

Regularized Image Recovery in Scattering Media

Yoav Y. Schechner, *Member, IEEE*, and
Yuval Averbuch

Abstract—When imaging in scattering media, visibility degrades as objects become more distant. Visibility can be significantly restored by computer vision methods that account for physical processes occurring during image formation. Nevertheless, such recovery is prone to noise amplification in pixels corresponding to distant objects, where the medium transmittance is low. We present an adaptive filtering approach that counters the above problems: While significantly improving visibility relative to raw images, it inhibits noise amplification. Essentially, the recovery formulation is regularized, where the regularization adapts to the spatially varying medium transmittance. Thus, this regularization does not blur close objects. We demonstrate the approach in atmospheric and underwater experiments, based on an automatic method for determining the medium transmittance.

Index Terms—Color, polarization, vision in bad weather, inverse problems, dehazing.

1 INTRODUCTION

IMAGING in turbid media such as fog, haze, and water suffers from poor visibility. In such media, the radiance from a scene point is significantly altered: It decays with distance and is compounded by stray light, which increases with distance. Recovering visual information in poor visibility is important for various machine vision and user-assisted applications such as oceanic engineering and mapping, surveillance, navigation, and flight safety. There is, thus, a growing interest in the analysis of such scenes [1], [7], [10], [14]. Several methods have been proposed to restore good visibility. Some of them use frames of the same scene taken under different weather conditions [10]. Others use active illumination hardware [6], [8], [21]. Effective inversion of image degradation can be performed based on analysis of frames taken with different states of a polarizing filter, as shown in haze [17], [19] and underwater [16], [21]. Recovery can also be performed using a model of the distance map of the scene, coupled to estimation of the atmospheric parameters. This parameter estimation may be obtained using statistical priors, which account for color variations and camera noise [14].

Whatever the method for compensating for the effects of the medium, the need arises for suppression of noise in the recovered image. The object radiance decays exponentially with the object distance. Thus, noise is amplified when attempting to invert the degrading scattering effects. Hence, noise in the restored images depends on distance. Far objects suffer from noise much more than close objects. We wish to suppress the generated noise. Due to the physical nature of the noise amplification, the noise suppression should be adaptive to the distance map. This map may be based on various machine vision methods [6], [8], [13], [14] or be set interactively by a user [11], [20]. It may also be estimated by analysis of frames taken through a polarizing filter [16], [17]. This paper proposes an adaptive digital filtering approach which achieves this goal. It restores visibility while adaptively suppressing noise, following changes of object distances. The approach is successfully applied to color images taken in atmospheric haze and underwater.

- The authors are with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel. E-mail: yoav@ee.technion.ac.il, averbuch@tx.technion.ac.il.

Manuscript received 18 Aug. 2005; revised 7 Apr. 2006; accepted 28 Nov. 2006; published online 3 Apr. 2007.

Recommended for acceptance by G. Finlayson.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0448-0805. Digital Object Identifier no. 10.1109/TPAMI.2007.1141.

2 TURBIDITY, INVERSION, AND NOISE

To make the paper self-contained, we briefly mention the imaging model and its known inversion. For simplicity, we mainly refer to processes occurring in the atmosphere, but the analysis suits other environments, particularly underwater. When using incoherent light in a scattering medium, the measured intensity at pixel position (x, y) is

$$I(x, y) = S(x, y) + A(x, y), \quad (1)$$

where S is the *signal* originating from the object. Here, A is stray illumination, termed *airlight* [17] in the context of atmospheric imaging,¹ scattered into the line of sight (LOS). The signal is related to the radiance of the object² in sight $L_{\text{object}}(x, y)$ by

$$S(x, y) = L_{\text{object}}t(z), \quad (2)$$

where $t(z)$ is the medium transmittance. Here, $z = z(x, y)$ is the object distance. For media having homogeneous extinction characterized by the coefficient β ,

$$t(z) = \exp(-\beta z). \quad (3)$$

Scene recovery has been achieved [10], [14], [16], [17] by a simple inversion of (1) and (2),

