisy scene appearance. For this reason, underwater vehicles usually have to approximate objects at close ranges in order to increase their visibility. This causes the so-called *near-lighting* effect, where the incident illumination on the scene is non-uniform.

Overall, image degradation is strong in murky water and de-features the captured images and limits the effectiveness of disparity-based methods [4,5] that require special post-processing of the captured images in order to work. Photometric methods on the other hand, such as Photometric Stereo, are based on modeling the image formation and the cause of image degradation and optimizing this for the useful scene orientation and albedo.

However, due to the murky water effects Photometric Stereo becomes complex and non-linear. Apart from the scene orientation and albedo, the scene depth and attenuation coefficient of the medium are also unknown. In order to solve this ambiguous problem, previous approaches have resorted either to simplified models (distant-lighting) that are valid only in limited scenarios or to external equipment for calibrating the additional unknown model variables. The goal of our work is to tackle the *uncalibrated* Photometric Stereo problem in murky water. Contrary to pure air works such as Papadhimitri and Favaro [6] where the term uncalibrated refers to unknown direction of light sources, in our case the term uncalibrated refers to the unknown scene distance and the medium attenuation for robotic inspection underwater where the position of the sources is fixed with respect to the camera.

Our revised Photometric Stereo optimization comprises the standard photometric consistency term between the measured brightness and the predicted brightness according to the image formation model, and two additional constrains. The first exploits the presence

☆ This paper has been recommended for acceptance by Michael Goesele.

* Corresponding author.
E-mail addresses: c.tsiotsios@imperial.ac.uk (C. Tsiotsios), a.davison@imperial.ac.uk (A. Davison), tk.kim@imperial.ac.uk (T. Kim).

http://dx.doi.org/10.1016/j.imavis.2016.10.005
0262-8856/

of pixels that correspond to Local Diffuse Maxima (LDM) regions in the image, i.e. pixels with a local intensity maximum due to shading [7,8]. Specifically, LDM regions correspond to scene points whose normal vector coincides with the direction of the incident illumination. For these pixels the measured brightness is maximized. Thus, we can evaluate whether the estimated orientation for pixels that correspond to local maxima in the image coincides with the estimated incident illumination. When the estimated normal map is correct, these two directions should coincide.

Our second constraint exploits the prior information that the values of the object albedo lie within a limited range. Specifically, due to its clear physical characteristic, the albedo varies from 0 for totally dark scene points to 1 for totally white and thus solutions outside this range should be penalized. At the same time, due to the statistics of natural scenes and man-made objects, it is much more likely that the albedo of the imaged scene lies somewhere in between the extreme values of 0 and 1. Thus, we employ a cost function that penalizes solutions that correspond to unlikely values for the estimated albedo.

We show that our Photometric Stereo approach can recover the normals of the scene when the scene distance and the medium attenuation are unknown. We compare its effectiveness with the result obtained using calibrated distant and near lighting. We perform several numerical simulations considering different scattering, depth and noise levels, and real experiments in the environment of a big water tank using different objects and scattering levels. To the best of our knowledge, we also present the first uncalibrated result for Photometric Stereo in real murky port water using a remotely operated vehicle.

2. Background and challenges

2.1. Light propagation model

The image formation model in murky water was derived analytically in [9,10], and simpler versions have been used later for scene reconstruction [11–14]. Specifically, the image formation model can be expressed using three medium effects: attenuation, back-scattering, and forward-scattering. Attenuation causes the light beams to degrade as they travel through the medium. Back-scattering causes some of the incident illumination on the particles to get scattered toward the sensor causing contrast loss. Forward-scattering corresponds to the scene-reflected light that gets scattered as it travels from the scene to the camera and eventually falls onto neighboring pixel positions on the sensor causing resolution loss.

In this work we solve analytically only for the direct component. Backscatter is an additive component that can be estimated and subtracted using the captured images. Here, we employ the method of Tsiotsios et al. [14] which is based on the assumption that backscatter is saturated with distance, and hence it becomes scene-distance independent. We also assume Lambertian reflectance and neglect forward-scattering as in the previous works of Narasimhan et al. [12], Tsiotsios et al. [14], Treibitz and Schechner [11], and Zhang and Negahdaripour [13]. We discuss the limitations of our assumptions and potential improvements in Section 7.

2.2. Photometric Stereo in murky water

Due to the complex and non-linear form of Photometric Stereo in murky water, previous works have resorted either to simplified models, or to calibration in order to estimate some of the unknown model variables.

Specifically, in [12,14,15] it was assumed that the incident illumination on the scene is distant, considering that the source-scene distance is sufficiently larger than the object size. In that case illumination on all scene points is assumed to be constant (Fig. 1). The distant-lighting approximation leads to a simplified model for the direct component since the light vector for all pixels is constant (per source) and can be calibrated using a white matte sphere [12,14] or estimated automatically [8,16].

When the ratio between the distance and the object size is small, the incident illumination across the surface differs significantly. Neglecting this variation by assuming distant-lighting causes significant errors (Fig. 1). Specifically, for near-lighting, the parts of the scene that are closer to the light sources receive significantly stronger illumination (such as the center of the sphere in Fig. 1). Neglecting this variation reflects error to the estimated normals since Photometric Stereo erroneously attributes the respective variation in the measured intensity to variation in orientation only. This causes low-frequency errors to the estimated shape [17].

The effect of near-lighting is stronger in murky water than in pure air. Apart from attenuation due to inverse-square law, light suffers further attenuation due to the interaction with the particles. Thus, the incident illumination on the scene exhibits an additional dependence on the attenuation coefficient of the medium. Fig. 2 shows the average reconstruction error using a synthetical sphere in simulations when distant-lighting is considered. In pure air, this error depends only on the ratio between the object size and the source-scene depth. In murky water it also depends on the level of water attenuation.

When no model simplifications are used, the incident illumination has to be defined per-pixel and hence it depends on the attenuation coefficient and the 3D position of the scene point that are both unknown in a typical underwater situation. Previous approaches

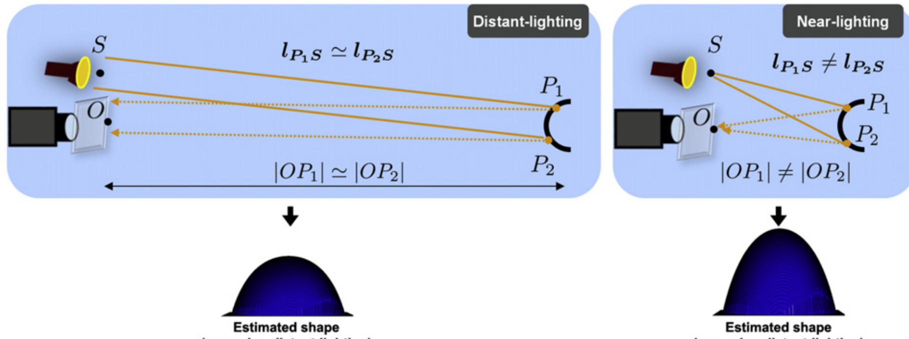

Fig. 1. When the ratio between the distance and the object size is large, distant-lighting approximation is used which assumes that the illumination on the object surface is constant. When the ratio is small, neglecting the near-lighting effect causes notable low-frequency reconstruction error.

have shown that if the average camera-scene distance and the attenuation coefficient are known a priori, the near-light Photometric Stereo can be optimized effectively. Specifically, in [18] additional equipment was used to estimate the total attenuation coefficient and the camera-scene distance. Then, an iterative scheme was used to solve the near-light equations for the unknown normal map, also used in [19].

Our reformulated Photometric Stereo cost function includes a term that exploits the pixels with local intensity maxima. This idea was previously used for pure air Photometric Stereo in [8], for resolving the bas-relief ambiguity when distant-lighting is considered. In [7] it was used in near-lighting conditions in pure air for refining the estimate of the normal map. However, the camera-scene distance was known or calibrated there while we tackle the uncalibrated problem in murky water where both the distance and the attenuation coefficient are unknown.

Our second cost term is based on a likely value range of the scene albedo. In [20] it was assumed that the albedo values of natural scenes follow a normal distribution and this was used for image restoration. In [21] the efficiency of the normal distribution was compared with a uniform and a learnt distribution using real images. It was shown that the normal distribution performs equally with the learnt distribution and it leads to optimization that is less dependent on the initial values. In our work we also employ a simple normal distribution for our photometric reconstruction both in simulations and real experiments. Other distributions such as log-normal [22] which were shown to approximate effectively the statistics of natural images can be also evaluated. Ideally, the albedo distribution would be learnt from real ground-truth data of typical scenes and objects of real underwater environments. Such prior information which is hard to obtain in murky water, was used in [23] for shape-from-shading in pure air.

3. Image formation model

3.1. Light attenuation

When a light beam with irradiance E_0 travels a distance d in a scattering medium it gets attenuated [9] by:

$$E_d = E_0\, e^{-cd}. \tag{1}$$

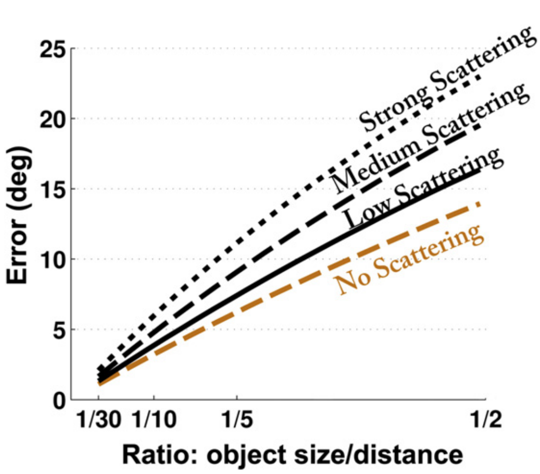

Fig. 2. Reconstruction error due to distant-lighting assumption for different ratios between the object size and the scene distance. The error also depends on the scattering level.

