

## Attenuating Natural Flicker Patterns

Yoav Y. Schechner and Nir Karpel

Dept. of Electrical Engineering  
Technion - Israel Inst. Technology  
Haifa 32000, ISRAEL

yoav@ee.technion.ac.il , karpeln@tx.technion.ac.il

**Abstract**—Waves on the water surface create spatio-temporal illumination patterns underwater. Concave regions on the surface diverge light rays refracting into the water, while convex regions create convergence of rays (caustics). Therefore, the natural illumination of underwater objects is spatially varying. Moreover, in shallow water this nonuniform intensity distribution varies significantly in time, as the water surface changes with the wave motion. In this paper we present a method that attenuates these illumination patterns, and results in an image which appears as if taken under much more stable and uniform illumination. The method is based on just a few consecutive frames taken of the scene. These frames are analyzed by a non-linear algorithm which preserves consistent image components while filtering out fluctuations. The use of the nonlinear algorithm alleviates the need for long acquisition periods and is therefore fast. We demonstrate its effectiveness and efficiency in underwater experiments.

### I. INTRODUCTION

A familiar sight in shallow water is the rapid change of bright illumination patterns on the seabed. These spatiotemporal patterns modulate the reflectance and texture of the bottom. They may disrupt photography and human interpretation of the scene. Moreover, this flicker may seriously hinder computer vision, since some of those algorithms assume consistent illumination in space and time.

In this work we discuss a way to eliminate this visual phenomenon. The method of our choice is inspired by a recent computer vision algorithm for removal of shadows in open-air (topside) scenes. It thus provides an additional demonstration of the power of computer vision in non-acoustic underwater imaging applications [2][3][5][9][12][13][15][20][23][25][26].

### II. NATURAL FLICKER

Consider a short exposure time. During this brief period, the water surface is generally not flat, but rather wavy. Hence, refraction of natural light through this surface is inhomogeneous [21], as depicted in Fig. 1. Concave regions on the surface diverge light rays refracting underwater, while convex regions [19][21] create convergence (wave focusing [18]). This results in three dimensional (3D) patterns of variable irradiance in the water volume [8][17].

The refracting surface is far from being a perfect lens. Hence, the brightest regions in this 3D pattern are not simple focus points. Rather, singularities of light intensity are formed in *caustic networks* [14]. Any object surface (e.g., the seabed) slices through these 3D patterns. Therefore, the object illumination is spatially varying, with bright stripes caused by caustics [8].

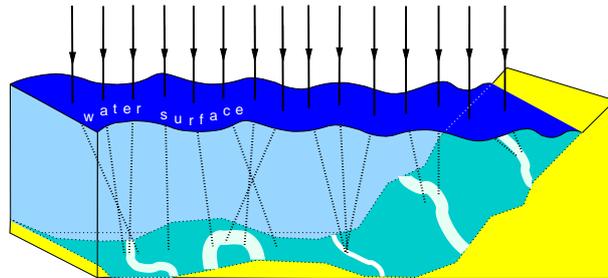


Fig. 1. Due to waves on the water surface, the refraction of natural sky illumination is spatially varying. This creates 3D patterns of variable light flux and caustics underwater, and 2D illumination patterns on the illuminated objects. These patterns vary in time due to the dynamics of the surface waves.

Moreover, the water surface changes with the wave motion [19]. Hence, the light patterns fluctuate in time [11][22]. These temporal variations are complex. Surface waves having long wavelengths move at different speeds than short surface waves [6][7][10]. This affects the caustic networks, because long surface waves focus light deeper into the water than the focusing by short surface waves. Moreover, the patterns do not relate linearly to the water wave spectrum: better fit to empirical data [22] is obtained by accounting for nonlinear interactions of light ray refraction angles relative to the surface state. This complexity is compounded by spatial “low pass filtering” effects caused by water turbidity [22].

### III. THE DOMAIN OF THE PROBLEM

Ref. [22] normalizes the amplitude (or standard deviation) of irradiance fluctuations by the mean of the irradiance at a given depth. The normalized fluctuations decay quickly with the underwater depth at which an object is illuminated. They decrease exponentially with the object depth, or even faster (depending on the sea model) [22]. The effective depth to which fluctuations are significant is scaled by the downwelling irradiance coefficient. Thus, the scaled intensity of the flicker is empirically and numerically [22] found to depend on the water turbidity: the significance of fluctuations in clear water light is maintained deeper than in turbid water. Thus, the problem is relevant for objects in relatively shallow depths, or in moderate depths when the water is clear.