$$\hat{L}_{\text{object}}(x, y) = [I(x, y) - A(x, y)]/t(x, y), \quad (4)$$

requiring estimates of $t(x, y)$ and $A(x, y)$. There are several methods for recovering these two functions. Assuming uniform illumination over the LOS [17] (which is reasonable in natural lighting), we may derive the transmittance if we know the airlight using

$$\hat{t}(x, y) = 1 - A(x, y)/A_{\infty}, \quad (5)$$

where A_{∞} is a global parameter. In haze and water, $A(x, y)$ has been estimated by analyzing frames acquired through a polarizer [16], [17], [19], [21], corresponding to two orientations of the polarizing³ filter. Let $I^{\text{min}}(x, y)$ be the frame having the least amount of airlight. The other frame, $I^{\text{max}}(x, y)$, is brighter due to an increased airlight component. Then,

$$\hat{A}(x, y) = [I^{\text{max}}(x, y) - I^{\text{min}}(x, y)]/p, \quad (6)$$

where $p \in [0, 1]$ is the airlight degree of polarization. Both A_{∞} and p do not vary significantly across narrow fields of view and can be extracted based on the image data [17], [19].

Inversion methods such as (4) can be very effective in recovering scene contrast. Contrast between two points $v = 1, 2$ is usually defined [4] as $|I_1 - I_2|/|I_1 + I_2|$, where I_v is the intensity at point v . This definition is sometimes referred to as *visibility* [3]. In [16], the contrast definition was generalized to regions in color images, as

$$C(I) = \frac{\sqrt{\frac{1}{N} \sum_{\chi=R,G,B} \sum_{x,y} [I^{\chi}(x, y) - \bar{I}^{\chi}]^2}}{\sum_{\chi=R,G,B} \bar{I}^{\chi}}. \quad (7)$$

Here, χ indexes the chromatic channel, N is the number of pixels in the image region, and $\bar{I}^{\chi} = (1/N) \sum_{x,y} I^{\chi}(x, y)$ is the mean region intensity at each channel. Using this definition, we can refer to the contrast (visibility) in the raw image $C(I^{\text{min}})$ and in the recovery $C(\hat{L}_{\text{object}})$ and, thus, quantitatively assess the improvement. An example of an experimental implementation [17] of the

1. Underwater it is referred to as veiling light, path radiance [4], spacelight, or backscatter [16].

2. Motivated by studies which show that blur by scattering is *not* the dominating source of visibility degradation [4], [16], we do not address this blur at this stage.

3. Analysis of sequences of polarization filtered images has proved useful in various other computer vision problems [2], [9], [18], [22].



Fig. 1. Top: An image of a hazy scene, corresponding to the best state of mounted polarizer. Bottom: Dehazed image, using the basic algorithms of [17]. While contrast is significantly improved, noise is amplified in pixels corresponding to long distances.

method (4)-(6) is shown⁴ in Fig. 1. Subjectively, the visibility is indeed better. Moreover, in the area surrounding the front row of buildings, C improved by a factor of ≈ 2 from 0.23 to 0.47. The effect is much stronger in the distant ridge (on the upper left part of the image), which is hardly seen in the raw image but is revealed by dehazing. There, C changes by a factor of ≈ 11 from 0.018 to 0.19.

Despite the effectiveness of such inversion methods regarding contrast recovery, they are rather noise sensitive. Since (4) inverts the exponential attenuation (3), it drastically amplifies even a negligible acquisition noise in pixels corresponding to distant objects. Quantitatively, for white irradiance-independent image noise of standard-deviation (STD) σ , the STD of the recovered image [16] is

$$\sigma_{\hat{L}_{\text{object}}} \sim \sigma / \hat{t}. \quad (8)$$

Since $1/\hat{t} \in [1, \infty)$, noise amplification is unbounded. Hence, \hat{t} significantly effects⁵ noise amplification. Indeed, as seen in Fig. 1, the dehazed image is very noisy in pixels corresponding to distant objects.

