

Underwater Single Image Color Restoration Using Haze-Lines and a New Quantitative Dataset

Dana Berman, Deborah Levy, Shai Avidan, and Tali Treibitz

Abstract

Underwater images suffer from color distortion and low contrast, because light is attenuated while it propagates through water. Attenuation under water varies with wavelength, unlike terrestrial images where attenuation is assumed to be spectrally uniform. The attenuation depends both on the water body and the 3D structure of the scene, making color restoration difficult. Unlike existing single underwater image enhancement techniques, our method takes into account multiple spectral profiles of different water types. By estimating just two additional global parameters: the attenuation ratios of the blue-red and blue-green color channels, the problem is reduced to single image dehazing, where all color channels have the same attenuation coefficients. Since the water type is unknown, we evaluate different parameters out of an existing library of water types. Each type leads to a different restored image and the best result is automatically chosen based on color distribution. We also contribute a dataset of 57 images taken in different locations. To obtain ground truth, we placed multiple color charts in the scenes and calculated its 3D structure using stereo imaging. This dataset enables a rigorous quantitative evaluation of restoration algorithms on natural images for the first time.

Index Terms



-
- *D. Berman and S. Avidan are with the School of of Electrical Engineering, Tel Aviv University, Israel.
E-mail: ttreibitz@univ.haifa.ac.il*
 - *D. Levy and T. Treibitz are with the Hatter Department of Marine Technologies, Charney School of Marine Sciences, University of Haifa, Israel.*

Underwater Single Image Color Restoration Using Haze-Lines and a New Quantitative Dataset

1 INTRODUCTION

UNDERWATER images often lack contrast and depict inaccurate colors due to the scattering and absorption of light as it propagates through the water. Yet color and contrast are extremely important for visual surveys in the ocean. For example, enhanced images can improve automatic segmentation, increase the accuracy of feature matching between images taken from multiple viewpoints, and aid in navigation.

The attenuation of light depends both on the light's wavelength and the distance it travels [1]. The wavelength-dependent attenuation causes color distortions that increase with an object's distance. In addition, the scattering induces a distance-dependent additive component on the scene which reduces contrast. These phenomena cannot be globally corrected since the color degradation depends on the distance of the object from the camera. Moreover, the attenuation parameters are affected by seasonal, geographic, and climate variations. These variations were categorized into different optical water types by Jerlov [2].

Unfortunately, existing single underwater image enhancement techniques under-perform as they do not take into account the diverse spectral properties of water. In addition, their evaluation is generally based on a handful of images, and is mostly qualitative. A few methods have been evaluated using no-reference image quality metrics, which operate only on the luminance and cannot measure the color correction.

In this paper we suggest a method to recover the distance maps and object colors in scenes photographed under water and under ambient illumination, using just a single image as input. Our recovery takes into account the different optical water types and is based on a more comprehensive physical image formation model than the one previously used. Restoration from a single image is desirable because water properties temporally change, sometimes within minutes. In addition, no additional equipment, such as a tripod or filters, is required.

A variety of methods have been developed for the closely related single image dehazing problem, in which images are degraded by weather conditions such as haze or fog. Under the assumption of wavelength-independent attenuation, single image dehazing is an ill-posed problem with three measurements per pixel (the R,G,B values of the input image) and four unknowns (the R,G,B values of the object and its distance from the camera). The transmission is the fraction of the scene's radiance that reaches the camera, and is related to the distance via the attenuation coefficient.

Under water, where the assumption of wavelength-independent attenuation does not hold, there are theoretically three unknown transmission values per pixel, one per channel, yielding six unknowns with only three mea-

surements. However, the color-dependent transmission is related to the distance via the attenuation coefficients. Based on this relation we reduce the problem to estimation of four unknowns per pixel as before, with two new *global* parameters, the ratios between the attenuation coefficients of the color channels. We estimate these parameters using an existing library of water types, and based on the color distribution of the image after correction. We utilize the fact that using a wrong water type leads to distorted colors. Our results demonstrate a successful single image restoration of underwater scenes using a comprehensive physical image formation model. Thus, we are able to recover more complex 3D scenes than previous methods and, in addition, estimate the optical water type. Fig. 1 depicts the proposed method.

Since no public datasets with ground truth data for underwater images are available, we took upon ourselves to acquire an *in situ* dataset with two channels of ground-truth data. We acquired linear images with a stereo rig of high-end SLR cameras and in each scene placed several identical color card in different distances from the camera. The color cards are used to evaluate scene recovery in *multiple distances*. The stereo images are used to recover the true distance from the camera to evaluate transmission estimation. This dataset currently contains 57 scenes is available publicly. We conducted a thorough quantitative analysis, comparing the results of multiple algorithms to the ground truth data. Our algorithm is competitive with other state-of-the-art methods.

2 RELATED WORK

The participating medium causes color shifts and reduces contrast in images. The degradation depends both on the water properties and the 3D structure, which can be estimated by analyzing *multiple* images that satisfy certain conditions. E.g., Schechner and Karpel [4], [5] take two images with orthogonal polarizer angles and utilize the partial polarization of light to restore the visibility. Multiple images of the same object taken from different *known* viewpoints [6], [7], [8] are used to estimate attenuation coefficients and recover the scene. These methods have limited applicability since they impose constraints on the imaging conditions.

Scattering media limits visibility in terrestrial images as well, when imaging in fog, haze, or turbulence [9]. *Single* image dehazing methods, e.g. [10], [11], [12], take as input a single hazy image and estimate the unknown distance map and the clear image of the scene. Different priors are used in order to solve the ill-posed problem, under the assumption of color-independent transmission, which is reasonable for terrestrial images yet is violated under water. The dark channel prior (DCP) [11] assumes that within small image

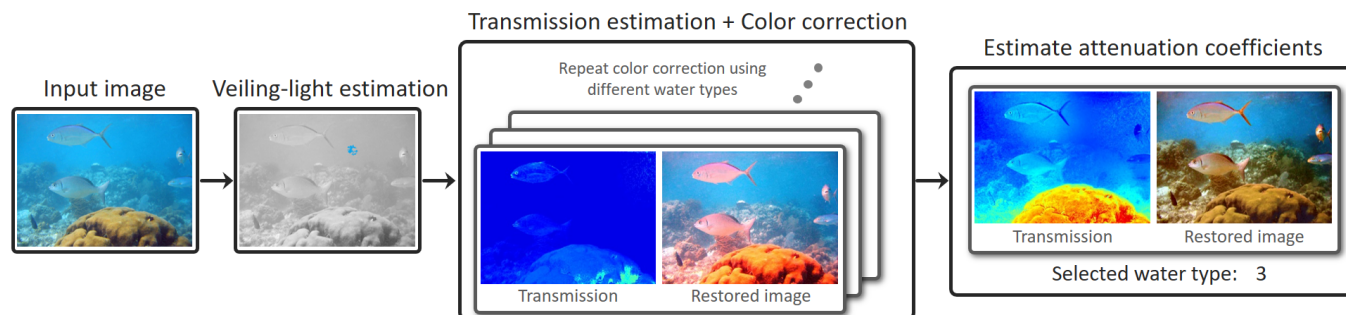


Fig. 1. The proposed color restoration and transmission estimation method: First, the veiling-light (the ambient light that is scattered into the line of sight) is estimated. The pixels whose color is the veiling light are shown in color. Then, the transmission estimation and color restoration are repeated for multiple water types that have different optical characteristics, as described in Section 3.2. Finally, the best result is selected automatically based on the gray-world assumption. Photo from [3].

patches, at least one pixel has a low value in some color channel, and uses the minimal value to estimate the distance. The haze-lines prior [10] is based on the observation that the number of colors in natural images is small, and that similar colors appear all over the image plane.

Some variations of DCP were proposed for the underwater domain [3], [13], [14], [15], [16]. Perez *et al.* [17] recently performed an evaluation of DCP variations for underwater autonomous vehicles. Carlevaris-Bianco *et al.* [3] assume a color-independent transmission, and propose a variant of DCP based on the difference between the maximum of the red channel and the maximum of the blue and green channels in each patch. They claim this value is inversely related to transmission, since red is attenuated at a higher rate than blue or green. Drews *et al.* [14] apply DCP to the blue and green channels only, since the red channel is often attenuated rapidly and cannot be used to estimate the transmission. They get improved transmission maps, but still assume a uniform transmission across channels for recovery. Chiang and Chen [13] recover the transmission using standard DCP. They assume the recovered transmission is the transmission of the red channel, as under water this channel has the lowest transmission. They use fixed attenuation coefficients measured for open ocean waters to recover the image based on the estimated red transmission. Lu *et al.* [16] estimate the transmission using the darker of the blue and red channels. They use the same fixed water attenuation as in [13] to recover the scene. Galdran *et al.* [15] suggest the Red-Saturation prior, incorporating information from the inverse of the red channel and the saturation in order to estimate the transmission. They use a spectrally uniform transmission since the water type is not easy to determine. Instead, they add an additive veiling-light term to the color restoration.

Emberton *et al.* [18] detect and segment regions of pure veiling-light and deliberately set their transmission value to be high, while using the estimation of Drews *et al.* [14] for the rest of the scene. This is done to avoid enhancing artifacts in the veiling-light regions. To handle the spectral dependency of the attenuation, they classify the water type to blue, turquoise, or green-dominated, and apply white-balance to the image before estimating the spectrally-uniform transmission. Since this step is global, it cannot compensate for the distance-dependent attenuation. In addition, the classification to three different water types is too coarse compared to common optical classification schemes.

Despite the abundance of DCP based methods, the underlying assumption does not hold in many underwater scenarios: bright foreground sand has high values in all color channels and might be mistaken to have a low transmission despite being close to the camera. Moreover, the background water has a dominant color (hence at least one color channel is low), and many of the mentioned methods inaccurately estimate the transmission there to be high.