Next, we need to know how fast the patterns vary. Temporally, the patterns are wide-band. Their spectral bandwidth is several Hz, while their peak frequency is in the order of 1Hz. Theoretic models [22] supported by

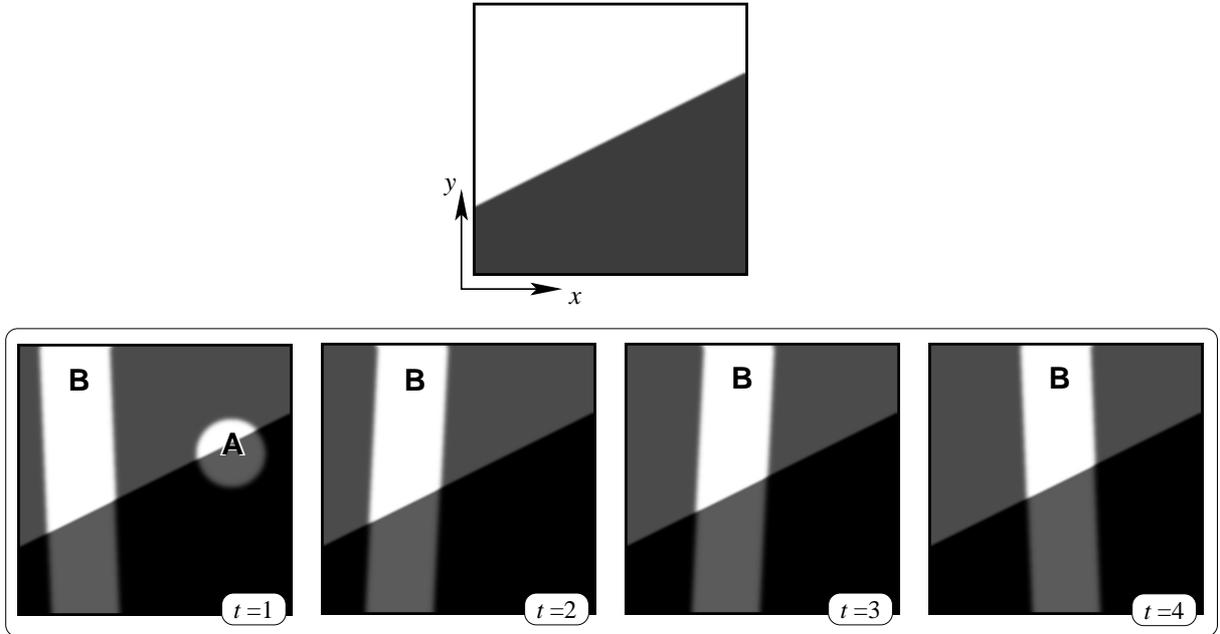


Fig. 2. [Top] The reflectance of a simulated seabed. [Bottom] Simulated consecutive raw frames.

empirical evidence show that the peak frequency behaves as

$$\omega_{\max} \propto \frac{1}{\sqrt{z}}, \quad (1)$$

where  $z$  is the object's depth under the water surface. For objects less than 2 meters deep, the peak frequency [22] of the illumination fluctuations is typically 1-3 Hz. When the objects are deeper, the fluctuations are slower, with an effective period of several seconds for the dominant spectral features of the fluctuations.

It follows from these orders of magnitude, that spatiotemporal patterns can be wiped out if the camera exposure time is set to several seconds. If images are extracted from a video stream, their exposure time is  $\mathcal{O}(30\text{ms})$ . Thus, if images of such a stream are averaged over an equivalent period of several seconds, then the spatial illumination patterns can be attenuated. However, such a brute force approach is slow. It requires the scene and camera to be stationary for seconds. This diminishes the usefulness of video cameras. It also greatly narrows the scope of dynamic scenes (e.g., of animals) that can be studied.

Thus, the domain of the problem is in cases where the temporal resolution should be at  $\mathcal{O}(0.1\text{seconds})$ , i.e., with the use of just a few quick frames. It is in this domain that a non-trivial algorithm is needed for attenuating the natural spatio-temporal illumination patterns.

#### IV. FAILURE OF TRIVIAL METHODS

Following the motivation given in Sec. III, we concentrate on cases where just a few consecutive frames are taken of the scene, at short intervals. In this section, we demonstrate scenarios in which trivial methods, such as averaging of frames, provide insufficient results. Let  $L_t(x, y)$  be consecutive raw frames, where  $t$  is the temporal (frame) index, while  $x$  and  $y$  are the spatial coordinates.

Two trivial processing methods are temporal averaging and temporal median filtering. In the first method, the calculated image is

$$L_{\text{mean}}(x, y) = \sum_{t=1}^T L_t(x, y), \quad (2)$$

where  $T$  is the number of raw frames. The second method calculates the median temporal value of each individual pixel<sup>1</sup>, independent of its neighbors:

$$L_{\text{median}}(x, y) = \text{median}_t L_t(x, y). \quad (3)$$

The latter method is more robust to spurious lighting conditions, since median filtering ignores outliers. It thus forces the output to rely on the intensity values that are most temporally consistent.