4. For display, the images shown have undergone the *same* standard, automatic contrast stretching (see also [16], [17]). This was done only toward the display. The recovery algorithms were run on raw, unstretched data.

5. A small systematic error in \hat{t} may occur due to slight deviations from the airlight model or an inaccurate distance prior [14]. According to (4), this creates a systematic bias in \hat{L}_{object} , i.e., a time invariant amplification. The method presented here may yield this bias, as in other methods [10], [14], [16], [17]: It may exist in any method compensating for t , whether spatial regularization is attempted or not.

The above paragraph describes how $\sigma_{\hat{L}_{\text{object}}}$ is affected by the spatially varying t . Yet, how does noise in \hat{t} itself vary? Consider first the noise variance of \hat{A} ,

$$\sigma_{\hat{A}}^2 = \left(\frac{\partial \hat{A}}{\partial I_{\text{max}}} \right)^2 \sigma_{I_{\text{max}}}^2 + \left(\frac{\partial \hat{A}}{\partial I_{\text{min}}} \right)^2 \sigma_{I_{\text{min}}}^2. \quad (9)$$

Using (6),

$$\sigma_{\hat{A}}^2 = [\sigma_{I_{\text{max}}}^2 + \sigma_{I_{\text{min}}}^2] / p^2. \quad (10)$$

Hence, the noise in the estimated airlight is similar to that of the raw frames, up to a global scale. It is not amplified in a spatially varying way. A similar analysis leads to the same conclusion regarding the estimates \hat{S} and \hat{t} . For this reason, standard processing, e.g., space invariant smoothing, is sufficient to suppress noise of these variables.

3 RESTORATION USING ADAPTIVE REGULARIZATION

Noise is suppressed by regularizing the problem. Here, the estimate \hat{L}_{object} should minimize a cost function φ , which is typically composed of fitting and regularization terms

$$\varphi(L_{\text{object}}) = (\text{Fitting} + \text{Regularization}). \quad (11)$$

Consider first regularization written in a standard form

$$\text{Regularization} = \lambda \|DL_{\text{object}}\|^2, \quad (12)$$

where \mathcal{D} is the 2D Laplacian operator and λ is the weight of the regularization term. This term penalizes the cost function for unsmooth results. This enforces a smooth solution and may indeed reduce noise at pixels corresponding to distant objects. The drawback of such a standard approach is degradation of resolution of objects close to the camera by the blur induced by this digital processing. This blur is unjustified since only a little noise exists in pixels corresponding to close objects. It becomes obvious that the regularization should adapt to the medium transmittance at each pixel.

We wish to avoid any regularization in pixels corresponding to close pixels, for which $t \rightarrow 1$. To impose this adaptivity, we use a local weight in the regularization formulation that depends explicitly on t . Consider a term of the form

$$\text{Regularization} = \lambda \|\mathcal{WDL}_{\text{object}}\|^2. \quad (13)$$

Here, \mathcal{W} is a weighting operator. It depends explicitly on the transmittance t at each pixel, hence, it is implicitly adaptive to the object distance z .

To perform the minimization, the images can be converted to column-stack vectors: \mathbf{s} denotes the signal and \mathbf{l} denotes the object radiance. Then, the Laplacian operator \mathcal{D} is expressed as a matrix operation over a vector. Regularization weighting is achieved by a diagonal matrix \mathbf{W} . We use a simple matrix \mathbf{W} , whose elements are defined as

$$w_{ii} = (1 - t_{ii})^2, \quad (14)$$

and t_{ii} expresses \hat{t} at each corresponding pixel i . Note that $t_{ii} \in [0, 1]$. The proposed weighting emphasizes the regularization of pixels corresponding to distant objects (where $t_{ii} \rightarrow 0$) and turns off the regularization at close objects (where $t_{ii} \rightarrow 1$).