When the light beam is not collimated, inverse-square light fall-off causes attenuation as well. In this case the attenuated light component at a distance d away from a point-source is:

$$E_d = \frac{E_0}{d^2}. \tag{2}$$

3.2. Direct component

The direct component corresponds to the light that travels from the source to the scene, that then gets reflected and reaches the sensor (Fig. 3). The light beam that is emitted from the source is degraded both due to the medium attenuation (Eq. (1)) and due to inverse-square law (Eq. (2)). Thus, the light that reaches the scene is

$$E_{scene} = \frac{I_k}{|PS_k|^2}\, e^{-c|PS_k|}, \tag{3}$$

where I_k is the radiant intensity of the light source, c is the total attenuation coefficient of the medium, and $|PS_k|$ is the distance (vector magnitude) between the light source at position S_k and the scene point at position P. Then, part of this light is reflected from the surface patch. Given that the incident light direction from the source is denoted by the unit vector $\hat{\boldsymbol{l}}_{\boldsymbol{PS_k}}$, and the orientation of the surface patch is denoted by the unit normal vector $\hat{\boldsymbol{n}}$, the reflected amount of light is given by: $E_{\text{reflected}} = E_{\text{scene}}\, \varrho\, \hat{\boldsymbol{l}}_{\boldsymbol{PS_k}} \cdot \hat{\boldsymbol{n}}$. Here, $\hat{\boldsymbol{l}}_{\boldsymbol{PS_k}}$ denotes the direction of the incident illumination and thus it equals the unit vector $\frac{\bar{PS_k}}{|PS_k|}$. The albedo ϱ and the unit vector $\hat{\boldsymbol{n}}$ are usually combined into a single vector $\boldsymbol{n} \equiv \varrho\hat{\boldsymbol{n}}$, and the reflected light can be rewritten as $E_{\text{reflected}} = E_{\text{scene}}\, \hat{\boldsymbol{l}}_{\boldsymbol{PS_k}} \cdot \boldsymbol{n}$.

Then, the scene-reflected light is attenuated along the distance $|OP|$ as it travels from the scene point to the sensor: $E_{reflected} e^{-c\,|OP|}$. The final light component on the pixel is

$$E_k = \frac{I_k}{|PS_k|^2}\, e^{-c(|PS_k|+|OP|)}\, \hat{\boldsymbol{l}}_{\boldsymbol{PS_k}} \cdot \boldsymbol{n}. \tag{4}$$

The sensor typically measures an additional backscatter component B_k (details in Sections 2 and 7), which can though be estimated directly from the images and subtracted.

4. Calibrated near-light Photometric Stereo

Photometric Stereo aims at estimating the normals and albedo for every pixel given the measured direct component from every source. Let us examine the unknown variables.

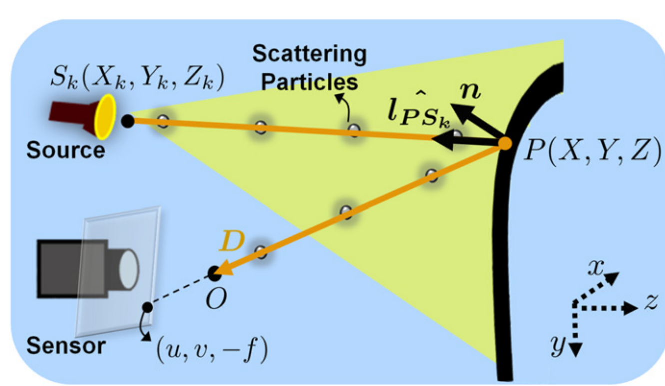

Fig. 3. The direct light component depends not only on the scene orientation and albedo, but also on the scene depth and the attenuation coefficient since light is degraded as it travels through the medium.

The coordinates of the scene point can be expressed with respect to the pixel coordinates as $P = (X, Y, Z) = \left(\frac{uZ}{f}, \frac{vZ}{f}, Z\right)$. Here f is the focal length of the camera which is considered known a priori and (u, v) are the pixel coordinates on the sensor (Fig. 3). The coordinates of every light source $S_k = (X_k, Y_k, Z_k)$ are also known since the sources are normally fixed with respect to the camera on a robotic platform underwater. Then, all of the unknown vectors in Eq. (4) are a function of the scene depth Z only:

$$|PS_k| = \sqrt{\left(X_k - \frac{uZ}{f}\right)^2 + \left(Y_k - \frac{vZ}{f}\right)^2 + (Z_k - Z)^2}, \tag{5}$$

$$|OP| = \sqrt{\left(\frac{uZ}{f}\right)^2 + \left(\frac{vZ}{f}\right)^2 + Z^2}, \tag{6}$$

$$\hat{\boldsymbol{l}_{\boldsymbol{PS_k}}} = \frac{\left(X_k - \frac{uZ}{f}, Y_k - \frac{vZ}{f}, Z_k - Z\right)}{\sqrt{\left(X_k - \frac{uZ}{f}\right)^2 + \left(Y_k - \frac{vZ}{f}\right)^2 + (Z_k - Z)^2}}. \tag{7}$$

Overall, Eq. (4) depends on the scene depth Z and the normal vector $\boldsymbol{n}$ (per-pixel unknowns), and the total attenuation coefficient c of the medium (global unknown). First, we describe how the normal map can be estimated when the total attenuation coefficient is calibrated and the mean camera-scene distance is known a priori.

4.1. Attenuation coefficient

We estimated the total attenuation coefficient c of the medium using a simple calibration step without the use of external hardware as in [18]. A flat white matte canvas was inserted into the murky water at a known depth away from the camera and perpendicular to its optical axis. First, the backscatter component was estimated and subtracted using the method proposed in [14].

The unit normal vector of all scene points on the canvas equals $\hat{\boldsymbol{n}} = (0, 0, -1)$ as it coincides with the optical axis of the camera sensor, and the albedo of the white points can be assumed to be roughly $\varrho \simeq 1$. Since the canvas is at known depth, the values of $|PS_k|$, $|OP|$, and $\hat{\boldsymbol{l}_{\boldsymbol{PS_k}}}$ in Eq. (4) are also known. Then the only unknown is the total attenuation coefficient which can be estimated by minimizing the cost function:

$$c' = \arg\min_c \sum^{noP} \left(E_k - \frac{I_k}{|PS_k|^2} e^{-c(|PS_k| + |OP|)} \hat{\boldsymbol{l}_{\boldsymbol{PS_k}}} \cdot \boldsymbol{n}\right)^2. \tag{8}$$

Here noP denotes the number of pixels. Only one source is enough for solving Eq. (8) in the least-square sense.

4.2. Scene distance

As soon as the total attenuation coefficient c of the medium is estimated, the problem comes to the estimation of the per-pixel normal vector and depth. In previous PS approaches in pure air [6,7], it was shown that these can be estimated via an iterative optimization scheme as soon as the average scene depth Z_c (the average distance between the camera and the object surface) is known.

It is initially assumed that all points correspond to the same depth Z_c (distant-lighting approximation). In this case all scene points receive constant-known illumination (Eq. (7)) and a first version of the object's surface normals is estimated. Then, the estimated normal vectors are integrated in order to recover the height map of the object. Using this and the known average scene distance Z_c, the per-pixel depth and lighting vector can be estimated and Photometric Stereo yields an improved version of the normal vectors as the varying depth of the scene points is taken into account. Then the lighting vectors are estimated again based on the new recovered shape. This procedure is iterated until the algorithm converges to the final object shape, i.e. when the difference between the estimated normal maps for two successive iterations is below a given value. Algorithm 1 overviews this method.

Algorithm 1. Iterative near-light Photometric Stereo.

initialization;
$P \leftarrow (0, 0, Z_c)$;
$\boldsymbol{l}_{\boldsymbol{PS_k}} \leftarrow (X_k, Y_k, Z_k - Z_c),\ k \in [1, noS]$;
$\boldsymbol{L} \leftarrow [\boldsymbol{l}_{\boldsymbol{PS_1}} \dots \boldsymbol{l}_{\boldsymbol{PS_k}}]$;
repeat
  $\boldsymbol{n} \leftarrow \boldsymbol{L}^{+}\boldsymbol{E}$ (Least-square estimate of normals) ;
  $H \leftarrow$ Integration of $\boldsymbol{n}$ [Frankot and Chellappa (1988)] ;
  $Z \leftarrow Z_c - H$;
  $P \leftarrow (\frac{uZ}{f}, \frac{vZ}{f}, Z)$;
  $\boldsymbol{l}_{\boldsymbol{PS_k}} \leftarrow (X_k - \frac{uZ}{f}, Y_k - \frac{vZ}{f}, Z_k - Z),\ k \in [1, noS]$;
  $\boldsymbol{L} \leftarrow [\boldsymbol{l}_{\boldsymbol{PS_1}} \dots \boldsymbol{l}_{\boldsymbol{PS_k}}]$;
until $\boldsymbol{n}_{i+1} - \boldsymbol{n}_i < 10^{-6}$;

Fig. 4 shows the estimated sphere shape using the described iterative near-light scheme. The initial estimate is obtained using the distant-lighting assumption which neglects the depth variation of the surface, and it suffers from low-frequency errors (Section 2.2). This is then used for estimating the per-pixel depth and incident illumination which yield a new-improved estimate of the object shape. The procedure converges after a few iterations (Fig. 4.b).

5. Uncalibrated solution

In the previous section it was shown that the normal map of the scene can be estimated when the values of the total attenuation coefficient c and the scene distance Z_c are known a priori. In this section we tackle the uncalibrated near-light problem.

5.1. Photometric consistency

Consider the Photometric Stereo system of equations. The measured intensities are obtained using different light sources. Then, given the image formation model, the normal vector is estimated by minimizing the difference between the measured and the modeled brightness values for every pixel. This is usually described as photometric consistency, since the objective function dictates that the estimated normal predicts a brightness according to the image formation model that is consistent with the measured brightness for every source:

$$g_{phot}(\boldsymbol{n}) = \sum_{k=1}^{noS} \left(E_k - E'_k(\boldsymbol{n})\right)^2, \tag{9}$$

where $E'_k(\boldsymbol{n})$ is the brightness that is predicted by the image formation model and E_k is the measured brightness from a light source k. In the calibrated scenario the only unknown is the normal map $\boldsymbol{n}$ in every iteration and hence the final linear system of equations can be solved using 3 sources. As soon as at least 4 sources are used the problem is over-determined (as the normal has only 3 unknowns) and the cost value corresponds to the residual of the least-squares solution.