To illustrate the limitations of those methods, we first performed a simulation. The top of Fig. 2 depicts the reflectance distribution on a surface (a simulated seabed) made of a bright patch and a dark patch. Had the illumination been uniform and constant, the acquired images  $L_t(x, y)$  should have been proportional to this reflectance distribution. However, we are interested in simulating spatiotemporal variations of illumination. This is illustrated in the simulated raw frames<sup>2</sup> shown on the bottom of Fig. 2. The illumination is not uniform, and includes two bright features (caustics). Bright feature A appears briefly in only one of the frames. On the other hand, bright feature B moves and evolves between consecutive frames.

<sup>1</sup>No spatial median filtering is performed.

<sup>2</sup>For clarity of display, the brightness of all the displayed pictures in this paper underwent the same standard contrast enhancement (stretching). This applies to derivative-images as well. When displaying color images, their hue and color saturation components were untouched. The recovery algorithms, of course, used the raw (not brightness enhanced) images.

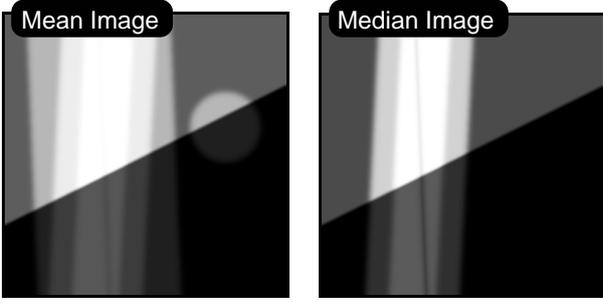


Fig. 3. The mean of the raw images shown in Fig. 2 and the median these images. Both results are not good, and do not recover the reflectance distribution shown on the top of Fig. 2.

The time interval between frames is small, thus the motion is slow. Hence, the bright and dark illumination patterns do not have enough time to get smeared all over the field of view. For this reason, the result of (2) is poor, as shown on the left of Fig. 3. The result of (3) is somewhat better, because feature A has been eliminated. This occurred since it was inconsistent in most frames. Nevertheless, as seen on the right of Fig. 3, the temporal median failed to eliminate the non-uniformity of the illumination, since evidence of feature B are still dominant. This occurs because for this feature *the movement between frames is small relative to the support of the feature*.

#### V. EFFECTIVE NON-UNIFORMITY ATTENUATION

The problem described above resembles problems of inhomogeneous illumination in open-air (topside) scenes. There, the cause is not refraction by the environmental medium, but simply *shadows* created by occlusion of light sources. In those scenarios, the nonuniform illumination varies very slowly in time, with the daily motion of the sun and clouds across the sky. A very effective “shadow removal” method has recently been proposed [16][24] for those cases. We thus propose to exploit it for our purposes.

Assume for the moment that veiling light (backscatter) is not a significant component in the range of object distances from the camera. In that case, the scene radiance is the product<sup>3</sup> of the spatiotemporal irradiance on the objects  $I_t(x, y)$  and the temporally constant reflectance  $R(x, y)$

$$L_t(x, y) = I_t(x, y) \cdot R(x, y) . \quad (4)$$

Taking the natural log of (4),

$$l_t(x, y) = i_t(x, y) + r(x, y) , \quad (5)$$

where

$$l \equiv \ln(L) , \quad (6)$$

while  $i \equiv \ln(I)$  and  $r \equiv \ln(R)$ . This operation enables the application of linear filtering, as convolution, to the illumination and the reflectance. Now, denote the derivatives along the  $x$  and  $y$  image coordinates as  $\partial_x l_t$  and  $\partial_y l_t$ , respectively:

$$\partial_x l_t(x, y) = \partial_x i_t(x, y) + \partial_x r(x, y) , \quad (7)$$

<sup>3</sup>There is a proportion factor between the scene radiance and image irradiance that depends on the imaging system, but does not depend on the illumination conditions. For this reason we leave this factor out.

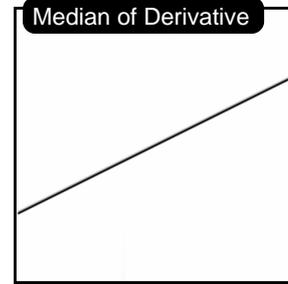
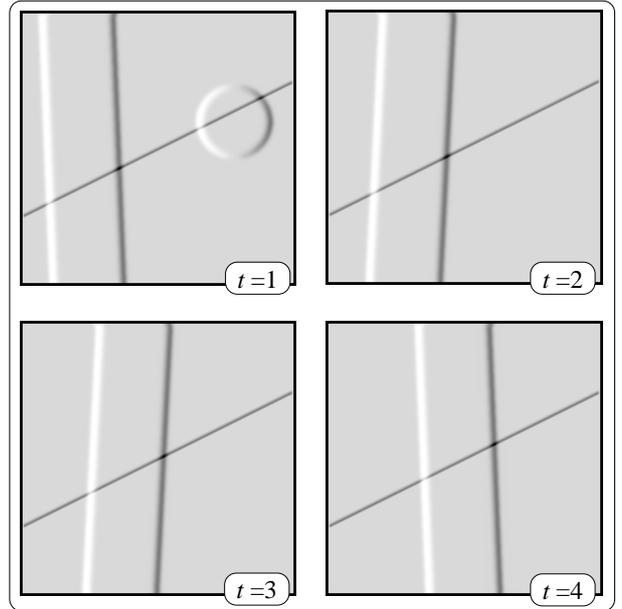