Peng *et al.* [19] estimate the scene distance via image blurriness, which grows with the distance due to scattering. They disregard the spectral dependence of the transmission. Peng and Cosman [20] combine the blurriness prior with [3] and assume open ocean waters. While this prior is physically valid, it has limited efficiency in textureless areas.

The above mentioned restoration methods aim for a physics-based recovery of a scene's colors, while estimating its 3D structure in the process. Other methods aim for a visually pleasing result, e.g. [21], [22], [23], [24], [25], [26], but have not shown color consistency that is required for scientific measurements.

Convolutional networks have been recently used for dehazing [27], [28], and also for underwater restoration [29]. However, their training relies on purely synthetic data and thus highly depends on the quality of the simulation models, which wrongly neglect the spectral dependency of the attenuation, leading to poor results. We believe that the severe lack of data and the challenge in conducting *in situ* experiments are withholding progress in this direction. Two recent attempts to cope with this problem include synthesizing data and acquiring images in test tanks. Blasinski *et al.* [30] released tools to realistically simulate the appearance of underwater scenes in a variety of water conditions. Unfortunately, the simulation accuracy decreases with increasing depth and chlorophyll concentration. To the best of our knowledge, it hasn't been utilized for learning-based underwater image enhancement methods yet. Duarte *et al.* [31] published a dataset for evaluating underwater image restoration methods. However, these images were acquired in a test tank using milk, and do not represent real-world conditions. In addition to the theoretical contributions, a major goal of this research is to create and publish an extensive dataset of underwater images with ground-truth of the 3D scene structure.

Recently, it was shown that the commonly used underwater image formation model [5] is inaccurate and introduces reconstruction errors [32]. A revised model was




Fig. 2. Left: Approximate attenuation coefficients (β) of Jerlov water types. Data taken from [1], based on measurements in [39]. Solid lines mark open ocean water types while dashed lines mark coastal water types. Middle: β ratios of R, G, B channels for each water type, based on the wavelengths: 475nm, 525nm, and 600nm, respectively. Right: simulation of the appearance of a perfect white surface viewed at depth of 1-20m in different water types (reproduced from [40]).

suggested, with a larger number of parameters. A method using this model [33] has shown superior results, however, it requires knowledge of the 3D scene structure. So far, no single image method that uses the revised model has been presented. Therefore, here we use the commonly used image formation model and show that even with it results can still be improved.

The wavelength dependent absorption properties of water are used in [34] to recover depth. However this method requires an active illumination system and is valid only for short distances, since near-infrared light is rapidly attenuated in water. Known geometry is used in [35] to plan the imaging viewpoints under water that result in high contrast.

Morimoto et al. [36] estimate the optical properties of layered surfaces based on non-linear curves in RGB space, which is a similar physical model. Nonetheless, user scribbles are required in order to detect the curves.

3 BACKGROUND

3.1 Image Formation Model

We follow the model developed in [5]. In each color channel $c \in \{R, G, B\}$, the image intensity at each pixel is composed of two components, attenuated signal and veiling-light:

$$I_c(\mathbf{x}) = t_c(\mathbf{x})J_c(\mathbf{x}) + (1 - t_c(\mathbf{x})) \cdot A_c, \quad (1)$$

where bold denotes vectors, \mathbf{x} is the pixel coordinate, I_c is the acquired image value in color channel c , t_c is the transmission of that color channel, and J_c is the object radiance that we wish to restore. The global veiling-light component A_c is the scene value in areas with no objects ($t_c = 0, \forall c \in \{R, G, B\}$). Eq. (1) applies to linear captured data, prior to in-camera processing such as color-space conversion, tone mapping, and compression. Therefore, I refers to the image obtained from the raw file after minimal processing such as demosaicing and black current subtraction [37], [38].

The transmission depends on object's distance $z(\mathbf{x})$ and the water attenuation coefficient for each channel β_c :

$$t_c(\mathbf{x}) = \exp(-\beta_c z(\mathbf{x})) . \quad (2)$$

In the ocean, the attenuation of red colors can be an order of magnitude larger than the attenuation of blue and green [1]. Therefore, as opposed to the common assumption in single image dehazing, the transmission $\mathbf{t} = (t_R, t_G, t_B)$ is wavelength-dependent.

3.2 Water Attenuation

The attenuation of light under water is not constant and varies with geographical, seasonal, and climate related events. In clear open waters, visible light is absorbed at the longest wavelengths first, appearing deep-blue to the eye. In near-shore waters, sea water contains more suspended particles than the central ocean waters, which scatter light and make coastal waters appear less clear than open waters. In addition, the absorption of the shortest wavelengths is stronger, thus the green wavelength reaches deeper than other wavelengths.

Jerlov [2] developed a frequently used classification scheme for oceanic waters, based on water clarity. The Jerlov water types are I, IA, IB, II and III for open ocean waters, and 1 through 9 for coastal waters. Type I is the clearest and type III is the most turbid open ocean water. Likewise, for coastal waters, type 1 is clearest and type 9 is most turbid. Fig. 2(Left) depicts the attenuation coefficients' dependence on wavelength, while Fig. 2(Right) shows an RGB simulation of the appearance of a perfect white surface viewed at different depths in different water types. The common notion that red colors are attenuated faster than blue/green only holds for oceanic water types.

When capturing an image using a commercial camera, three color channels R, G, B are obtained. Thus, we are interested in three attenuation coefficients: $(\beta_R, \beta_G, \beta_B)$ in order to correct the image. We use the Jerlov water types to constrain the space of attenuation coefficients in the RGB domain. We show in Sec. 4.2 that the three attenuation coefficient themselves are not required for transmission estimation, but rather their ratios (two variables). Fig. 2(Middle) shows the ratios of the attenuation coefficients: β_B/β_R vs. β_B/β_G of Jerlov water types for wavelengths of peak camera sensitivity according to [41] (475nm, 525nm, and 600nm for B, G, R , respectively). Essentially, we approximate the cameras' spectral response as a Dirac delta function, similarly to [20], [36]. This approximation is camera-agnostic, which is advantageous since the true response functions are rarely published.

4 COLOR RESTORATION

We show in the following that if we were given the two *global* attenuation ratios, then the problem can be reduced to single image dehazing (i.e., single attenuation coefficient across all channels) for which good algorithms exist. The

question now becomes how to estimate these two parameters? Our solution is to evaluate every possible water type, as defined by Jerlov, and pick the best one. Each water type defines a known, and fixed, pair of attenuation ratios that we can use. Once we have evaluated all possible water types (10 in total) we use a standard natural image statistics prior to pick the best result. In particular, we found the Gray-World assumption prior to work well for our needs. The method is illustrated in Fig. 1 and summarized in Alg. 1.

4.1 Veiling-Light Estimation

First we describe how we estimate the veiling-light, which is required for de-scattering. We assume an area without objects is visible in the image, in which the pixels' color is determined by the veiling-light alone. Such an area should be smooth and not have texture, which is an important feature for veiling light estimation, e.g. [18], [20]. This assumption often holds when the line of sight is horizontal. It does not hold when photographing a reef wall up close, or when the camera is pointed downwards. However, in these cases, the distance of objects from the camera usually varies less than in horizontal scenes, and a simple contrast stretch is likely to be sufficient.

In order to detect the pixels that belong to the veiling-light, we apply a local linear contrast stretch (CLAHE) and then generate an edge map of the scene using the Structured Edge Detection Toolbox [42] with pre-trained model and default parameters. We limit this step to the top 20% of the image. We then threshold the edge map and look for the largest connected component. The pixels belonging to the largest component are classified as veiling-light pixels ($\mathbf{x} \in VL$). The veiling-light \mathbf{A} is the average color of those pixels. This is demonstrated on the veiling-light estimation step of Fig. 1, where only the pixels $\mathbf{x} \in VL$ are shown in color.

4.2 Transmission Estimation

Combining and rearranging Eqs. (1,2) yields for the blue channel:

$$A_B - I_B = e^{-\beta_B Z} \cdot (A_B - J_B) \quad (3)$$

and similarly for the red channel:

$$A_R - I_R = e^{-\beta_R Z} \cdot (A_R - J_R) \quad (4)$$

Raising Eq. (4) to the power of $\frac{\beta_B}{\beta_R}$ yields:

$$(A_R - I_R)^{\frac{\beta_B}{\beta_R}} = e^{-\beta_R Z \cdot \frac{\beta_B}{\beta_R}} (A_R - J_R)^{\frac{\beta_B}{\beta_R}} = t_B (A_R - J_R)^{\frac{\beta_B}{\beta_R}} \quad (5)$$

Denote the ratios between the attenuation coefficients:

$$\beta_{BR} = \beta_B / \beta_R, \quad \beta_{BG} = \beta_B / \beta_G \quad (6)$$

Then, in this medium-compensated space we achieve a form similar to Eq. (1), with one unknown transmission per-pixel, common to all color channels:

$$\begin{bmatrix} (I_R(\mathbf{x}) - A_R)^{\beta_{BR}} \\ (I_G(\mathbf{x}) - A_G)^{\beta_{BG}} \\ (I_B(\mathbf{x}) - A_B) \end{bmatrix} = t_B(\mathbf{x}) \begin{bmatrix} (J_R(\mathbf{x}) - A_R)^{\beta_{BR}} \\ (J_G(\mathbf{x}) - A_G)^{\beta_{BG}} \\ (J_B(\mathbf{x}) - A_B) \end{bmatrix} \quad (7)$$

This form is similar to the Haze-Lines [10] formulation. Therefore, we similarly cluster the pixels to Haze-Lines and

obtain an initial estimation of the transmission of the blue channel \tilde{t}_B . Since the value $(I_c - A_c)$ might be negative, we avoid numerical issues when raising to the power β_{Bc} by raising the absolute value and keeping the sign.

In [10] it was assumed that there is a haze-free pixel in each Haze-Line. However, the attenuation coefficients measures by Jerlov indicate that even scene points that are located at a distance of only one meter from the camera have a t_B of about 0.9, depending on water type. Thus, we multiply the initial transmission estimation by 0.9 (similarly to [19]).