3.1 Color Formulation

In the context of *color*, define the column stack representation of the object radiance in the red, green, and blue channels as \mathbf{l}_R , \mathbf{l}_G , and \mathbf{l}_B , respectively. These vectors are linearized representations of light energy, i.e., compensating for a possible radiometric nonlinearity of the camera. Penalizing for deviation from the model described in Section 2, we use

$$\text{Fitting} = \|\mathbf{s}_R - \mathbf{T}_R \mathbf{l}_R\|^2 + \|\mathbf{s}_G - \mathbf{T}_G \mathbf{l}_G\|^2 + \|\mathbf{s}_B - \mathbf{T}_B \mathbf{l}_B\|^2, \quad (15)$$

where \mathbf{s}_R , \mathbf{s}_G , and \mathbf{s}_B are the signals at each color channel. The matrices \mathbf{T}_R , \mathbf{T}_G , and \mathbf{T}_B are diagonal, whose diagonal elements express \hat{t} at each corresponding pixel and color.

For regularization, we process the images in a color space which takes advantage of perceptual properties of color spatial resolution, such as the NTSC YIQ color space. There, the effective spatial resolution [5] of the chromatic content can be significantly lowered relative to the raw image resolution with very little perceived distortion. This enables aggressive denoising by spatial blur. The NTSC conversion, as applied to the linearized RGB values is given by

$$\begin{bmatrix} \mathbf{l}_Y \\ \mathbf{l}_I \\ \mathbf{l}_Q \end{bmatrix} = \begin{bmatrix} \alpha_Y & \beta_Y & \gamma_Y \\ \alpha_I & \beta_I & \gamma_I \\ \alpha_Q & \beta_Q & \gamma_Q \end{bmatrix} \begin{bmatrix} \mathbf{l}_R \\ \mathbf{l}_G \\ \mathbf{l}_B \end{bmatrix}, \quad (16)$$

where \mathbf{l}_Y , \mathbf{l}_I , and \mathbf{l}_Q are the values of each of the components in the YIQ space. Details of the elements of the matrix used in (16) are given in [5]. This representation expresses the luminance in the Y component separately from the chromatic components I and Q.

It is important to note that we apply this transformation *only* in the regularization term. We do not apply it to the fitting term. The reason is that fitting is derived from the original inverse problem (4), which is wavelength dependent through all its parameters: The image irradiance I , the airlight A , and the transmittance t . The relationship between these variables is physical, rather than perceptual, and is

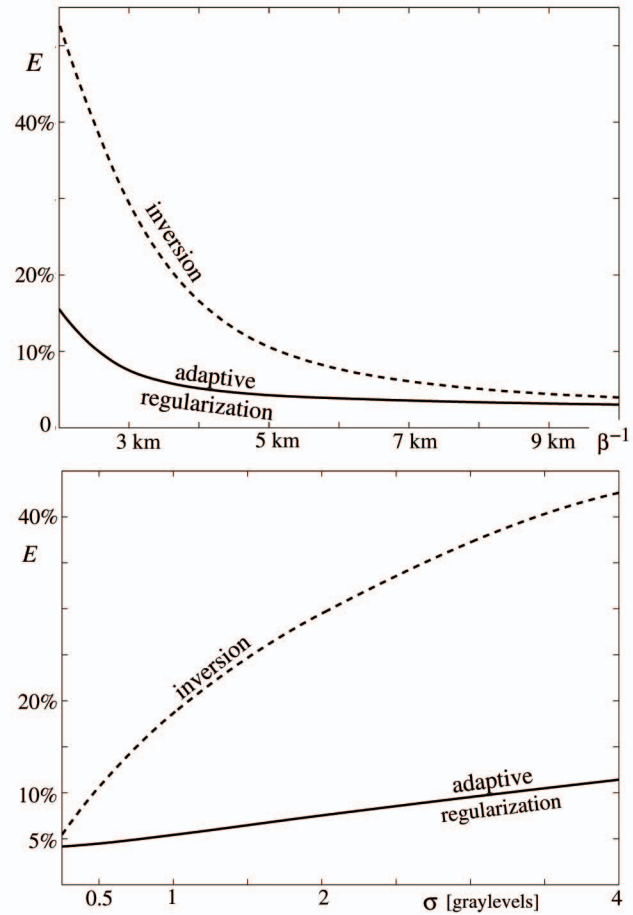


Fig. 2. Top: A simulated hazy and somewhat noisy image of a scene having a spatially varying distance map. Next to it is a regularized dehazing result. Middle: The normalized RMS error of the recovered image, as a function of the attenuation distance, for $\sigma = 2$ graylevels. Bottom: The error as a function of the raw image noise, for $\beta = 1/3 \lfloor \frac{1}{k_{\text{min}}} \rfloor$.