Fig. 4. [Top] The  $x$ -derivative applied to the raw images shown in Fig. 2. The results are sparse, as most pixels have low absolute values. [Bottom] Result of the algorithm performing a temporal median over spatial derivatives. The result is consistent with the derivative of the reflectance distribution function shown on the top of Fig. 2.

$$\partial_y l_t(x, y) = \partial_y i_t(x, y) + \partial_y r(x, y) . \quad (8)$$

The assumption made by [16][24] is that image derivatives associated with cast shadows are *sparse*, i.e.,  $\partial_x i_t(x, y)$  and  $\partial_y i_t(x, y)$  have very small values for most image pixels. We make an analogous assumption here. The spatial support of image regions having a high value of  $|\partial_x i_t(x, y)|$  or  $|\partial_y i_t(x, y)|$  is small. Thus, in contrast to Sec. IV, even a small motion of caustic features is *large relative to the small support of the derivative features*.

To see this, we applied (7) to the raw images shown in Fig. 2. To return to the intensity derivative range from the log domain, define  $\tilde{L}_t^{\partial_x} \equiv \exp(\partial_x l_t)$ . The top images in Fig. 4 show the resulting  $\tilde{L}_t^{\partial_x}$  for all frames  $t = 1 \dots 4$ . In these images, dark pixels indicate highly negative values of  $\partial_x l_t$  while bright pixels indicate highly positive values. It is clear from these images, that high absolute values of image derivatives are confined to small areas.

Thanks to this sparsity, temporal median filtering of these fields is very effective in eliminating the inconsistent illumination-derivative features, while maintaining the

temporally constant derivatives of the reflectance. We wish to estimate the log-reflectance used in (5). The estimate of its derivatives is thus

$$[\partial_x \hat{r}(x, y)]_{\text{median}} = \text{median}_t \partial_x l_t(x, y) , \quad (9)$$

$$[\partial_y \hat{r}(x, y)]_{\text{median}} = \text{median}_t \partial_y l_t(x, y) . \quad (10)$$

To demonstrate how effective this method is, define  $\tilde{R}_{\text{median}}^{\partial_x} \equiv \exp([\partial_x \hat{r}]_{\text{median}})$ . The result is shown in the bottom of Fig. 4. Indeed, the features associated with the spatiotemporal illumination patterns have been mostly eliminated. On the other hand, the method maintained the derivative associated with the reflectance edge shown on the top of Fig. 2.

Now that we have “clean” derivative fields in  $x$  and  $y$ , they should be integrated to yield an estimated  $\hat{r}(x, y)$ , leading subsequently to the reflectance image  $\hat{R}(x, y)$ . Note that for  $N$  pixels, we have  $2N$  constraints, as expressed by (9,10). A solution in the least-squares sense is obtained by a pseudo-inverse [24]. It is given by the following process. First calculate the term  $D(x, y)$

$$D(x, y) = [\partial_x \hat{r}(x, y)] * d_x^{\text{rev}} + [\partial_y \hat{r}(x, y)] * d_y^{\text{rev}} , \quad (11)$$

where  $*$  denotes convolution. Here  $d_x^{\text{rev}}$  and  $d_y^{\text{rev}}$  are the reversed derivative kernels, e.g., if the derivative kernel  $\partial_x \equiv [1 \ 0 \ -1]/2$ , then  $d_x^{\text{rev}} = [-1 \ 0 \ 1]/2$ . Then, the estimate we seek is

$$\hat{r}(x, y) = W * D(x, y) , \quad (12)$$

where the linear operator  $W$  is the one which solves

$$W * [d_x^{\text{rev}} * \partial_x + d_y^{\text{rev}} * \partial_y] = \delta(x, y) , \quad (13)$$

while  $\delta$  is the delta-function. Note that  $W$  is independent of the image sequence, and is thus pre-computed. For ease of implementation, expressions (12,13) are calculated in the frequency domain [24]. Finally, the estimated reflectance is

$$\hat{R}(x, y) = \exp[\hat{r}(x, y)] . \quad (14)$$

## VI. ADDITIONAL ASPECTS

### A. Stability

The definition of  $W$  in (13) is problematic. The Fourier transforms of derivative operators have a zero value at the zero-frequency (DC) component. Hence, the definition of this component of  $W$  is ill posed. To stabilize this numerical calculation, the value of this component should be set to an arbitrary finite value, e.g., 1.