A bound on the transmission arises from the fact that $J_c \geq 0, \forall c \in \{R, G, B\}$. We substitute this bound in Eq. (1) and obtain a lower bound t^{LB} on the transmission of the blue channel, t_B , taking into account the different attenuation coefficients of the different color channels:

$$t^{LB} := \max \left\{ 1 - \frac{I_B}{A_B}, \left(1 - \frac{I_G}{A_G} \right)^{\beta_{BG}}, \left(1 - \frac{I_R}{A_R} \right)^{\beta_{BR}} \right\} \quad (8)$$

We detect pixels \mathbf{x} with no scene objects based on their Mahalanobis distance from the distribution of the veiling-light pixels: $D_M(\mathbf{I}(\mathbf{x}))$ (from here on, when referring to the Mahalanobis distance, it is with respect to the distribution of intensities of veiling-light pixels). We set the transmission of such pixels to be the lower bound t^{LB} calculated in Eq. (8). However, a binary classification of the pixels to $VL, \bar{V}\bar{L}$ often results in abrupt discontinuities in the transmission map, which are not necessarily distance discontinuities. Therefore, we use a soft-matting and calculate the transmission as follows:

$$t_B(\mathbf{x}) = \begin{cases} t^{LB}(\mathbf{x}) & D_M(\mathbf{I}(\mathbf{x})) \leq \overline{D_M} + \sigma_M \\ \tilde{t}_B(\mathbf{x}) & D_M(\mathbf{I}(\mathbf{x})) \geq D_M^{max} + \sigma_M \\ \alpha(\mathbf{x}) \cdot t^{LB}(\mathbf{x}) + (1 - \alpha(\mathbf{x})) \cdot \tilde{t}_B(\mathbf{x}) & \text{otherwise} \end{cases} \quad (9)$$

where $\overline{D_M} = \frac{1}{|VL|} \cdot \sum_{\mathbf{x} \in VL} D_M(\mathbf{I}(\mathbf{x}))$ is the average Mahalanobis distance of the veiling-light pixels, $D_M^{max} = \max_{\mathbf{x} \in VL} \{D_M(\mathbf{I}(\mathbf{x}))\}$ is their maximal Mahalanobis distance and σ_M is the standard deviation. $\alpha(\mathbf{x})$ is the matting coefficient for pixels that cannot be classified to object/ water with high probability: $\alpha(\mathbf{x}) = \frac{D_M(\mathbf{I}(\mathbf{x})) - \overline{D_M} - \sigma_M}{D_M^{max} - \overline{D_M} - \sigma_M}$, yielding a relatively steep transition between VL and $\bar{V}\bar{L}$.

Finally, we regularize the transmission using Guided Image Filter [43], with a contrast-enhanced input image as guidance.

4.3 Scene Recovery

Once t_B is estimated, we can compensate for the color attenuation using the following:

$$J_c = A_c + \frac{I_c - A_c}{e^{-\beta_c Z}} = A_c + \frac{I_c - A_c}{t_B^{\beta_c / \beta_B}}, \quad \forall c \in \{R, G, B\} \quad (10)$$

Eq. (10) compensates for the intensity changes that happen in the path between the object and the camera. In addition, the ambient illumination is attenuated by the water column from the surface to the imaging depth, resulting in a colored global illumination. We are interested in restoring the colors as if they were viewed under white light, without a color cast. Since this effect is global in the scene, we correct it by performing white balance on the result. This global

Algorithm 1: Underwater image restoration**Input:** $I(\mathbf{x})$ **Output:** $\hat{J}(\mathbf{x}), \hat{t}(\mathbf{x})$

- 1: Detect veiling light pixels using structured edges (Sec. 4.1) and calculate A
- 2: **for** each (β_{BR}, β_{BG}) values of water types (Fig. 2 middle) **do**
- 3: **for** each $c \in \{R, G, B\}$ **do**
- 4: $\tilde{I}_c(\mathbf{x}) = \text{sign}(I_c(\mathbf{x}) - A_c) \cdot \text{abs}(I_c(\mathbf{x}) - A_c)^{\beta_{Bc}}$
 (β_{Bc} is defined in Eq. 6)
- 5: **end for**
- 6: Cluster pixels to 500 Haze-Lines as in [10] and estimate an initial transmission \tilde{t}_B
- 7: Apply soft matting to \tilde{t}_B with lower bound (Eq. 9)
- 8: Regularization using guided image filter, with a contrast-enhanced input as guidance
- 9: Calculate the restored image using Eq. 10
- 10: Perform a global WB on the restored image
- 11: **end for**
- 12: Return the image that best adheres to the Gray-World assumption on pixels $\mathbf{x} \notin VL$

operator works well only because the distance-dependent attenuation has already been compensated for.

Finally, since Eq. (1) applies to the linear captured data, we convert the linear image to sRGB using a standard image processing pipeline, including color-space conversion and gamma curve for tone mapping [38].

4.4 Estimation of Attenuation Coefficient Ratios

Using the wrong coefficients (β_{BR}, β_{BG}) leads to restorations with skewed colors and wrong transmission maps. We use this insight to determine the most appropriate water type. We perform the restoration multiple times using different ratios of attenuation coefficients, corresponding to different Jerlov water types, and choose the best result automatically based on the Gray-World assumption. The attenuation coefficients' ratios are shown in Fig. 2(middle).

According to the Gray-World assumption [44], the average reflectance of surfaces in the world is achromatic. It has been used in the past for estimating attenuation coefficients underwater using known distances [6]. However, a significant portion of images taken under water often contain water without any objects. The Gray-World assumption obviously does not hold there. Therefore, we apply the Gray-World assumption only at image regions that contain objects, i.e., those that were not identified as veiling-light pixels. Thus, among all results using different water types, we choose the image with the smallest difference between the average values of the red, green and blue channels.

We considered several other measures (e.g. maximal contrast [12]), but found that the simple Gray-World assumption gave the best results and therefore we focus on this measure.

4.5 Applicable Distance Range

This method, similarly to all physics-based methods, does not aim to recover areas with zero signal. In most underwater environments the red light is attenuated first and therefore is the limiting factor. The range for which the




Fig. 3. **Quantitative evaluations in the literature.** (a) Diver holding a board with six colored patches, from left to right: before diving, diving at a depth of 5m, and diving at a depth of 15m. (b-c) Water tank experiments: due to the tank's size (90cm \times 45cm \times 45cm), the objects are 30 cm deep and the distance between the objects and the camera is approximately 60cm. Deep sea soil is added to the seawater to increase scattering, at a concentration of 20mg/liter (b) and 100mg/liter (c). (d) Example of an underwater image taken in a swimming pool [23], where similar photos were taken with different cameras to evaluate white balancing techniques.

red signal is above zero depends on the water properties and the camera's noise properties and settings. This topic is thoroughly discussed in [45]. In general it is shown there the signal exists up to ~ 3 attenuation lengths (where attenuation length is defined as βz). This distance (sum of horizontal and vertical) ranges between 18m to 5m in the Jerlov water types I – 9.

5 DATABASE

5.1 Experimental Set-up and Method

We wish to quantitatively evaluate the proposed method and compare it to other methods. Since it is impossible to obtain medium-free *in situ* images, comparisons in the existing literature are either based on photos of a color card taken from a short range, e.g. [13], [16], [23] (shown in Fig. 3), or based on non-reference image quality metrics [46], [47]. Photos of color cards, even if they are taken at different depths, e.g. [13], are usually taken from a short range and do not contain objects at different distances from the camera as is often the case in natural scenes. Nevertheless, restoring colors of objects that are at a single distance from the camera is a much easier problem than the challenging color reconstruction required for a complex 3D scene. Similarly, images taken in water tanks [16] or swimming pools [23] depict objects close to the camera, and exhibit different absorption and scattering properties than natural bodies of water. Alternatively, several no-reference image quality metrics have been developed for comparing natural underwater scenes. UCIQE [47] metric is designed to quantify the nonuniform color cast and low-contrast that characterize underwater images, while UIQM [46] addresses three underwater image quality criterions: colorfulness, sharpness and contrast. These metrics are somewhat heuristic and have limited applicability, as shown in [48].

Therefore, we decided to gather *in situ* data at different seasons, depths, and water types (both in tropical waters and in murkier coastal waters). All of the scenes were illuminated by natural light only. We made sure all scenes




Fig. 4. **An example scene from the dataset.** [Left] The image from the left camera. The sizes of two identical color chart were measured based on the 3D reconstruction, and are marked on the enlarged areas on the left. The true size of the chart is: $12\text{cm} \times 18.5\text{cm}$. At a distance of 6m from the camera, the reconstruction accuracy is about $\sim 1.5\text{cm}$, which is quite accurate considering the precision of existing datasets [49]. [Center] The distances of the objects from the left camera. [Right] The same scene with the color cards masked.



Fig. 5. **Creating the dataset.** [Top left] The stereo rig used for acquisition. [Bottom left] The color charts and calibration checkerboard. [Right] Dive sites' location.

have a complex 3D structure with objects at the visibility limit.

Our database contains 57 stereo pairs from four different sites in Israel, two in the Red Sea (representing tropical water) and two in the Mediterranean Sea (temperate water). In the Red Sea the sites were a coral reef ('Katzaa') which is 10-15 meters deep (15 pairs) and a shipwreck ('Satil'), 20-30 meters deep (8 pairs). In the Mediterranean Sea both sites were rocky reef environments, separated by 30km, Nachsholim at 3-6 meters depth (13 pairs), and Mikhmoret at 10-12 meters depth (21 pairs). The dive sites' geographical location is displayed in Fig. 5[Right]. For brevity in the following we denote them Red 1, Red 2, Med 1, Med 2, respectively.