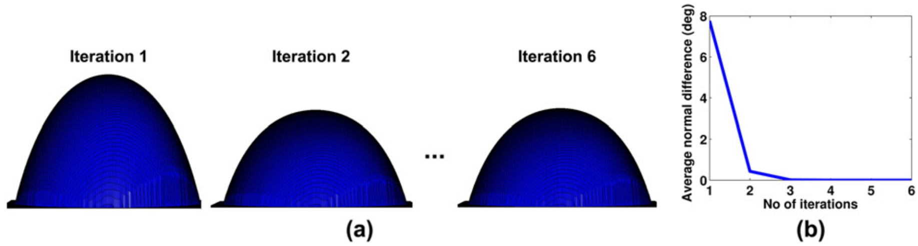

Fig. 4. (a) The reconstructed shape of a sphere using simulations and Algorithm 1. Initially the shape suffers from error due to the distant-lighting assumption. Then the result is gradually refined by updating the depth value per pixel. (b) The algorithm converges within a few iterations.

In the uncalibrated case, Z_c and c are also unknown apart from the normal map $\boldsymbol{n}$. Thus, the reformulated objective function of the photometric consistency is given as:

$$g_{phot}(\boldsymbol{n}, Z_c, c) = \sum_{k=1}^{noS} \left(E_k - E'_k(\boldsymbol{n}, Z_c, c)\right)^2. \tag{10}$$

Every potential combination of Z_c and c in Eq. (10) leads to the respective calibrated case. Thus, the photometric cost value for every (Z_c, c) corresponds to the residual of the least-squares solution as described above.

Fig. 5 shows the estimated values of the photometric cost function for different Z_c and c. Specifically, we simulated the image formation model for a sphere object using 8 light sources and some Gaussian sensor noise. This example corresponds to scene depth $Z_c = 0.8$ m and $c = 1$ m^{-1} (indicated by the black dot). As these are the true values, we expect the photometric objective function to be minimized at that point. The problem though is ambiguous as there are several potential solutions that correspond to a very small cost.

We performed several numerical experiments for different distance, scattering, and object characteristics. The photometric objective function was ambiguous in all cases regardless of the number of sources. Next, we describe how the problem can be determined using additional constraints.

5.2. LDM prior

We take advantage of surface patches whose normal vector coincides with the lighting direction in order to determine uncalibrated near-light PS in murky water. As described in [7,8,24], for continuous diffuse surfaces with a constant albedo, such patches create a local maximum in the measured intensity. Specifically, for these pixels the dot product between the incident illumination and the normal vector ($\hat{\boldsymbol{l}} \cdot \hat{\boldsymbol{n}}$) is 1 as their directions coincide (Eq. (4)), and the scene-reflected light is maximized (Fig. 6). In order to emphasize that the intensity maximum for such regions is attributed to orientation-lighting direction coincidence only and not to a complex reflectance, they are described as Local Diffuse Maxima (LDM).

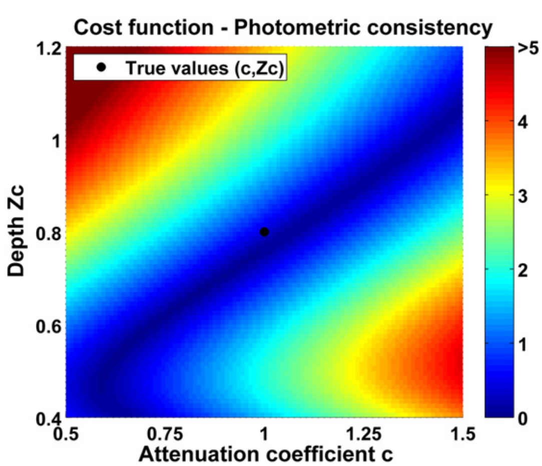

Fig. 5. The cost of the photometric objective function for different values of the scene distance Z_c and the attenuation coefficient c. The true values that were used in the simulations correspond to the dark dot in the graph. The problem is ambiguous as there are different combinations of the unknown variables that predict brightness close to the measured one.

In order to examine whether pixels with a local intensity maximum correspond always to LDM regions we need to investigate the image formation model (Eq. (4)). The measured brightness at every pixel is a function of the normal vector which comprises the albedo ϱ and the direction of the unit normal vector $\hat{\boldsymbol{n}}$, the scene depth Z, and the attenuation coefficient c. c is constant for all pixels in the image since we are dealing with the uniform murky water medium. Thus, a local maximum in the measured brightness within the image space can be attributed only to change in ϱ, $\hat{\boldsymbol{n}}$ and Z.

Iwahori et al. [24] assumed that the object has a constant albedo ϱ and thus an LDM can be attributed only to variation in surface normal $\hat{\boldsymbol{n}}$ and depth Z. Specifically they examined whether a local intensity maximum can be created because the distance between the source and a scene point is minimum but its orientation doesn't coincide with the lighting direction. However they proved that this is not possible within a local neighborhood of a smooth surface, since the minimum distance between the light source and the surface is always at a point that is perpendicular to the surface.

Favaro and Papadhimitri [8] showed that LDM pixels can be effectively detected even when the object has a varying albedo. This was achieved by simply discarding pixels that had a locally maximum intensity in more than one images. As the illumination direction changes from one image to the other in Photometric Stereo, a pixel

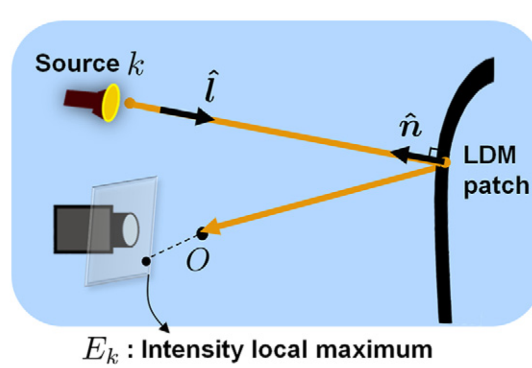

Fig. 6. The measured intensity for scene points whose orientation coincides with the direction of the incident illumination is locally maximum (Local Diffuse Maxima) as the dot product between $\hat{\boldsymbol{l}}$ and $\hat{\boldsymbol{n}}$ is 1.

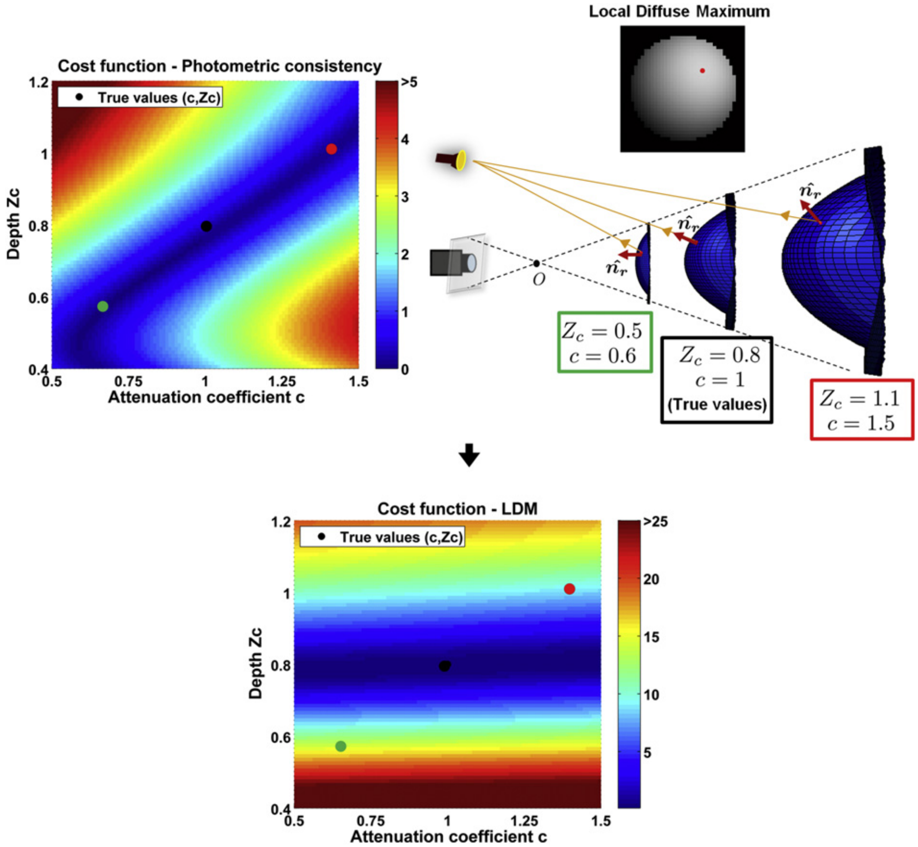

Fig. 7. Different Photometric Stereo solutions that have a low photometric cost are shown in the first row. For the true solution (black), the pixel with a local intensity maximum in one of the captured images corresponds to a normal vector estimate that coincides with the lighting direction. For the erroneous solutions (green, red), the directions of these vectors differ. The cost function which measures the level of agreement between the normal and lighting directions for LDM pixels is shown below. Erroneous solutions with a low photometric cost (green, red) are now penalized.

should not have a maximum intensity in more than one images when it corresponds to an LDM region as its normal vector can be oriented toward only one of the light sources. Otherwise, the maximum is attributed to change in albedo and it can be discarded as outlier.