A consequence of this issue is that the DC component of  $\hat{r}(x, y)$  cannot be determined, if recovery is based solely on (9)-(13). We will further deal with this issue in the next sections.

As typical in inverse filtering approaches, the Fourier transform of  $W$  has very high values near a spectral singularity, as appears here in the DC component. Thus (12) greatly amplifies the low frequency components of  $D(x, y)$ . One may suspect that this leads to amplified noise in  $\hat{r}(x, y)$ . Fortunately, however, this is not the case. Thanks to the temporal median filtering, the signals  $\partial_x \hat{r}(x, y)$  and

$\partial_y \hat{r}(x, y)$  defined in (9,10) have *less* noise than the derivatives of the original signals  $l_t(x, y)$ . Thus the integration of the filtered derivatives as realized by (11,12) results in lower noise.

### B. Boundary Conditions

The method (11,12) essentially seeks a solution to a set of partial differential equations: we need to find  $\hat{r}(x, y)$  that satisfies equations dealing with its derivatives, given by (9,10). As in any differential equation problem, there is a need to set boundary conditions in order to obtain a unique solution. This is achieved by the following recipe.

- 1) For each raw frame  $t$ , calculate  $l_t(x, y)$  as in (6).
- 2) Pad  $l_t(x, y)$  from all sides by a margin of pixels having a value of zero. The width of these margins is at least as large as the support of the convolution kernels used subsequently. This sets the boundary condition prior to spatial filtering.
- 3) Operate (7)-(9) on the padded frames.
- 4) Perform the recovery (11,12).
- 5) Crop the margin pixels.

The convolution (11) in step 4 above is an ordinary convolution, which implicitly uses zeros as boundary values for the image. It thus enforces the original boundary values on the output.

### C. Processing of Color

The method is based on image derivative operations as in (7,8). The derivative operation erases any information about the mean value (DC) of the log-images, consequently leading to an ill posed value for this component in the estimated log-reflectance. This implies an ambiguous scale factor in the estimated reflectance (14).

There is a need to disambiguate the solution for the following reasons. If the images are in gray-scale and intended solely for display, then the recovery of the intensity scale is not critical, since the contrast of the display can be stretched at will. The situation is different, however, in color. The method described in Sec. V is applied to each color channel independently. It is then important to recover the energy of each color channel, so that the colors of the scene will not be distorted. For example, the green channel should not be amplified or biased relative to the blue channel. Moreover, even in grayscale images, a proper estimate of the intensity scale is important for image-based photometric calculations [20].

We determine the scale factor of (14) based simply on the total amount of light detected in the raw images. For example, for the red channel, define

$$e_t^{\text{red}} = \sum_{x,y} L_t^{\text{red}}(x, y) \quad E^{\text{red}} = \frac{1}{T} \sum_{t=1}^T e_t^{\text{red}} . \quad (15)$$

Then, take the result  $R^{\text{red}}(x, y)$  obtained by (14) and correct it by scaling:

$$\hat{R}_{\text{corrected}}^{\text{red}}(x, y) = \alpha^{\text{red}} R^{\text{red}}(x, y) , \quad (16)$$

where

$$\alpha^{\text{red}} = \frac{E^{\text{red}}}{\sum_{x,y} \hat{R}^{\text{red}}(x, y)} . \quad (17)$$

Analogous operations are performed over each color channel.

We note that  $e_t$  for each color channel (as in 15) has only a slight dependence on  $t$ . The different surface states of the water at different times  $t$  hardly change the overall light energy penetrating into the water. This occurs despite the temporal changes of the spatial distribution of light. Therefore, the image resulting from this scaling operation represents well the scene under the illumination conditions existing during acquisition, excluding the flicker. This is contrary to topside shadows, for which this algorithm had originally been developed. In open-air shadows, there may be significant temporal variations of the total amount of the light falling on the scene in the field of view, as well as variations of the illumination color [4].

#### D. Veiling Light

The image formation model in (4) suits imaging through clear media, as air. Underwater, however, there exists backscatter  $B(x, y)$  of the natural illumination into the line of sight (veiling light) [20]. A more accurate model of the scene radiance is

$$L_t(x, y) = S_t(x, y) + B_t(x, y) , \quad (18)$$

where

$$S_t(x, y) = I_t(x, y) \cdot R(x, y) , \quad (19)$$

is the signal coming from the object. The backscatter varies in time since it is directly related to the illumination patterns in the 3D water volume. These variations are weak, because backscatter integrates the illumination distribution along the line of sight, thus smoothing out spatial illumination fluctuations along this path. This is contrary to  $S_t(x, y)$  whose values depend on the illumination at particular points in space. The backscatter increases with the distance of the objects from the camera [9][20].