To enable ground truth validation the quantitative evaluation is two-fold, and relies both on color and distances. First, we placed *several* identical waterproof color charts (from DGK Color Tools) at different distances from the camera (circled in purple in Fig. 5right), as opposed to previous methods that only show one card at a fixed distance. This enables us to make sure that the color restoration is consistent across different distances from the camera. Since in some scenes it was difficult to place the cards, a subset of 50 images in the dataset contain color cards.

In addition, all of the images were taken using a pair of DSLR cameras (Nikon D810 with an AF-S NIKKOR 35mm f/1.8G ED lens, encased in a Hugyfot housing with a dome port) on a rigid rig, as shown in Fig. 5(left). Based on stereo imaging we recover the distances of objects from the camera (except for occlusions) and we can quantitatively evaluate the transmission maps, which are estimated from a single image. Using raw images is crucial in order to process the linear intensity, since the signal of highly attenuated areas is stored in the least significant bits of each pixel. High resolution TIF image, raw images, camera calibration files, and the reconstructed scene distance maps can be downloaded from the project's webpage: http://csms.haifa.ac.il/profiles/tTreibitz/datasets/ambient_forwardlooking.

After arriving in each dive location and focusing the cameras, we took 20-30 different images of a checkerboard. See example in Fig. 5(right). Using those images and the MATLAB stereo calibration toolbox, we extracted both the intrinsic and extrinsic parameters of the stereo cameras. Given the intrinsic parameters, we corrected the lens distortions, to have pairs of images for which the pinhole camera model holds. The pinhole model holds since we use a dome port and a wide-angle lens, such that the lens entrance pupil is at the center of the dome, and there is no refraction at the interface between the water and the housing.

The next step was to find matching points in both images in order to triangulate the real-world data points. This task is quite difficult in the areas further from the camera, due to the low signal-to-noise-ratio. In order to find dense matching we used EpicFlow [50], and applied it in a bi-directional manner: from the left image to the right and vice-versa. We considered a match to be valid only for pixels that had an end-point-error of 5 pixels or less for images with resolution of 1827×2737 pixels. A resolution of over 5MPixels is good for practical use cases, and was chosen due to memory constraints using EpicFlow. The distance map was validated by measuring the size of the color charts in the scene, which is known to be $12.5\text{cm} \times 18\text{cm}$. Fig. 4 shows an example of an image from the dataset, with the dimensions of the color charts marked on the left, while the distance map of the scene is shown on the right.

6 EXPERIMENTS

6.1 Color Restoration - Qualitative comparison

Results are summarized in Figs. 6 and 8. The proposed color restoration method in Sec. 4 is based on a single image, and is evaluated as such.




Fig. 6. Comparison of single-image underwater enhancement techniques. For Ancuti [24] the reference image used for all results is shown as an insert in #R4376. In [24], [26] and ours the insert in #R3008 provides a zoom-in on the furthest chart. Order of dive sites from left to right: Med 1, Red 2, Med 2, Red 1.

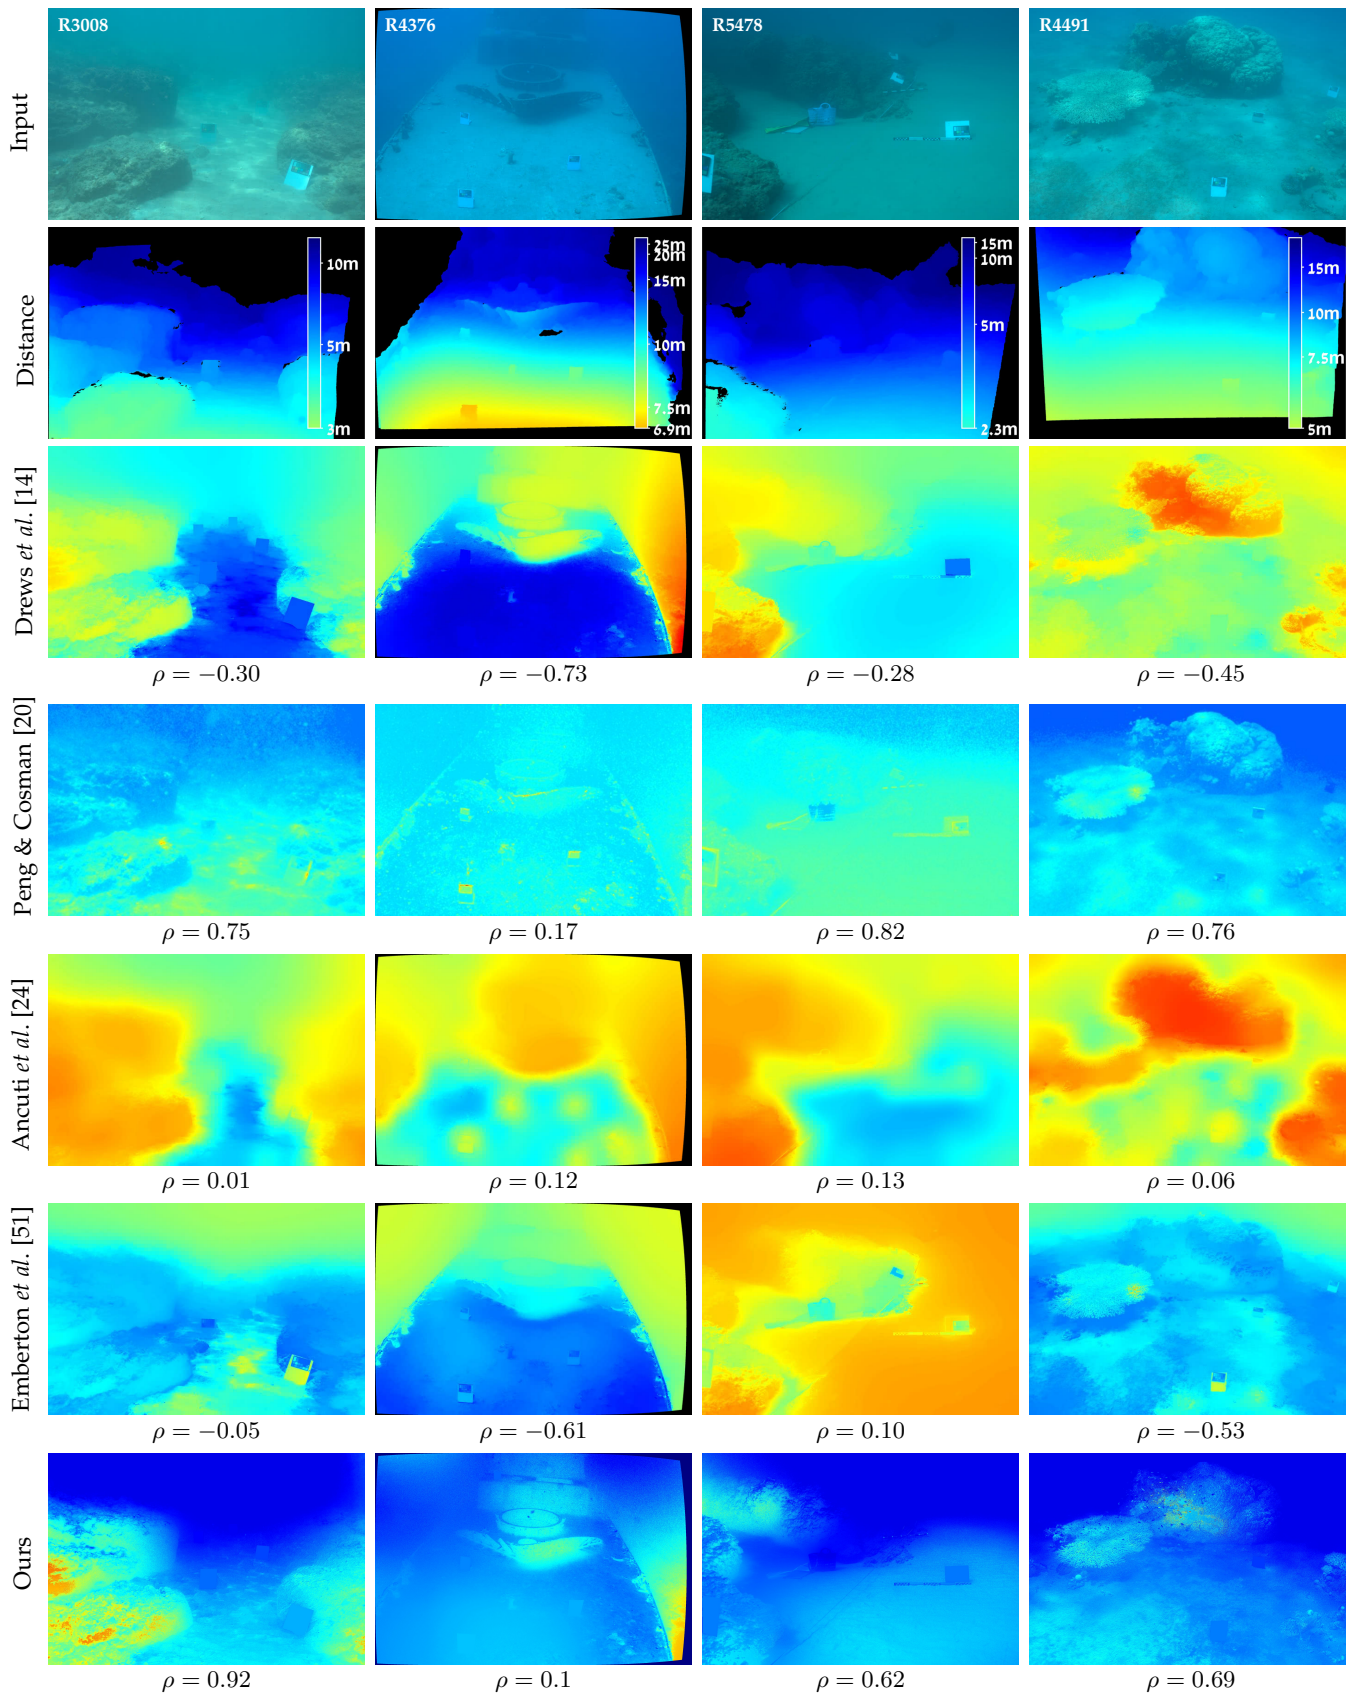


Fig. 7. True distances and estimated transmission maps for images shown in Fig. 6, shown only for methods that estimate a transmission map. The quality of the transmission is measured by the Pearson correlation coefficient ρ , which is calculated between the negative logarithm of the estimated transmission and the true distance.