Consider the uncalibrated near-light Photometric Stereo problem in murky water. Fig. 7 shows potential estimates that have a small photometric cost value (described in the previous section). The cost value of every solution is marked by a different color. For the true solution (black), the LDM pixel corresponds to a normal vector estimate that coincides with the incident illumination direction, as described above. For other erroneous solutions (green, red), these vectors differ. Thus, we formulate the following objective function

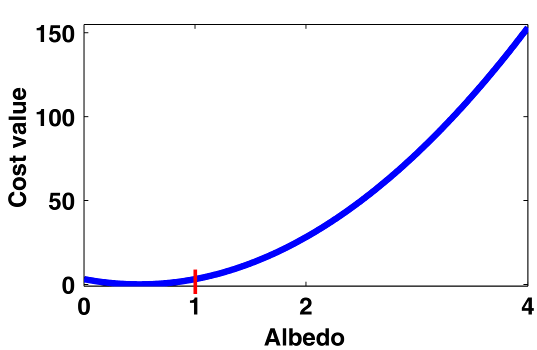

Fig. 8. Our employed cost function for the estimated albedo corresponds to the negative log of a normal distribution. The cost is increased rapidly beyond estimated albedos > 1 that have no physical meaning.

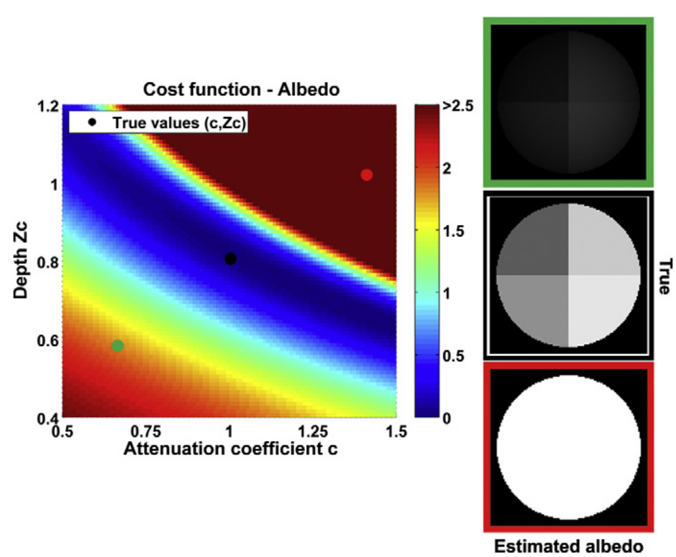

Fig. 9. The cost function for the estimated albedo. The true values of Z_c, c are denoted by the black color. Other erroneous solutions for Z_c, c with a small photometric cost (Fig. 5) correspond to very dark (denoted by green) or unnaturally high — greater than 1 (denoted by red) albedos and are penalized. In this way, the uncalibrated near-light Photometric Stereo problem is constrained using a prior distribution for the absolute value of the albedo.

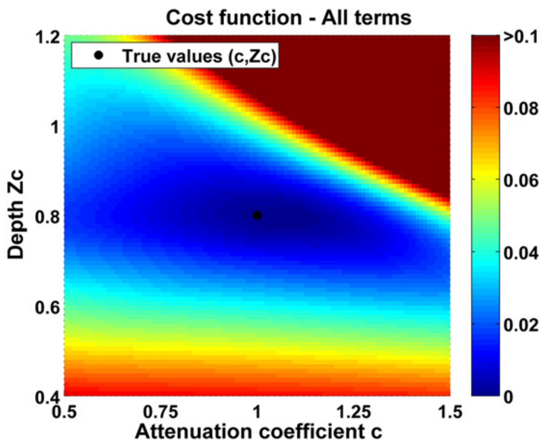

Fig. 10. The final cost function for uncalibrated near-light Photometric Stereo in murky water (Eq. (14)).

which measures the level of agreement between the estimated normal vector and the lighting direction for detected LDM pixels:

$$g_{LDM}(\boldsymbol{n}, Z_c, c) = \sum_{k=1}^{noS} \sum_{r=1}^{noLDM} \operatorname{acos}(\hat{\boldsymbol{n}}_{\boldsymbol{r}} \cdot \hat{\boldsymbol{l}}_{\boldsymbol{r}}). \tag{11}$$

Here *noLDM* is the number of detected LDM pixels in every image, and $\hat{\boldsymbol{n}}_{\boldsymbol{r}}$, $\hat{\boldsymbol{l}}_{\boldsymbol{r}}$ are the unit vectors of the normal and the incident illumination for the LDM pixel.

In this way, shape estimates for which the two directions do not coincide at the detected LDM pixels are penalized. Fig. 7 shows the estimated LDM cost function for the same example. The true solution has a very low cost while erroneous solutions that had a low photometric cost are now penalized.

5.3. Albedo prior

The physical characteristic of the albedo offers an additional constraint. Specifically, the albedo lies between 0 for totally dark objects, to 1 for totally white. Thus, any solution for Z_c, c and $\boldsymbol{n}$ that corresponds to $\varrho = |\boldsymbol{n}| > 1$ should be penalized. At the same time, the albedo of natural scenes [21] or man-made objects [23] exhibits statistical regularities. For example, it is very unlikely that an object is totally black (in which case Photometric Stereo is not possible anyway) or totally white. For this reason, we add another cost term to our optimization that penalizes estimated albedos according to a previously determined distribution.

Fig. 11. Images of different albedo maps used for the simulations.

As in [20,21] we assume that the albedo follows a normal distribution:

$$f_N(\varrho|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(\varrho-\mu)^2}{2\sigma^2}}, \tag{12}$$

where μ is the mean and σ the standard deviation of the distribution. Maximizing the likelihood that the estimated albedo follows this distribution is equivalent to minimizing the sum of the cost expressed as the negative log of Eq. (12):

$$g_{alb} = \frac{1}{noP} \sum^{noP} -\log f_N(\varrho|\mu, \sigma). \tag{13}$$

Fig. 8 shows our employed cost function for the absolute value of the estimated albedo. We used $\mu = 0.5$ and $\sigma = 0.2$ in both the simulations and real experiments in this work. As Fig. 8 shows, the cost is small and varies little within the range of $\varrho \in [0, 1]$, and then it is increased rapidly penalizing solutions that correspond to unnatural albedos. Assuming that there are ground-truth measurements about the true albedo statistics within the underwater environment, the prior distribution can be learnt as in [23].

Fig. 9 shows the estimated albedo cost function for the simulated example of the previous sections. The true solution (black) is assigned a small cost as it corresponds to estimated albedos that are likely to appear according to the global distribution. Under-estimated values of Z_c, c (green) correspond to under-estimated, very dark albedos, and over-estimated values of Z_c, c (red) correspond to over-estimated albedos that are even greater than 1. Thus, erroneous estimates of Z_c, c are penalized according to the cost function of Eq. (13).

5.4. Cost function

The final cost function consists of all terms; the standard photometric cost function (Eq. (10)), the cost function that penalizes the level of disagreement between the estimated orientation and incident illumination for LDM pixels (Eq. (11)), and the cost function that penalizes estimated albedos that lie outside a likely range of values (Eq. (13)):

$$g = g_{phot} + \lambda_1\, g_{LDM} + \lambda_2\, g_{alb}. \tag{14}$$

Using only the standard photometric consistency term (Fig. 5) is ambiguous, while employing the proposed cost function determines the problem estimating a unique maximum near the true solution (Fig. 10).

6. Results

We have performed a large number of experiments using simulations and real murky water, comparing the result of our method (uncalibrated near-lighting) with calibrated distant and near lighting. The values of λ_1 and λ_2 in Eq. (14) were selected so that the minimum and maximum values of all three cost terms lie within the same order of magnitude ($\lambda_1 = 2.5 \times 10^{-3}, \lambda_2 = 2 \times 10^{-2}$). The cost function was optimized using the MATLAB function fmincon which uses the interior-point method described in [25].

The LDM prior is based on pixels with local intensity maxima that are attributed to coincidence of the lighting direction and the normal vector. In order to reject outliers we employ the method of Favaro and Papadhimitri [8]. Specifically, maxima that are present in more than one images are rejected as they are attributed to changes in albedo (Section 5.2), the images are filtered with a low-pass filter in