As before, we applied (6) to the raw images. Due to  $B_t(x, y)$ , this operation does not turn reflectance and irradiance into additive quantities, contrary to (5). Nevertheless, when we applied the method (9)-(17) in real underwater scenarios, the results were visually pleasing, as we will show in the next section. Apparently, the method is robust to such deviations. However, we should now note that the image we obtain in the end is *not* an estimate of the reflectance  $R(x, y)$ , since the backscatter still influences the result. What we indeed derive is an estimate of the scene's appearance  $L(x, y)$ , including veiling light effects, yet attenuating the spatiotemporal illumination patterns.

## VII. EXPERIMENTS

We performed several experiments, both at sea and in a swimming pool. We took the images using the Nikon-D100 camera, since it has a linear radiometric response. The camera was placed in a housing, such that all of the images were taken underwater. Seven photographs of a pool bottom were taken in shallow water. Samples of these frames are shown in Fig. 5. The algorithm described in sections V and VI significantly improves the uniformity, as shown in Fig. 6. This result is superior to the mean image and the median image shown in Fig. 7.

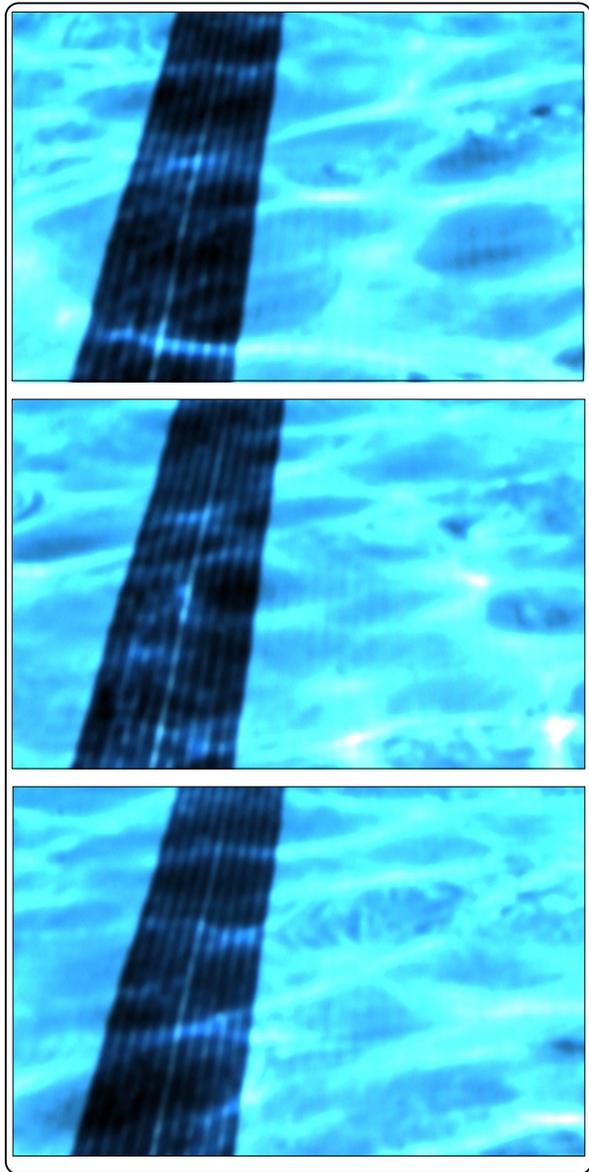


Fig. 5. Raw images of a swimming pool bottom. They exhibit spatiotemporal illumination patterns (caustic network). They are the result of due to light refraction effects caused by waves on the water surface.



Fig. 6. The spatiotemporal patterns are greatly attenuated using the algorithm that is based on temporal median filtering of spatial derivatives.