Fig. 8. Comparison of single-image underwater enhancement techniques. Order of dive sites from left to right: Med 2, Red 1, Med 2, Med 1.

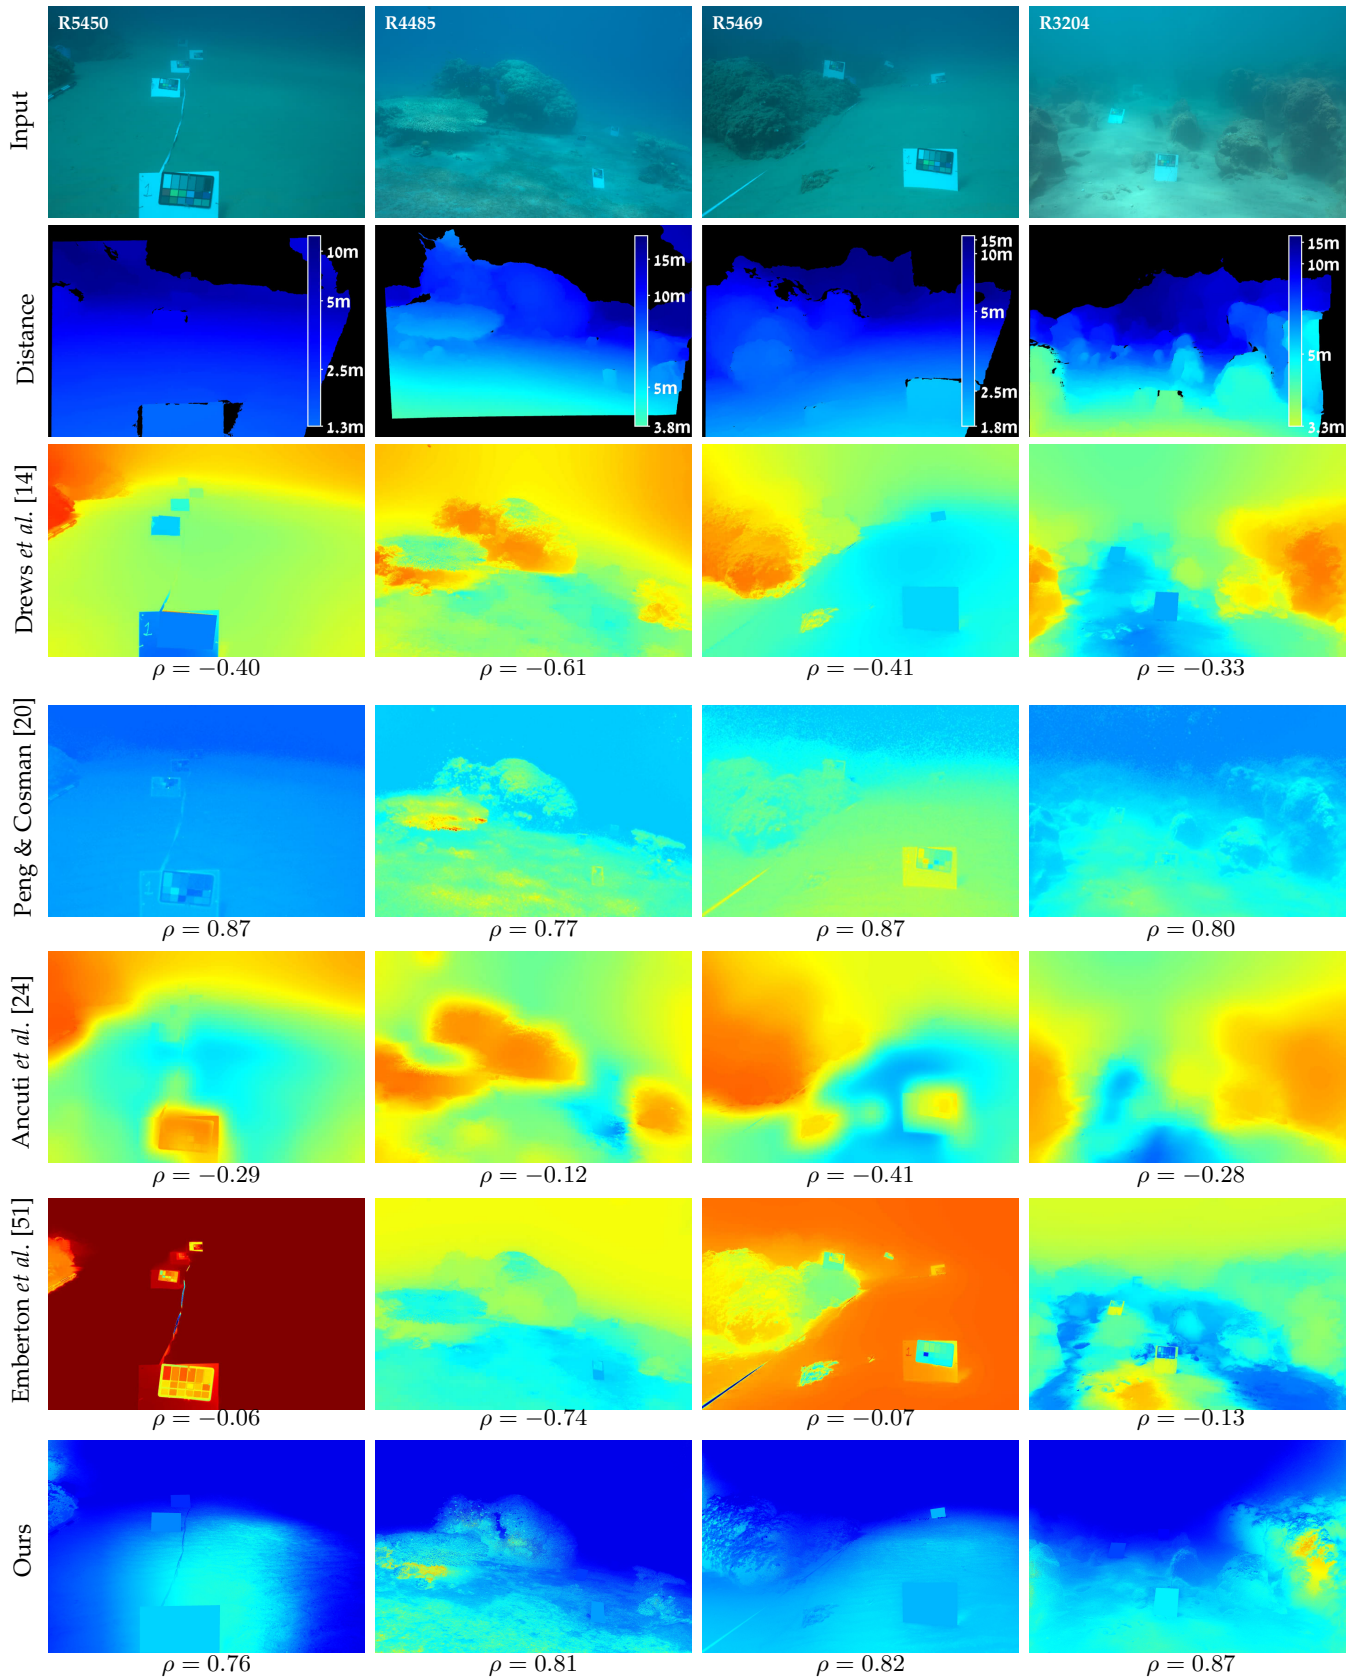


Fig. 9. True distances and estimated transmission maps for images shown in Fig. 8, shown only for methods that estimate a transmission map. The quality of the transmission is measured by the Pearson correlation coefficient ρ , which is calculated between the negative logarithm of the estimated transmission and the true distance.

TABLE 1

Each column shows the Pearson correlation coefficient ρ between the estimated transmission maps and the ground truth distance for a different image, as shown in Figs. 7, 9, and 12. Since the distance is calculated in meters, we calculate the correlation between the distance and $-\log(t)$.

Image	R3008	R4376	R5478	R4491	R5450	R4485	R5469	R3204	R3158	Med 1	Med 2	Red 1	Red 2
Drews <i>et al.</i> [14]	-0.28	-0.71	-0.40	-0.45	-0.46	-0.59	-0.43	-0.30	-0.26				
Peng & Cosman [20]	0.75	0.17	0.82	0.76	0.87	0.78	0.87	0.81	0.91	0.82	0.87	0.72	0.48
Ancuti <i>et al.</i> [24]	0.01	0.12	0.13	0.06	0.39	0.06	0.04	0.11	0.20				
Emberton <i>et al.</i> [51]	-0.04	-0.87	0.11	-0.54	-0.04	-0.74	-0.05	-0.12	0.06				
Ours	0.92	0.1	0.62	0.69	0.76	0.81	0.82	0.87	0.9	0.88	0.78	0.64	0.41

The color charts are used only for validation and are masked when fed to our algorithm. The transmission values of the masked pixels are determined based on neighboring values at the bottom of the chart. The stereo pair was *not* used for the color restoration, it was only used to generate the ground-truth distance map. This type of data enables us to quantitatively evaluate the accuracy of the estimated transmission. The code is available at: <https://github.com/danaberman/underwater-hl>.