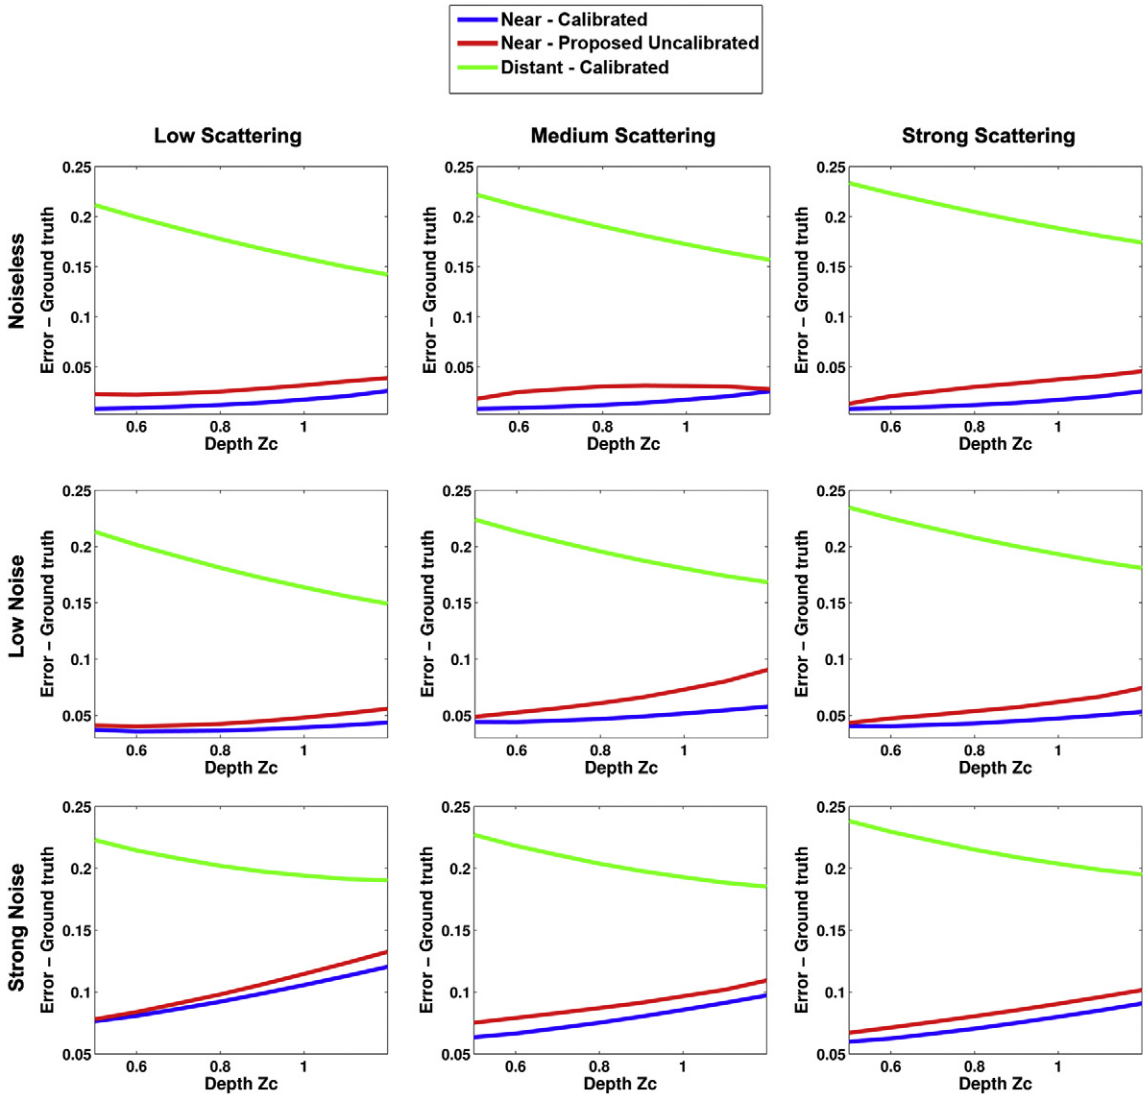

Fig. 12. Numerical simulations for different scattering, depth, albedo, and noise characteristics. Our proposed uncalibrated near-light Photometric Stereo method outperforms the distant-lighting solution and yields effective reconstruction, similar to the one obtained with the calibrated method in all considered scenarios.

order to take account of additive Gaussian noise before detecting the maxima, and maxima with a small or saturated intensity are rejected.

6.1. Simulations

In our numerical simulations we considered different scattering, distance and noise characteristics. Specifically, we simulated the image formation model in murky water considering that the object is a sphere with a big diameter (0.4 m) so that the near-light effect is strong. Three different scattering levels were tested ($c = 0.8, 1.3$, and 2 m^{-1} were used for low, medium, and strong scattering levels, respectively [26]) and in every level the sphere was imaged from a wide range of distances (from 0.5 to 1.2 m) and for different levels of sensor noise (additive Gaussian noise with $\sigma = 0, 0.01, 0.025$ was used for the no, low, and strong noise levels, respectively). In every case, it was assumed that the sphere has blocks with constant albedo. Both the blocks and the albedo value in every block were randomly selected considering uniform distributions. Specifically, the number of blocks varied between 1–100 and the albedo value of each block varied between 0.1–1(uniformly distributed). Fig. 11 shows different random instances of the simulated albedo. Our albedo cost function considers a normal prior distribution, and hence in this way we tested how well it performs when the real albedo distribution is different.

Fig. 12 shows the simulation results. The calibrated and uncalibrated near-lighting methods outperform distant-lighting. This is expected as we considered a big object with respect to the camera-scene distance, and thus the illumination was near-field. At the same time, our proposed method yielded similar reconstruction results with the calibrated method in all scenarios. The normal distribution proved effective even though the simulated true albedo was selected from a different distribution (uniform). The performance of both the calibrated and uncalibrated methods was slightly degraded for larger camera-scene distances. This can be explained by the fact that we considered a fixed sensor resolution regardless of the depth, i.e. for larger distances the sphere was imaged by a smaller amount of pixels. This increased the impact of noise to the reconstruction since

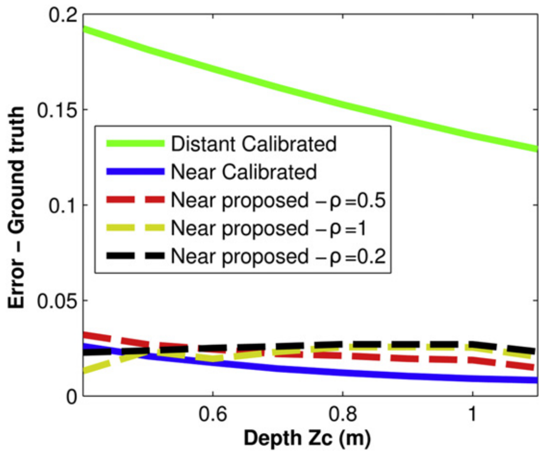

Fig. 13. Red, yellow and black dashed lines correspond to the result of our uncalibrated algorithm when the object has a true albedo of $\varrho = 0.5$ (gray), 1 (totally white), and 0.2 (very dark), respectively. Although our normal distribution albedo prior corresponds to $\mu = 0.5$, the result is similar in all cases.

some algorithmic parts such as the normal integration that yields the height map are benefited from high resolution.

In order to examine the impact of the assumed normal distribution for the albedo prior, we performed an additional experiment (Fig. 13), where the object had a single albedo equal to: 0.5 (dashed red), 1 (dashed yellow) or 0.2 (dashed black). Despite the fact that our normal distribution has a mean value of $\mu = 0.5$, the result when the true albedo of the object was exactly 0.5 (dashed red) was similar to the other cases (different due to small noise variation). This is reasonable, as the photometric consistency and the LDM prior are also used for optimization and furthermore the cost of the albedo prior only varies slightly for physically valid albedo values between $[0, 1]$ and penalizes mostly unnatural albedos > 1 (Fig. 8).

6.2. Water tank experiments

For the real experiments we used a controlled environment that can be seen in Fig. 14. This consists of a big water tank and a metal platform for mounting the light sources and the camera. The light sources were light-emitting diodes (LED) with a narrow beam angle to reduce backscatter [14,27]. A Nikon D60 camera with a AF-S Nikkor $35mmf/1.8G$ lens was immersed into the water, enclosed in an underwater housing. In order to simulate the scattering conditions, clean tap water was diluted with different amounts of milk as in [4,12,14]. The distance between the platform and the centroid of the objects was manually adjusted to 0.7 m. We selected objects with different albedos (for example the shell has a non-uniform albedo and is significantly brighter than the container object), but we used the same albedo distribution in all cases (Section 5.3). The object parts were manually cropped in the captured images. In all experiments the optimization algorithm converged to a solution within approximately 1 min for 800 × 600 pixels images, using an Intel Core i5-2410M CPU @ 2.3 GHz and a MATLAB implementation using the interior-point based fmincon function. In order to reduce the impact of shadows and highlights, we rejected pixels with intensity < 0.03 and > 0.95.

Figs. 15–17 show the reconstruction results for all objects in different scattering levels. Our proposed uncalibrated near-lighting solution can be compared with the reconstruction result using a depth sensor in pure air [28], and the Photometric Stereo output considering calibrated distant and near-lighting. In all cases the backscatter component was estimated and subtracted first using the method of Tsiotsios et al. [14], by imaging a black canvas at the object distance.

In all experiments it can be noticed that due to the small ratio of the scene depth with respect to the object size, employing the distant-lighting assumption introduces strong error to the reconstruction. This is stronger for the low-frequencies of the estimated shape and it is evident when compared with the reconstruction using a depth sensor in pure air. Our uncalibrated approach takes account of the near-field illumination and estimates a detailed reconstruction similar to that obtained using the calibrated method. The low-frequency error that is evident in distant-lighting is mitigated, but the high-frequency details are rich. In Figs. 15 and 16 we can see the error-difference maps between: a) the calibrated near and distant lighting, and b) the calibrated near and our proposed uncalibrated near lighting approach. The difference map corresponds to the per-pixel difference in normal orientation degrees. The calibrated distant-lighting differs significantly with the calibrated near-lighting reconstruction, while our proposed algorithm differs only by a few degrees across the map.

Tables 1 and 2 compare the calibrated and estimated values for Z_c and c using our uncalibrated algorithm. Although the calibrated values are not strictly ground-truth data and they are also subject to approximation errors, comparing them with our estimated values indicates the ability of our algorithm to yield similar solutions with the calibrated case without any prior knowledge about the imaging characteristics. The difference of our average estimated scene distance for all cases/objects differs by 0.025 m from the real distance, while the max difference (shell object, high scattering) was 0.099 m. The respective average difference for the estimated attenuation coefficient for low scattering is 0.11 m^{-1} and the max difference is 0.382 m^{-1}, while the average and max differences for high scattering are 0.2319 m^{-1} and 0.546 m^{-1}. Fig. 18 shows the respective difference maps for the case when the estimated c value differed most from the calibrated value (corresponding to the container experiment for high scattering). Our algorithm yields a result very similar to the near-calibrated case.

6.3. Port water experiment

A Photometric Stereo system was installed in a remotely operated vehicle and a barrel object was imaged in real port waters in Porto, Portugal. Fig. 19 shows the platform operating in murky water. The camera (Lumenera Le-165 with a Tamron 219-HB lens) was enclosed in a waterproof housing with a flat port and four LED sources were mounted around it in a symmetric arrangement. The camera and the lights were synchronized so that every frame was taken having only one of the LED sources on. A fifth frame with all sources off was captured and subtracted from the rest of the frames, as this corresponds to the additive environment light [2]. The backscatter was approximated using the method of Tsiotsios et al. [14,29] assuming that backscatter is saturated. The ROV was driven to sit on the ground in order to avoid any image misalignments between successive frames.