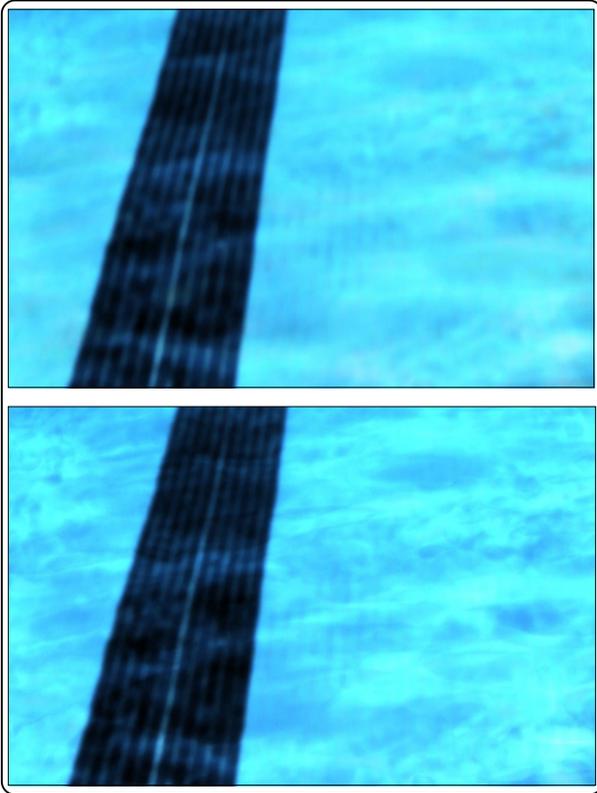


Fig. 7. Trivial processing of the raw images shown in Fig. 5. The result of averaging frames [Top] or taking their temporal median [Bottom] has only a limited success.

We then performed a similar experiment in the sea. Three raw images were acquired, shown in Fig. 8. The result of the algorithm described in sections V and VI is shown in Fig. 9. The result appears as if taken under a flat water surface or using a very long exposure, despite the fact that only three short-exposures were used.

### VIII. CONCLUSIONS

The method we propose for attenuating natural flicker patterns proved effective in several experiments. There is still a place for additional experiments and a study of the limitations of this method. Perhaps even better solutions will come up. We believe that this method can add up to existing enhancement tools for underwater photography [1][3][15][20], with applications to marine biology research, systems for search and rescue [13], and autonomous vehicles operating in shallow water.

### ACKNOWLEDGMENTS

We are grateful to Avishai Adler and Ori Bryt for their insights. We also wish to thank Sarit Shwartz and Yoav Fhiler for enabling the photography in the Technion's swimming pool. Yoav Schechner is a Landau Fellow - supported by the Taub Foundation, and an Alon Fellow. The work was supported by the Israeli Science Foundation (grant No. 315/04) and by the Ollendorff Center in the Elect. Eng. Dept. at the Technion. Minerva is funded through the BMBF.

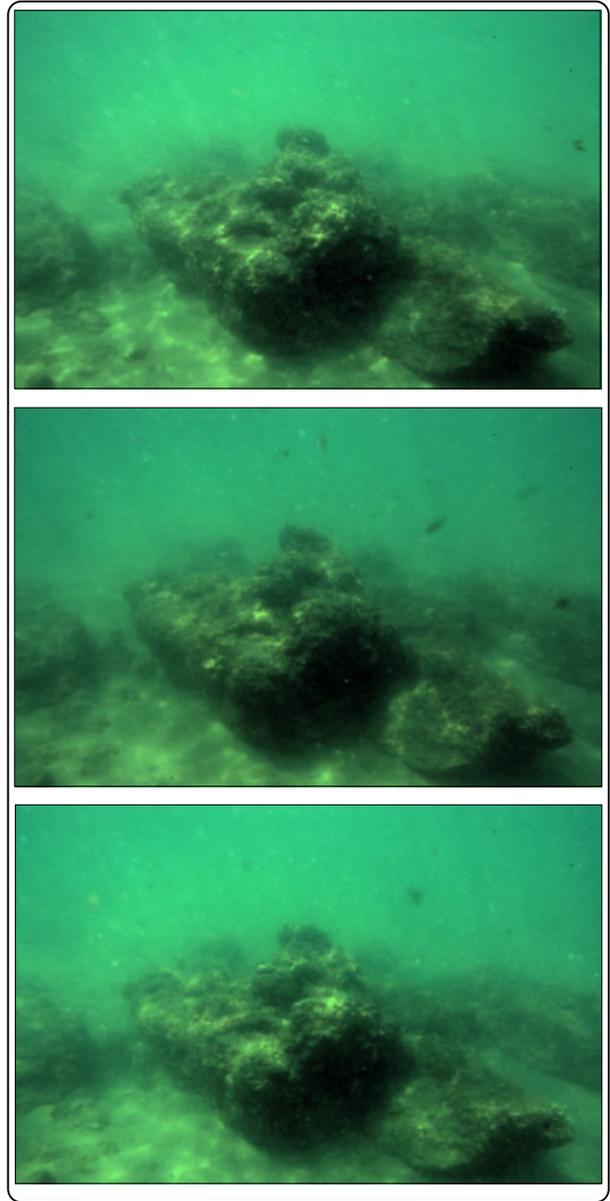


Fig. 8. The raw images taken in the sea at shallow water. They exhibit spatiotemporal illumination patterns, due to light refraction effects caused by waves on the water surface.



Fig. 9. The spatiotemporal patterns are practically eliminated using the algorithm based on temporal median filtering of spatial derivatives. The image looks as if taken under a flat water surface.