Figs. 6 and 8 show the images we acquired on the top row, followed by the output of different underwater image enhancement methods. A naïve contrast stretch performed on each color channel separately is also included as a baseline. The results of Drews *et al.* [14] and Peng *et al.* [19] were generated using code released by the authors, results of Gao *et al.* [20], [26] generated by code sent in a private correspondence, and the results of Ancuti *et al.* [23], [24] and Emberton *et al.* [51] were provided by the authors. A recent thorough review [52] found that our method outperforms deep-learning methods and therefore we excluded comparison to them.

A qualitative comparison shows that methods [14], [20] are not able to consistently correct the colors in the images. These methods estimate the transmission and rely to some extent of the dark channel prior (DCP) which was proven empirically on outdoor scenes, but its assumptions do not always hold underwater [53]. The methods by Ancuti *et al.* [23], [24] focus on enhancing the image itself, rather than estimating an accurate transmission. Specifically, texture is extremely enhanced by [23], as evident by the details on the distant area of the shipwreck in Fig. 6. On the other hand, it also creates false textures on the foreground of the ship and in the sand in Fig. 6. The method by [26] produces results that look very good in the center of the scene but with halos at more distant areas. In [24], color transfer is used in conjunction with DCP-based transmission estimation to restore the color compensated image. The reference image for the color transfer is shown in the white inset on the left column of Fig. 6, and is identical for all of the images shown in Figs. 6 and 8. This method is rather effective in restoring the colors across different distances, as demonstrated for example by the color charts on the left column of Fig. 8.

6.2 Transmission Estimation - Quantitative Evaluation

Figs. 7 and 9 depict the ground-truth distances and the transmission maps for the images shown in Figs. 6 and 8, respectively. The true distance is shown in meters and is color-mapped, where black indicates missing distance values due to occlusions and unmatched pixels between the images. For example, the images shown here are from the right camera, and therefore many distance values are missing on the right border, since these regions were not visible in the

left camera. The transmission maps are shown only for the algorithms that estimate them [14], [20], [24], [51] and ours. These transmission maps are estimated from a single image and are evaluated quantitatively based on the true distance map of each scene. The transmission is a function of both the distance z and the unknown attenuation coefficient β_c : $t_c = e^{-\beta_c z}$. Therefore, we measure the correlation between the true distance z_{GT} and the negative logarithm of the estimated transmission ($-\log(t_c) = \beta_c z$). The Pearson correlation coefficient $\rho = \frac{\text{cov}(z_{GT}, -\log(t))}{\sigma_{z_{GT}} \sigma_{\log(t)}}$ is noted below each transmission map, and shows that the proposed method estimates a much more accurate transmission. In addition, the data is summarized in Table 1. Note that the images were undistorted to correct the lens distortion during the stereo calculation, and this is taken into account in the calculation of ρ .

The correlation coefficient has a value between -1 and $+1$, where $+1$ is a perfect correlation, indicating a good transmission estimation, while lower values indicate that the 3D structure of the scene was not estimated correctly. A negative value implies either a wrong estimation of the veiling light or that the used prior is not valid for that scene.

Some of the methods, e.g. Emberton *et al.* [51], do not aim for a physically valid transmission, as they set its value in veiling-light regions to be intentionally high in order to suppress noise enhancement. However, we still expect to see a correlation between the true distances and the transmission, since in veiling light regions the true distance are often missing. This is caused by the lack of texture in those regions, making it difficult to find matches between the left and right images and therefore impossible to triangulate. As Table 1 shows, in most methods the transmission is far off, except of [20] that slightly outperforms our method on average. However, the scene correction by [20] is extremely color-biased.

6.3 Color Restoration - Quantitative Evaluation

All of the scenes in Figs. 6 and 8 contain color charts at different distances from the camera. Fig. 10 shows a quantitative evaluation of the color restoration accuracy for two different images out of the entire set, while Table 2 lists the values for all of the images. The input images are shown on the left of Fig. 10, with the color charts marked. On the right, the bars depict the average reproduction angular error $\bar{\psi}$ (unit: degrees) [54] in RGB space between the grayscale patches and a pure gray color, for each chart, image, and method. For a given chart, the average angular error is defined as:

$$\bar{\psi} = \frac{1}{6} \cdot \sum_{\mathbf{x}_i} \cos^{-1} \left(\frac{\mathbf{I}(\mathbf{x}) \cdot (1, 1, 1)}{\|\mathbf{I}(\mathbf{x})\| \cdot \sqrt{3}} \right), \quad (11)$$

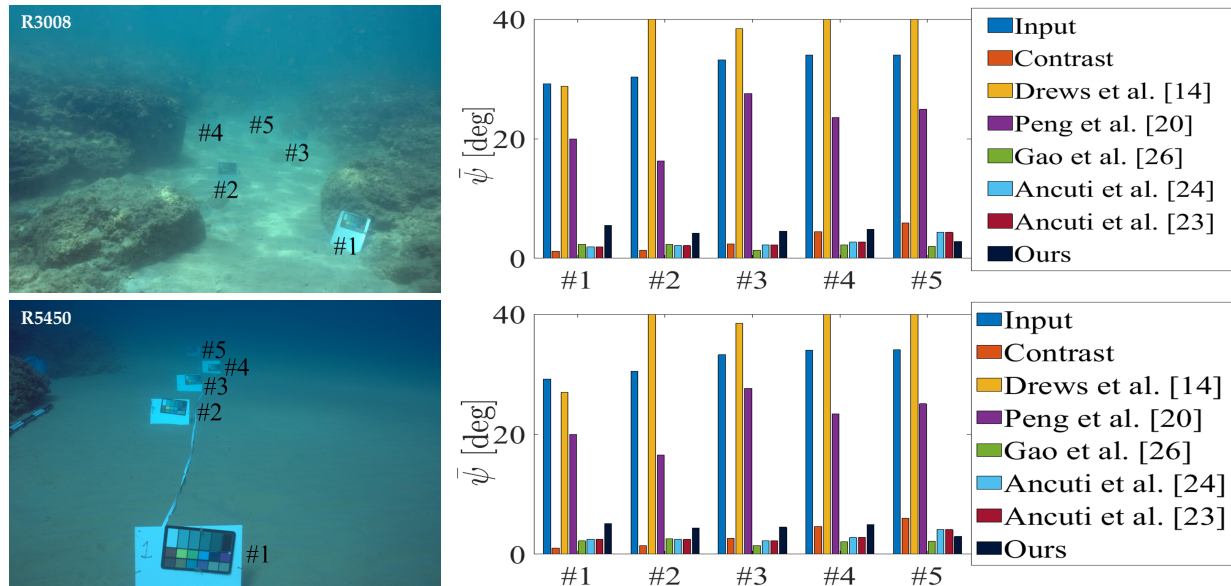


Fig. 10. Average angular reproduction error (in degrees) of gray-scale patches, measured for different color restoration techniques (shown in Figs. 6 and 8). The color charts are numbered in increasing order, starting with the closest to the camera, and are annotated on the images on the left column. The right column shows the average angular reproduction error $\bar{\psi}$ for each chart and each method. Lower is better.

where x_i mark the coordinates of the grayscale patches in the image plane (there are 6 of those). Lower angles indicate a more accurate color restoration. By calculating the angle we eliminate the influence of the global illumination on the brightness of each color patch and are able to calculate a robust measure on the neutral patches. The charts' order is by increasing distance from the camera, i.e., chart #1 is closest to the camera.

By comparing the input image to a naïve contrast stretch (the two leftmost bars for each chart), we see a good demonstration of the distance dependent degradation: while the contrast stretch angle (orange bar) is often the lowest for the closest chart (#1), it increases with distance, demonstrating that a distance-dependent correction is required. Table 2 summarizes the results. Note that the *contrast stretch often outperforms* most methods, especially on the closer charts. This is an important observation as many papers do not perform the comparison to a contrast stretch, while testing mainly on scenes with nearby objects. Table 2 separates the scores per table in selected images, and shows averages on the entire dataset on average of the charts, and only on the farthest one. Looking at the chart average, *contrast stretch outperforms all methods* as it performs extremely well on the nearby ones. Looking at the farthest one tells a different story, where our method is competitive with two visually-pleasing methods [24], [26] and outperforms all physics-based methods. The method in [24] requires a manually selected reference input image, while [26] produces some halos in the results. We suspect that some advantage of the visually pleasing methods partially stems from evaluation on the gray patches. This is seen for example, in the zoom-ins on #R3008 in Fig. 6. The farthest color charts in the results of [24], [26] clearly have worse contrast and color than ours although they yield a lower numerical error. This is because the veiling light makes the entire chart gray. We plan to devise a color metric on the charts in the future.

In addition, as expected, the color restoration significantly aids feature matching, demonstrated in Fig. 11.




Fig. 11. Color restoration significantly aids feature matching. [Top] SIFT feature matching performed on an input image. [Bottom] Running the same feature matching on our result yields a 5-fold increase in correct matches.

6.4 Water Type Effect

Here we want to emphasize the importance of estimating the water type. Fig. 12 shows on the leftmost column (top row) an image taken in murky waters at the Mediterranean Sea. The center column shows our output, which automatically estimated the water type to be C5, whereas the rightmost column shows the output assuming the image was taken in open ocean waters (type I). As evident, the color restoration with the wrong attenuation ratios leads to distorted colors, and the transmission estimation is much less accurate compared to the selected water type. Note, that the estimated Jerlov type are not always the expected ones and they are not always consistent within the same location. We suspect this is because we approximated the camera response to be a delta function (see Sec. 3.2). This was shown to be inaccurate in [40] and we plan to further explore this topic. Nevertheless, as demonstrated here, allowing for different water types has a significant positive effect on the results.

TABLE 2

The average reproduction angular error $\bar{\psi}$ in RGB space between the gray-scale patches and a pure gray color, for all each chart in each image, and all methods, including the input and global contrast stretch (labeled Cont.). The last rows summarize average scores over the entire dataset per dive site. 'Av' represents average over color cards, 'Last' represent the score on the farthest card. Lower is better.