Fig. 20 shows one of the four captured images of the barrel object. The water condition was very murky and the resulting image degradation was strong. The near-lighting effect is also evident, since the size of the barrel is big and it was imaged from a small distance. This is also reflected in the reconstructed shape when distant-lighting was employed, which exhibits a low-frequency peak in the middle part (Section 2). This effect is significantly mitigated using our proposed near-lighting algorithm without prior knowledge

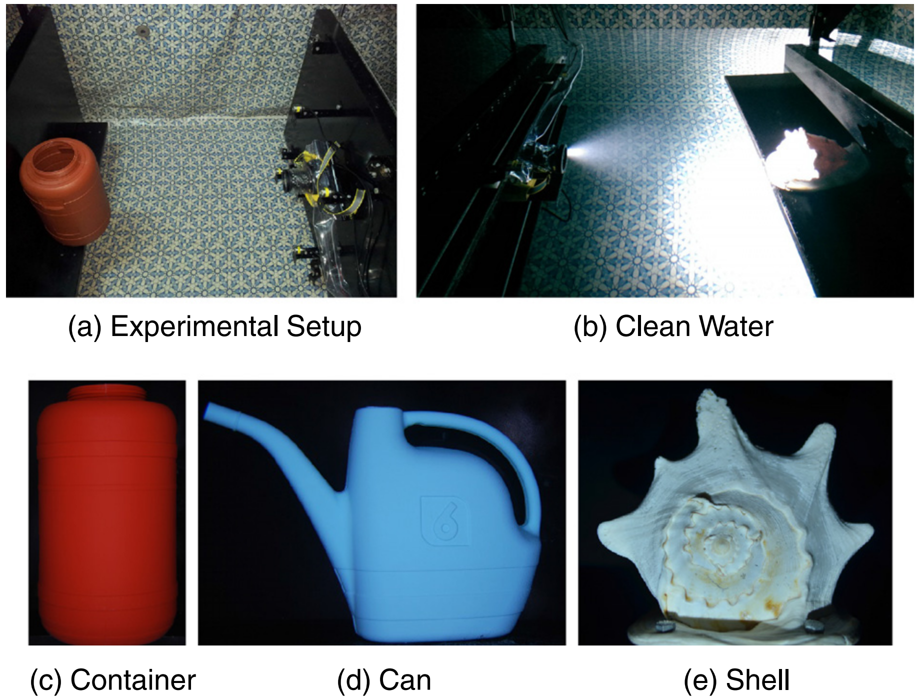

Fig. 14. First row: the experimental setup consisting of a camera and four LED sources in a big water tank. Second row: the objects that were used for the experiments.

about the distance or the scattering level. Although we noticed no strong specularities (Fig. 20), since the barrel corresponds to complex reflectance some shape artifacts are still evident possibly due to weak highlights. In Section 7 we discuss the limitations of our method and potential extensions. The estimated values for Z_c and c were equal to 1.06 m and 2.63 m^{-1} for this experiment.

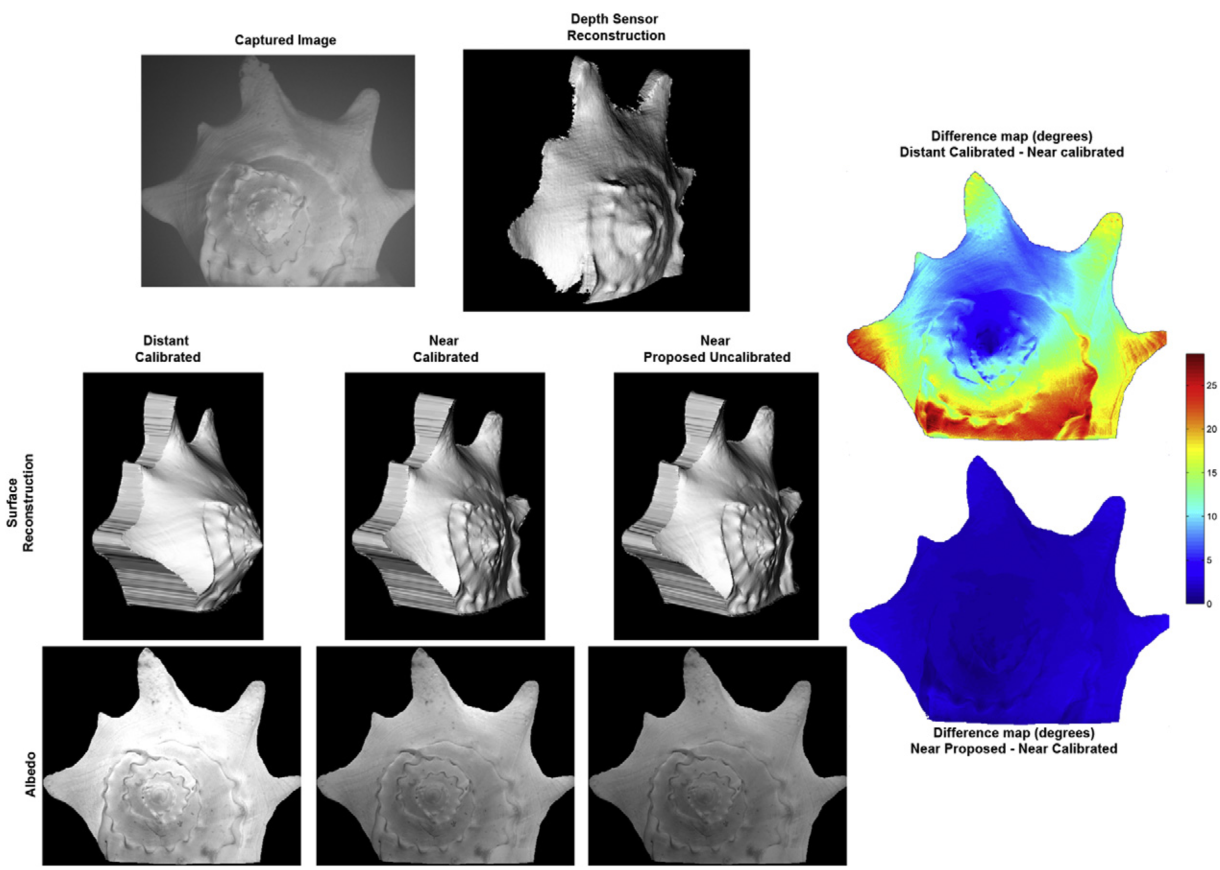

Fig. 15. Top left: one of the four captured images using our Photometric Stereo system in murky water. Top right: the reconstructed surface in pure air using a depth sensor [28]. Middle row: the reconstructed and relit surfaces using distant-lighting, calibrated near-lighting, and uncalibrated near-lighting, respectively. Bottom row: the estimated albedo using the respective methods. Far right column: the difference map in normal map degrees between near-calibrated and distant-calibrated, and near-calibrated and our proposed uncalibrated result, respectively.

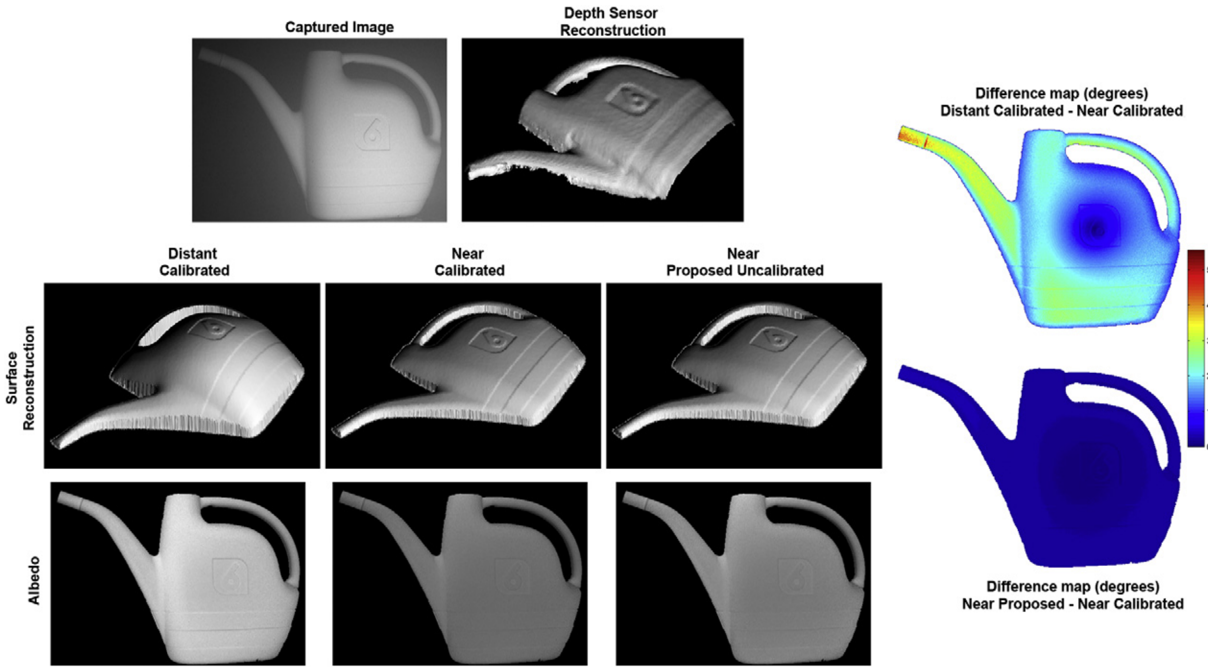

Fig. 16. Reconstruction results for the water can object.

7. Limitations and future work

Backscatter was estimated using the method described in [14] which is based on the assumption that backscatter is saturated and becomes scene-depth independent. The validity of this method has been numerically evaluated for a wide range of imaging conditions and scene depths in [29]. In unknown conditions where the validity of this assumption is not guaranteed a-priori, the method of Tsiotsios et al. [30] can be used to adopt the most valid backscatter approximation method. Otherwise, backscatter can be compensated using polarizers on the sources and the camera [11], eliminated directly for objects that fluoresce [31], or degraded using a large separation between the camera and the sources [27,32]. Our work is complementary with such approaches since backscatter is an additive component to the sensor.

Forward-Scattering effects have been neglected in our work as previous photometric studies indicated that the effect of the resulting resolution loss is small compared with contrast loss from backscatter and attenuation [13,33-35]. However the recent work of Murez et al. [31] showed that in strong scattering conditions, compensating for forward-scattering decreases significantly the reconstruction error. Including the forward-scattering compensation into our model is an interesting future direction, especially in conditions of very murky port water.