## REFERENCES

- [1] J. Åhlén and D. Sundgren, "Bottom reflectance influence on a color correction algorithm for underwater images," *Proc. Scandinavian Conf. Image Analysis*, pp. 922-926, 2003.
- [2] T. Boulton, "DOVE: Dolphin omni-directional video equipment," *Proc. Int. Conf. Robotics & Autom.* pp. 214-220, 2000.
- [3] R. Eustice, O. Pizarro, H. Singh, and J. Howland, "UWIT: Underwater Image Toolbox for optical image processing and mosaicking in MATLAB," *Proc. IEEE Int. Sympos. on Underwater Technology*, pp. 141-145, 2002.
- [4] G. D. Finlayson, S. D. Hordley and M. S. Drew, "Removing shadows from images," *Proc. European Conf. Computer Vision*, pp. 823-836 *Lecture Notes in Computer Science* vol. 2353, 2002.
- [5] G. L. Foresti, "Visual inspection of sea bottom structures by an autonomous underwater vehicle," *IEEE Trans. Syst. Man and Cyber, Part B*, vol. 31, 691-705 (2001).
- [6] A. Fournier and W. T. Reeves, "A simple model of ocean waves," *Proc. SIGGRAPH*, pp. 75-84, 1986.
- [7] M. N. Gamito and F. K. Musgrave, "An accurate model of wave refraction over shallow water," *Computers & Graphics*, vol. 26, pp. 291-307, 2002.
- [8] K. Iwasaki, Y. Dobashi and T. Nishita, "An efficient method for rendering underwater optical effects using graphics hardware," *Compt. Graphics Forum*, vol. 21, pp. 701-711, 2002.
- [9] J. S. Jaffe, "Computer modeling and the design of optimal underwater imaging systems," *IEEE J. Oceanic Engineering*, vol. 15, pp. 101-111, 1990.
- [10] J. V. Jelley, "Sea waves: their nature, behaviour, and practical importance," *Endeavour, New Series* vol. 13, pp. 148-156, 1989.
- [11] N. G. Jerlov, *Marine Optics*, Ch. 6, Elsevier, Amsterdam, 1976.
- [12] D. M. Kocak and F. M. Caimi, "Computer vision in ocean engineering," *The Ocean Eng. Handbook*, editor F. El-Hawari, Ch. 4.3, CRC Press, 2001.
- [13] J. M. Lavest, F. Guichard and C. Rousseau, "Multiview reconstruction combining underwater and air sensors," *Proc. IEEE Int. Conf. Image Processing* vol. 3, pp. 813-816, 2002.
- [14] D. K. Lynch and W. Livingston, *Color and light in nature*, 2<sup>nd</sup> ed., Cambridge University Press, 2001.
- [15] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall and M. Bolas, "Synthetic aperture confocal imaging," *Proc. SIGGRAPH* (2004).
- [16] Y. Matsushita, K. Nishino, K. Ikeuchi and M. Sakauchi, "Shadow elimination for robust video surveillance," *Proc. IEEE Worksh. Motion and Video Comp.*, pp. 15-21, 2002.
- [17] M. G. J. Minnaert *Light and color in the outdoors*, 1974. Translated by L. Seymour, Springer-Verlag, New York, 1993.
- [18] C. D. Mobley, *Light and Water: Radiative Transfer in Natural Waters*, Ch. 4 (Academic Press, San-Diego 1994).
- [19] H. Murase "Surface shape reconstruction of a nonrigid transparent object using refraction and motion," *IEEE Trans. of Pattern Analysis and Machine Intell.* vol. 14, pp. 1045-1052, 1992.
- [20] Y. Y. Schechner and N. Karpel "Recovering scenes by polarization analysis," *Proc. MTS/IEEE Oceans* (2004).
- [21] P. Y. Ts'o and B. A. Barski, "Modeling and rendering waves: wave-tracing using beta-splines and reflective and refractive texture mapping," *ACM Trans. on Graphics*, vol. 6/3, pp. 191-214, 1987.
- [22] R. E. Walker, *Marine light field statistics*, Ch. 10, John Wiley, New York, 1994.
- [23] D. Walther, D. R. Edgington and C. Koch, "Detection and tracking of objects in underwater video," *Proc. IEEE Int. Conf. Computer Vision*, vol. 1, pp. 544-549, 2004.
- [24] Y. Weiss, "Deriving intrinsic images from image sequences," *Proc. Int. Conf. on Computer Vision*, Vol. II, pp. 68-75, 2001.
- [25] L. B. Wolff, "Polarization vision: a new sensory approach to image understanding," *Image & Vision Computing* vol. 15, pp. 81-93, 1997.
- [26] X. Xu and S. Negahdaripour, "Automatic optical station keeping and navigation of an ROV: sea trial experiment," *Proc. Oceans*, vol.1, pp. 71-76, 1999.