Method	Input	Cont.	[14]	[20]	[26]	[24]	[23]	[51]	Ours
R3008 #1	29.18	1.65	26.74	1.47	3.24	2.25	5.00	5.14	6.23
R3008 #2	30.48	3.10	43.74	1.88	1.95	2.35	3.63	8.25	4.98
R3008 #3	33.24	4.47	38.20	6.69	2.03	2.27	4.39	5.88	5.25
R3008 #4	34.03	6.48	40.44	21.64	4.175	2.71	3.96	5.80	5.42
R3008 #5	34.08	6.45	40.05	16.60	2.5	4.21	4.91	5.76	3.68
R3204 #1	27.43	2.89	12.81	3.17	2.19	2.79	13.00	6.64	5.36
R3204 #2	34.10	4.87	32.03	9.50	3.62	3.04	4.24	5.08	0.98
R3204 #3	34.11	6.83	35.92	30.43	3.41	4.91	6.14	7.20	2.69
R3204 #4	34.21	7.61	35.66	31.75	5.43	4.76	3.21	7.29	1.99
R3204 #5	34.21	8.13	37.30	35.38	5.28	5.83	4.88	7.62	3.35
R4376 #1	35.17	1.95	37.00	30.29	31.19	6.58	17.05	39.10	25.28
R4376 #2	35.21	2.18	40.27	29.65	33.69	6.85	19.87	37.73	23.80
R4376 #3	35.13	1.87	40.86	31.26	18.61	5.55	16.62	37.04	23.10
R4376 #4	35.96	4.60	39.33	28.61	5.22	11.25	8.47	38.07	14.04
R4485 #1	34.37	1.37	40.39	28.64	4.44	1.74	13.89	36.46	7.78
R4485 #2	34.34	4.75	46.21	35.62	3.76	4.18	9.02	37.91	4.43
R4485 #3	34.54	7.95	42.41	35.64	5.91	3.12	5.91	35.91	1.41
R4491 #1	34.37	1.36	41.76	19.43	8.24	6.30	13.45	33.42	6.88
R4491 #2	34.24	2.55	45.78	26.14	7.38	5.70	10.59	34.26	4.06
R4491 #3	34.29	5.67	44.35	36.45	14.5	9.11	7.54	34.11	1.52
R5450 #1	34.31	1.54	28.25	25.46	8.88	6.73	7.81	27.69	6.32
R5450 #2	34.29	1.78	35.53	33.76	5.16	5.97	13.37	24.74	5.15
R5450 #3	34.45	2.84	36.56	35.56	8.04	7.33	8.90	30.89	0.84
R5450 #4	34.21	4.09	36.66	35.65	11.36	7.98	6.42	32.85	4.6
R5450 #5	34.50	6.03	40.29	36.14	12.15	9.54	4.72	34.48	3.46
R5469 #1	34.28	2.30	20.11	20.15	4.86	4.92	10.17	30.29	5.48
R5469 #2	34.36	5.16	20.67	31.11	8.6	6.26	13.82	34.82	6.16
R5469 #3	34.38	4.80	27.36	32.87	6.73	7.38	7.47	34.38	7.17
R5469 #4	34.42	6.63	29.80	35.26	10.25	7.05	3.09	34.75	5.2
R5478 #1	34.34	1.20	23.51	21.15	3.2	4.80	12.11	31.35	5.94
R5478 #2	34.42	1.23	27.30	21.075	3.21	6.28	5.53	35.28	5.96
R5478 #3	34.50	2.94	32.78	31.33	8.91	9.22	5.27	35.49	3.14
Average	33.91	3.98	35.00	25.61	8.06	5.59	8.58	25.49	6.49
Med 1 Av	33.19	8.85		28.35	4.40				8.66
Med 1 Last	35.09	9.74		18.39	4.38				6.96
Med 2 Av	35.44	4.41		34.75	10.20				6.45
Med 2 Last	35.55	6.95		29.01	8.50				5.06
Red 1 Av	35.57	5.88		35.51	8.03				6.38
Red 1 Last	35.67	8.82		32.50	8.10				10.26
Red 2 Av	36.59	3.05		36.24	9.11				25.92
Red 2 Last	36.84	2.34		36.03	6.10				24.15

6.5 Veiling Light Stability

Fig. 13 depicts the results of the veiling light estimation colored by dive site. We see that the results from different dive sites are clustered and there is more similarity of the values within dive sites compared to other dive sites. This suggests consistency of the veiling-light estimation. Within dive sites the scenes differ in depth and viewing angle therefore we do not expect the veiling-light to be identical.

6.6 Failure Cases

Generally, reasons for failure can be divided into two: failing to comply to the model's assumptions, or low signal. Common cases that break the model are a non-horizontal LOS (e.g. in #R4376, Fig. 6) and non-uniform illumination. Illumination becomes non-uniform when imaging in the early morning or late evening when the sun is angled with respect to the sea surface, under shadowed objects (e.g., Fig. 14[Top]), or in shallow water when surface caustics create illumination patterns [55]. As discussed in Sec. 4.4




Fig. 12. Top two rows: An image taken in murky waters is enhanced both by automatically estimating the water type (center column) and by assuming a specific water type which is wrong (right column). Bottom two rows: the enhanced output of different methods for comparison.




Fig. 13. Veiling light stability. Each point represents normalized veiling light value $V_c/(V_R + V_G + V_B)$, to eliminate intensity effects. We see that there is clustering between the different dive sites, indicating some consistency of the veiling-light estimation.

signal is lost in deep scenes, e.g., in our Red 2 scenes (Fig. 14[Bottom]), and in objects that are extremely far. These failure cases are common to all physics-based methods.

7 CONCLUSIONS

We expanded the haze-lines model to handle wavelength-dependent attenuation. Specifically, we recover underwater scenes by taking into account various water bodies that were classified by oceanographers according to their optical properties. We showed that recovering transmission for each color channel separately adds just two global parameters to the problem. By considering a comprehensive physical




Fig. 14. Example failure cases of our method. [Top] Scene from Red 1, failed because of non-uniform illumination. [Bottom] Scene from Red 2, failed because of the scene depth.

model we were able to reconstruct scenes with a complex 3D structure and correct the colors of objects that are farther away. Our method is not the top ranked in our quantitative criteria. However, it is the best one considering their combination, i.e., we are able to provide both good color restoration and transmission estimation, as opposed to all other methods.

Unlike terrestrial images the imaging model in Eq. 1 is over-simplified [32], [40]. This causes more failure cases in all physics-based methods, and gives some advantage to visually-pleasing methods. Nevertheless, when the model holds, physics-based methods achieve better results. This also makes underwater scenes more difficult to simulate despite recent efforts in this direction [30]. As a result, the evaluation of underwater image enhancement methods is more difficult. This can also explain the relative unsuccess of learning-based techniques in this field [52]. We collected a new *in situ* dataset of underwater images with ground truth and hope that our stereo dataset with color charts will advance the field of underwater image enhancement. Though the current database is limited to four locations, it still presents a variety of challenging scenes in different water depths. We plan to further expand the dataset to more locations and hope the methodology presented in it will inspire scientist in other locations in the world to collect similar data.

Recovering nearby objects in underwater images is relatively easy and as demonstrated in our results often a simple contrast stretch does the job. The problem becomes more and more difficult while the objects are further away. This is where differences between different methods are more pronounced. Unfortunately, many previous works use underwater images downloaded from the internet for evaluation. These images are biased to have more nearby objects, as the lay person who acquired them does not have immediate access to reconstruction methods, and has no incentive to include further parts with limited visibility. Therefore we believe our dataset will significantly progress the field of underwater image enhancement as it will push the evaluation of methods towards more challenging scenes, to increase real-world implications for oceanic research.

ACKNOWLEDGMENTS

TT was supported by the The Leona M. and Harry B. Helmsley Charitable Trust, The Maurice Hatter Foundation, Israel Science Foundation grant #680/18, Ministry of Science, Technology and Space grant #3 – 12487, and the Technion Ollendorff Minerva Center for Vision and Image Sciences. DB was supported by The Mediterranean Sea Research Center of Israel and by Apple Graduate Fellowship.