Reflectance was assumed to be Lambertian in our work. This limits the effectiveness of our approach, however future work can evaluate how methods for tackling non-Lambertian objects that have been proposed for pure-air can be adapted to our work. To the best of our knowledge, there has been no effort to model complex reflectance for underwater photometry. Some studies have indicated that specularities have a low effect underwater since the refractive index of the medium is similar to that of the reflecting wet surfaces [34,36]. Otherwise, algorithmic approaches like Barsky and Petrou [37] can be used to detect pixels that correspond to shadows or highlights and omit them from the optimization framework. In our work, we used a naive outlier rejection method, where very bright or dark pixels were neglected. Other approaches like Wu et al. [38] do not detect particular pixels, but perform a global matrix factorization to reject outliers. Systematic approaches can also be used to reduce or detect specularities, i.e. using polarizers on the source and camera [33,39], or using a system with moveable light sources that can evaluate the effectiveness of the photometric model in unknown conditions [30].

Color has not been considered in our work, in the sense that we worked on a single color channel. This makes our work valid for multispectral cameras [35]. Since the attenuation coefficient c is wavelength-dependent, in the case of RGB imaging a different parameter value should be estimated for each channel.

Light arrangement in our work has been symmetric around the camera as in [12,14]. However, adopting different setups does not limit the assumptions and validity of our formulation and future work can investigate the optimal light configuration for Photometric Stereo in murky water.

8. Conclusions

When the object size is large compared with the camera-scene distance, the illumination on the scene is near-field. Specifically, the incident illumination on the imaged surface differs significantly according to the 3D position of every scene point with respect to the light source. This effect is stronger in murky water since the attenuation coefficient of the medium introduces an additional scene-dependent factor for the incident illumination. Neglecting this characteristic (when distant-lighting model is employed) results in erroneous shape estimation. On the other hand, optimizing the near-lighting model for all of its unknown variables is hard. In this work, we showed that additional constraints can be introduced that take advantage of pixels with local intensity maxima or prior information about the scene albedo. This leads to effective Photometric Stereo reconstruction even when the camera-scene distance and the total attenuation coefficient of the medium are unknown.

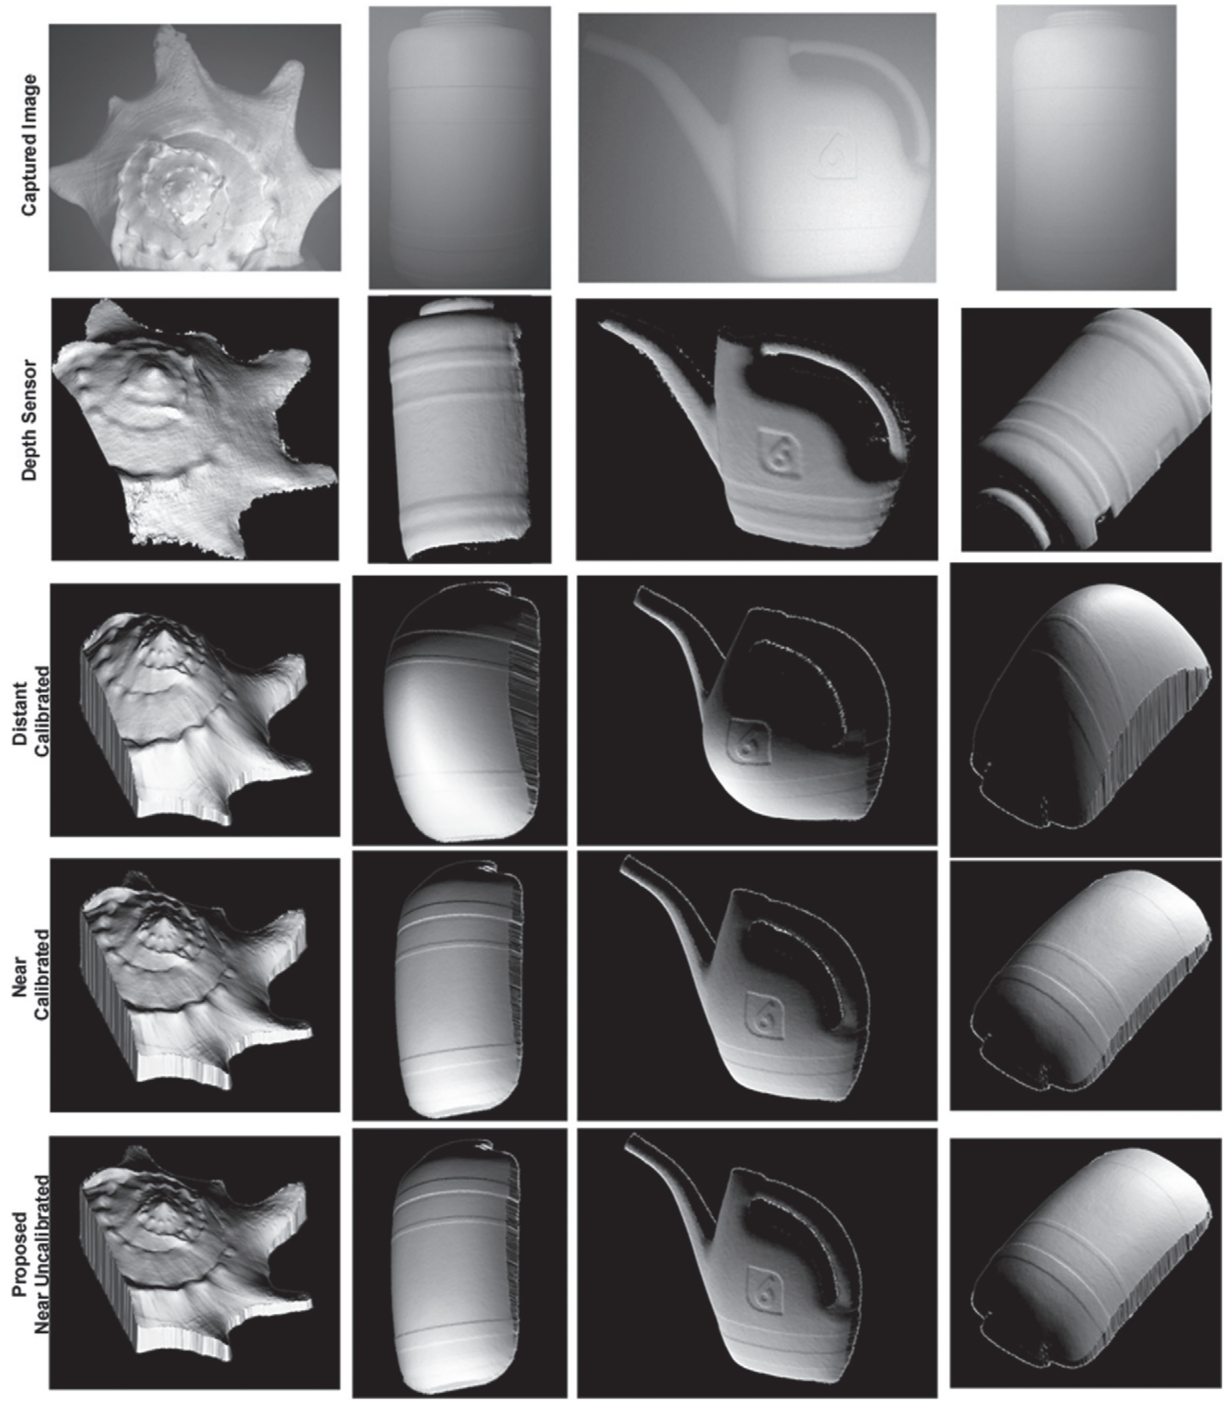

Fig. 17. Reconstruction results for different objects and scattering levels.

Acknowledgments

This work was supported by the contract 270180 of the European Communities FP7 (NOPTILUS). We are very grateful to Underwater Systems and Technologies Laboratory (LSTS) and OceanScan in Porto for the access to the ROV and the help with the port-water experiments.

Table 1
Estimated Z_c (m).

	Low scattering	High scattering ml
Calibrated	0.7	0.7
Can	0.744	0.664
Container	0.749	0.722
Shell	0.672	0.799

Table 2
Estimated c (m^{-1}).

	Low scattering	High scattering
Calibrated	1.371	1.944
Can	1.273	2.357
Container	1.753	2.49
Shell	1.417	1.68

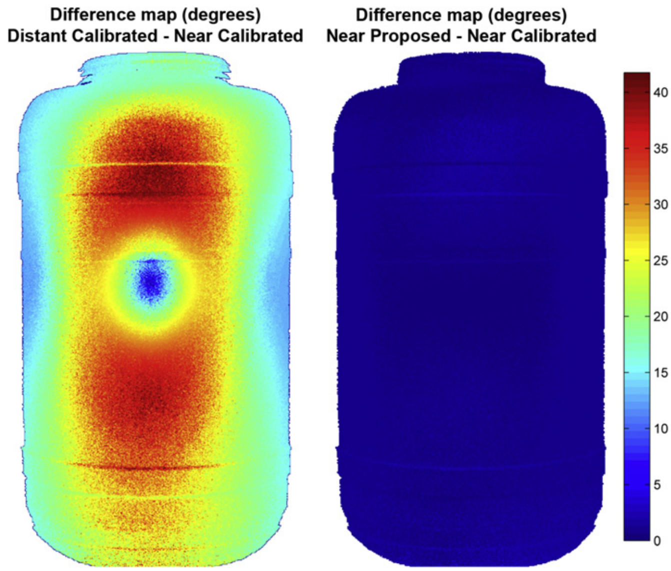

Fig. 18. The difference in normal map degrees between near-calibrated and distant-calibrated (left), and near-calibrated and our proposed uncalibrated algorithm (right).