REFERENCES

- [1] C. D. Mobley, *Light and water: radiative transfer in natural waters*. Academic press, 1994.
- [2] N. G. Jerlov, *Marine optics*. Elsevier, 1976, vol. 14.
- [3] N. Carlevaris-Bianco, A. Mohan, and R. M. Eustice, "Initial results in underwater single image dehazing," in *Proc. IEEE/MTS Oceans*, 2010.
- [4] Y. Y. Schechner and N. Karpel, "Clear underwater vision," in *Proc. IEEE CVPR*, 2004.
- [5] —, "Recovery of underwater visibility and structure by polarization analysis," *IEEE J. Oceanic Engineering*, vol. 30, no. 3, pp. 570–587, 2005.
- [6] M. Bryson, M. Johnson-Roberson, O. Pizarro, and S. B. Williams, "Colour-consistent structure-from-motion models using underwater imagery," in *Robotics: Science and Systems*, 2012.
- [7] —, "True color correction of autonomous underwater vehicle imagery," *J. of Field Robotics*, vol. 33, no. 6, pp. 853–874, 2015.
- [8] A. Yamashita, M. Fujii, and T. Kaneko, "Color registration of underwater images for underwater sensing with consideration of light attenuation," in *IEEE. Conf. Robotics and Automation*, 2007, pp. 4570–4575.
- [9] A. Schwartzman, M. Alterman, R. Zamir, and Y. Y. Schechner, "Turbulence-induced 2d correlated image distortion," in *Proc. IEEE ICCP*, 2017.
- [10] D. Berman, T. Treibitz, and S. Avidan, "Non-local image dehazing," in *Proc. IEEE CVPR*, 2016.
- [11] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," in *Proc. IEEE CVPR*, 2009.
- [12] R. T. Tan, "Visibility in bad weather from a single image," in *Proc. IEEE CVPR*, 2008.
- [13] J. Y. Chiang and Y.-C. Chen, "Underwater image enhancement by wavelength compensation and dehazing," *IEEE Trans. Image Processing*, vol. 21, no. 4, pp. 1756–1769, 2012.
- [14] P. Drews, E. Nascimento, F. Moraes, S. Botelho, and M. Campos, "Transmission estimation in underwater single images," in *Proc. IEEE ICCV Underwater Vision Workshop*, 2013, pp. 825–830.
- [15] A. Galdran, D. Pardo, A. Picón, and A. Alvarez-Gila, "Automatic red-channel underwater image restoration," *J. of Visual Communication and Image Representation*, vol. 26, pp. 132–145, 2015.
- [16] H. Lu, Y. Li, L. Zhang, and S. Serikawa, "Contrast enhancement for images in turbid water," *JOSA A*, vol. 32, no. 5, pp. 886–893, 2015.
- [17] J. Perez, P. J. Sanz, M. Bryson, and S. B. Williams, "A benchmarking study on single image dehazing techniques for underwater autonomous vehicles," in *Proc. MTS/IEEE OCEANS*, 2017.
- [18] S. Emberton, L. Chittka, and A. Cavallaro, "Hierarchical rank-based veiling light estimation for underwater dehazing," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2015.
- [19] Y.-T. Peng, X. Zhao, and P. C. Cosman, "Single underwater image enhancement using depth estimation based on blurriness," in *Proc. IEEE ICIP*, 2015.
- [20] Y.-T. Peng and P. C. Cosman, "Underwater image restoration based on image blurriness and light absorption," *IEEE Transactions on Image Processing*, vol. 26, no. 4, pp. 1579–1594, April 2017.
- [21] C. Ancuti, C. O. Ancuti, C. De Vleeschouwer, R. Garcia, and A. C. Bovik, "Multi-scale underwater descattering," in *Proc. ICPR*, 2016.
- [22] C. Ancuti, C. O. Ancuti, T. Haber, and P. Bekaert, "Enhancing underwater images and videos by fusion," in *Proc. IEEE CVPR*, 2012, pp. 81–88.
- [23] C. O. Ancuti, C. Ancuti, C. De Vleeschouwer, and P. Bekaert, "Color balance and fusion for underwater image enhancement," *IEEE Transactions on Image Processing*, vol. 27, no. 1, pp. 379–393, 2018.

- [24] C. O. Ancuti, C. Ancuti, C. Devleeschouwer, L. Neumann, and R. Garcia, "Color transfer for underwater dehazing and depth estimation," in *Proc. IEEE ICIP*, 2017.
- [25] H. Lu, Y. Li, and S. Serikawa, "Underwater image enhancement using guided trigonometric bilateral filter and fast automatic color correction," in *Proc. IEEE ICIP*, 2013.
- [26] Y. Gao, J. Wang, H. Li, and L. Feng, "Underwater image enhancement and restoration based on local fusion," *Journal of Electronic Imaging*, vol. 28, no. 4, p. 043014, 2019.
- [27] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng, "AOD-Net: All-in-one dehazing network," in *Proc. IEEE ICCV*, 2017.
- [28] W. Ren, S. Liu, H. Zhang, J. Pan, X. Cao, and M.-H. Yang, "Single image dehazing via multi-scale convolutional neural networks," in *European Conference on Computer Vision*, 2016.
- [29] Y.-S. Shin, Y. Cho, G. Pandey, and A. Kim, "Estimation of ambient light and transmission map with common convolutional architecture," in *Proc. MTS/IEEE OCEANS*, 2016.
- [30] H. Blasinski, T. Lian, and J. Farrell, "Underwater image systems simulation," in *Imaging and Applied Optics Congress*. OSA, 2017.
- [31] A. Duarte, F. Codevilla, J. D. O. Gaya, and S. S. Botelho, "A dataset to evaluate underwater image restoration methods," in *Proc. MTS/IEEE OCEANS*, 2016.
- [32] D. Akkaynak and T. Treibitz, "A revised underwater image formation model," in *Proc. IEEE CVPR*, 2018, pp. 6723–6732.
- [33] —, "Sea-thru: A method for removing water from underwater images," in *Proc. IEEE CVPR*, 2019, pp. 1682–1691.
- [34] Y. Asano, Y. Zheng, K. Nishino, and I. Sato, "Shape from water: Bispectral light absorption for depth recovery," in *Proc. ECCV*, 2016.
- [35] M. Sheinin and Y. Y. Schechner, "The next best underwater view," in *Proc. IEEE CVPR*, 2016, pp. 3764–3773.
- [36] T. Morimoto, R. T. Tan, R. Kawakami, and K. Ikeuchi, "Estimating optical properties of layered surfaces using the spider model," in *Proc. IEEE CVPR*, 2010.
- [37] D. Akkaynak, T. Treibitz, B. Xiao, U. A. Gürkan, J. J. Allen, U. Demirci, and R. T. Hanlon, "Use of commercial off-the-shelf digital cameras for scientific data acquisition and scene-specific color calibration," *JOSA A*, vol. 31, no. 2, pp. 312–321, 2014.
- [38] R. Sumner, "Processing raw images in matlab," <https://users.soe.ucsc.edu/~rcsumner/rawguide/RAWguide.pdf>, 2014.
- [39] R. W. Austin and T. J. Petzold, "Spectral dependence of the diffuse attenuation coefficient of light in ocean waters," *Optical Engineering*, vol. 25, no. 3, pp. 253 471–253 471, 1986.
- [40] D. Akkaynak, T. Treibitz, T. Shlesinger, R. Tamir, Y. Loya, and D. Iluz, "What is the space of attenuation coefficients in underwater computer vision?" in *Proc. IEEE CVPR*, 2017.
- [41] J. Jiang, D. Liu, J. Gu, and S. Susstrunk, "What is the space of spectral sensitivity functions for digital color cameras?" in *Proc. IEEE Workshop Applications of Computer Vision (WACV)*, 2013, pp. 168–179.
- [42] P. Dollár and C. L. Zitnick, "Structured forests for fast edge detection," in *Proc. IEEE ICCV*, 2013.
- [43] K. He, J. Sun, and X. Tang, "Guided image filtering," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 6, pp. 1397–1409, 2013.
- [44] G. Buchsbaum, "A spatial processor model for object colour perception," *Journal of the Franklin Institute*, vol. 310, no. 1, pp. 1–26, 1980.
- [45] T. Treibitz and Y. Y. Schechner, "Resolution loss without imaging blur," *JOSA A*, vol. 29, no. 8, pp. 1516–1528, 2012.
- [46] K. Panetta, C. Gao, and S. Agaian, "Human-visual-system-inspired underwater image quality measures," *IEEE Journal of Oceanic Engineering*, vol. 41, no. 3, pp. 541–551, 2016.
- [47] M. Yang and A. Sowmya, "An underwater color image quality evaluation metric," *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 6062–6071, 2015.
- [48] C. Li, J. Guo, and C. Guo, "Emerging from water: Underwater image color correction based on weakly supervised color transfer," *IEEE Signal processing letters*, vol. 25, no. 3, pp. 323–327, 2018.
- [49] F. Ma and S. Karaman, "Sparse-to-dense: Depth prediction from sparse depth samples and a single image," *CoRR*, vol. abs/1709.07492, 2017. [Online]. Available: <http://arxiv.org/abs/1709.07492>
- [50] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid, "Epicflow: Edge-preserving interpolation of correspondences for optical flow," in *Proc. IEEE CVPR*, 2015.
- [51] S. Emberton, L. Chittka, and A. Cavallaro, "Underwater image and video dehazing with pure haze region segmentation," *Computer Vision and Image Understanding*, 2017.
- [52] S. Anwar and C. Li, "Diving deeper into underwater image enhancement: A survey," *arXiv preprint arXiv:1907.07863*, 2019.
- [53] D. Berman, T. Treibitz, and S. Avidan, "Diving into haze-lines: Color restoration of underwater images," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2017.
- [54] G. Finlayson and R. Zakizadeh, "Reproduction angular error: An improved performance metric for illuminant estimation," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2014.
- [55] Y. Swirski, Y. Y. Schechner, B. Herzberg, and S. Negahdaripour, "Stereo from flickering caustics," in *Proc. ICCV*, 2009, pp. 205–212.



Dana Berman Received her BSc degree in Electrical Engineering and her BA in Physics from the Technion-Israel Institute of Technology in 2007, and her MSc and PhD in Electrical Engineering from Tel Aviv University in 2013 and 2018, respectively.



Deborah Levy Received her BSc degree in Electrical Engineering from the Technion-Israel Institute of Technology in 2019. Currently she is a PhD Candidate at the Department of Marine Technologies, School of Marine Sciences, University of Haifa. She is also an SSI divecon.



Shai Avidan received the PhD degree from the School of Computer Science, Hebrew University, Jerusalem, Israel, in 1999. He is currently a professor on the Faculty of Engineering, Tel-Aviv University, Tel-Aviv, Israel. In between, he worked for Adobe, Mitsubishi Electric Research Labs, and Microsoft Research. He has published extensively in the fields of object tracking in video and 3D object modeling from images. He is also interested in Internet vision applications such as privacy-preserving image analysis, distributed algorithms for image analysis, and image retargeting.



Tali Treibitz received her BA degree in Computer Science and her PhD degree in Electrical Engineering from the Technion-Israel Institute of Technology in 2001 and 2010, respectively. Between 2010-2013 she has been a post-doctoral researcher in the department of computer science and engineering, in the University of California, San Diego and in the Marine Physical Lab in Scripps Institution of Oceanography. Since 2014 she is heading the marine imaging lab in the school of Marine Sciences in the University of Haifa. She was the recipient of the Google Anita Borg Scholarship in 2009 and the Weizmann Institute of Science National Postdoctoral Award for Advancing Women in Science in 2010. She is also a PADI diving instructor.