References

[1] G. Foresti, Visual inspection of sea bottom structures by an autonomous underwater vehicle, IEEE Trans. Syst. Man Cybern. B Cybern. 31 (5) (2001) 691–705.
[2] T. Treibitz, B.P. Neal, D.I. Kline, O. Beijbom, P.L. Roberts, B.G. Mitchell, D. Kriegman, Wide Field-of-View Daytime Fluorescence Imaging of Coral Reefs, IEEE Oceans, 2013.
[3] 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, OCEANS-Bergen, 2013 MTS/IEEE, IEEE. 2013, pp. 1–8.
[4] R. Garcia, N. Gracias, Detection of interest points in turbid underwater images, OCEANS, 2011 IEEE-Spain, IEEE. 2011, pp. 1–9.
[5] A. Sarafraz, S. Negahdaripour, Y. Schechner, Performance assessment in solving the correspondence problem in underwater stereo imagery, OCEANS 2010, Sept. 2010. pp. 1–7.
[6] T. Papadhimitri, P. Favaro, Uncalibrated Near-Light Photometric Stereo, Proc. British Machine Vision Conference, 2014.
[7] J. Ahmad, J. Sun, L. Smith, M. Smith, Improving photometric stereo through per-pixel light vector calculation, Proc. British Machine Vision Conference, 2013.
[8] P. Favaro, T. Papadhimitri, A closed-form solution to uncalibrated photometric stereo via diffuse maxima, Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE. 2012, pp. 821–828.
[9] J.S. Jaffe, Computer modeling and the design of optimal underwater imaging systems, IEEE J. Ocean. Eng. 15 (2) (1990) 101–111.
[10] B.L. McGlamery, A computer model for underwater camera systems, SPIE Ocean Opt. 208 (1979) 221–231.
[11] T. Treibitz, Y. Schechner, Active polarization descattering, IEEE Trans. Pattern Anal. Mach. Intell. 31 (3) (2009) 385–399.
[12] S. Narasimhan, S. Nayar, B. Sun, S. Koppal, Structured light in scattering media, Computer Vision, 2005. 2005. Tenth IEEE International Conference on, vol. 1, 2005. pp. 420–427. Vol. 1.

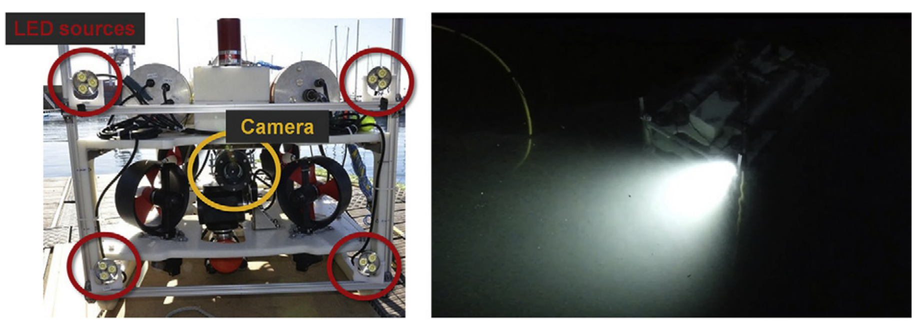

Fig. 19. The Photometric Stereo setup that was installed on a remotely operated vehicle. Four sources were synchronized with a camera so that successive frames were captured having only one light on at a time. The right image shows the system operating in murky water.

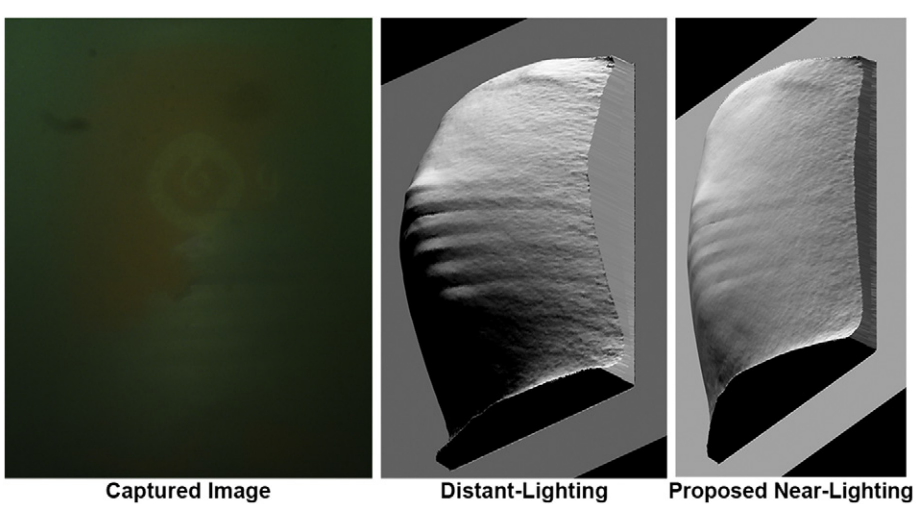

Fig. 20. Left: one of the four captured images using Photometric Stereo in very murky conditions in the port water of Porto, Portugal. Middle-Right: the reconstructed and re-lit surface using distant-lighting and our near-lighting algorithm.

[13] S. Zhang, S. Negahdaripour, 3-D shape recovery of planar and curved surfaces from shading cues in underwater images, IEEE J. Ocean. Eng. 27 (1) (2002) 100–116.
[14] C. Tsiotsios, M.E. Angelopoulou, T.-K. Kim, A.J. Davison, Backscatter Compensated Photometric Stereo with 3 Sources, Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE. 2014, pp. 2259–2266.
[15] Z. Wu, M. Jian, Y. Sun, Augmenting Photometric stereo with coaxial illumination under water, Int. J. Inf. Comput. Sci. 10 (18) (2013) 5981–5988.
[16] P. Tan, S.P. Mallick, L. Quan, D. Kriegman, T. Zickler, Isotropy, reciprocity and the generalized bas-relief ambiguity, Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE. 2007, pp. 1–8.
[17] D. Nehab, S. Rusinkiewicz, J. Davis, R. Ramamoorthi, Efficiently combining positions and normals for precise 3D geometry, ACM Trans.Graph. (TOG) 24 (3) (2005) 536–543.
[18] S. Negahdaripour, H. Zhang, X. Han, Investigation of photometric stereo method for 3-d shape recovery from underwater imagery, OCEANS'02 MTS/IEEE, vol. 2, 2002. pp. 1010–1017.
[19] N. Kolagani, J.S. Fox, D.R. Blidberg, Photometric stereo using point light sources, Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on, IEEE. 1992, pp. 1759–1764.
[20] J. Oakley, B. Satherley, Improving image quality in poor visibility conditions using a physical model for contrast degradation, IEEE Trans. Image Process. 7 (2) (Feb. 1998) 167–179.
[21] K. Tan, J.P. Oakley, Physics-based approach to color image enhancement in poor visibility conditions, J. Opt. Soc. Am. A 18 (10) (Oct. 2001) 2460–2467.
[22] T. Pouli, D.W. Cunningham, E. Reinhard, Image statistics and their applications in computer graphics, Eurographics, State of the Art Report, vol. 72, 2010.
[23] J.T. Barron, J. Malik, Shape, illumination, and reflectance from shading, IEEE Trans. Pattern Anal. Mach. Intell. (2015)
[24] Y. Iwahori, H. Sugie, N. Ishii, Reconstructing shape from shading images under point light source illumination, Pattern Recognition, 1990. Proceedings., 10th International Conference on, vol. 1, IEEE. 1990, pp. 83–87.
[25] R.H. Byrd, M.E. Hribar, J. Nocedal, An interior point algorithm for large-scale nonlinear programming, SIAM J. Optim. 9 (4) (1999) 877–900.
[26] J.S. Jaffe, Multi autonomous underwater vehicle optical imaging for extended performance, OCEANS 2007-Europe, IEEE. 2007, pp. 1–4.
[27] T. Treibitz, Y. Schechner, Turbid scene enhancement using multi-directional illumination fusion, IEEE Trans. Image Process. 21 (11) (2012) 4662–4667.
[28] R.A. Newcombe, A.J. Davison, S. Izadi, P. Kohli, O. Hilliges, J. Shotton, D. Molyneaux, S. Hodges, D. Kim, A. Fitzgibbon, KinectFusion: Real-time dense surface mapping and tracking, Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on, IEEE. 2011, pp. 127–136.
[29] C. Tsiotsios, M.E. Angelopoulou, A.J. Davison, T.-K. Kim, Effective Backscatter Approximation for Photometry in Murky Water, 2016. arXiv preprint arXiv:1604.08789
[30] C. Tsiotsios, T. Kim, A. Davison, S. Narasimhan, Model effectiveness prediction and system adaptation for photometric stereo in murky water, Comput. Vis. Image Underst. (2016)
[31] Z. Murez, T. Treibitz, R. Ramamoorthi, D. Kriegman, Photometric Stereo in a Scattering Medium, 2015.
[32] M. Gupta, S. Narasimhan, Y. Schechner, On controlling light transport in poor visibility environments, Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE. 2008, pp. 1–8.
[33] Y.Y. Schechner, N. Karpel, Recovery of underwater visibility and structure by polarization analysis, IEEE J. Ocean. Eng. 30 (3) (July 2005) 570–587.
[34] Y.Y. Schechner, N. Karpel, Clear underwater vision, Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on, vol. 1, IEEE. 2004, pp. I–536.
[35] University of Manchester. School of Electrical and Electronic Engineering, H. Mortazavi, J. Oakley, Mitigation of Contrast Loss in Underwater Images, University of Manchester. 2010, http://books.google.co.uk/books?id=L_tsMwEACAAJ.
[36] T.W. Cronin, N. Shashar, R.L. Caldwell, J. Marshall, A.G. Cheroske, T.-H. Chiou, Polarization vision and its role in biological signaling, Integr. Comp. Biol. 43 (4) (2003) 549–558.
[37] S. Barsky, M. Petrou, The 4-source photometric stereo technique for three--dimensional surfaces in the presence of highlights and shadows, IEEE Trans. Pattern Anal. Mach. Intell. 25 (10) (2003) 1239–1252.
[38] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, Y. Ma, Robust photometric stereo via low-rank matrix completion and recovery, Computer Vision-ACCV 2010, Springer. 2011, pp. 703–717.
[39] S.K. Nayar, X.-S. Fang, T. Boult, Removal of specularities using color and polarization, Computer Vision and Pattern Recognition, 1993. Proceedings CVPR'93., 1993 IEEE Computer Society Conference on, IEEE. 1993, pp. 583–590.