wavelength dependent. Using the YIQ color space in the fitting term would lose the physical relevance of the wavelength dependence. Thus, fitting is maintained in the RGB color space.

While we maintain the fitting term as in (15), we use the following regularization term:

$$\text{Regularization} = \lambda_Y \|\mathbf{WDL}_Y\|^2 + \lambda_C \|\mathbf{WDL}_I\|^2 + \lambda_C \|\mathbf{WDL}_Q\|^2, \quad (17)$$

where λ_Y is a weight of the luminance regularization and λ_C is a weight of the chrominance regularization. To make the image blur associated with regularization hardly apparent, we use $\lambda_C > \lambda_Y$: Selecting a relatively large λ_C suppresses noise in image components for which human vision tolerates blur, while keeping a small λ_Y minimizes blur in the component that is perceptually dominant. In (17), we used the weighting matrix \mathbf{W} as defined in (14). The weights are determined by the transmittance of the green



Fig. 3. The image corresponding to Fig. 1, recovered using adaptive regularization.

color channel in haze and the blue channel in water. These channels are used since they typically contain the highest light energies in the respective media, hence presumably yielding better distance maps for use in the weighting.

The bottom line is that we minimize (11) based on (15) and (17). This weighted regularization creates no blur at all in regions corresponding to close objects. There, the result is indistinguishable from that of simple inversion. On the other hand, the weighted regularization yields results that are less noisy than simple inversion at regions corresponding to long distances. This noise reduction is traded for image blur in those regions. This trade-off is controlled by the parameters λ_Y and λ_C . Thanks to the use of YIQ regularization, this blur is not severe, as real experiments show in Section 4.

3.2 Gradient and Hessian

To minimize (11), we apply standard optimization tools. Such optimization benefits from knowledge of the gradient and the Hessian of φ . We now derive these functions. Let us encapsulate the vectors \mathbf{l}_R , \mathbf{l}_G , and \mathbf{l}_B in a single column stack

$$\mathbf{p} = \begin{bmatrix} \mathbf{l}_R \\ \mathbf{l}_G \\ \mathbf{l}_B \end{bmatrix}. \quad (18)$$

Similarly, we may encapsulate the vectors \mathbf{s}_R , \mathbf{s}_G , and \mathbf{s}_B in a single column stack,

$$\mathbf{s}_{\text{RGB}} = \begin{bmatrix} \mathbf{s}_R \\ \mathbf{s}_G \\ \mathbf{s}_B \end{bmatrix}. \quad (19)$$

Define a matrix

$$\mathbf{T}_{\text{RGB}} = \begin{bmatrix} \mathbf{T}_R & \dots & 0 \\ \vdots & \mathbf{T}_G & \vdots \\ 0 & \dots & \mathbf{T}_B \end{bmatrix}. \quad (20)$$

The cost function defined by (15) and (17) is

$$\varphi(\mathbf{p}) = \|\mathbf{s}_{\text{RGB}} - \mathbf{T}_{\text{RGB}}\mathbf{p}\|^2 + \lambda_Y \|\mathbf{W}\mathcal{D}\Phi_Y\mathbf{p}\|^2 + \lambda_C \|\mathbf{W}\mathcal{D}\Phi_I\mathbf{p}\|^2 + \lambda_C \|\mathbf{W}\mathcal{D}\Phi_Q\mathbf{p}\|^2, \quad (21)$$

where the matrices Φ_Y , Φ_I , and Φ_Q are defined as

$$\Phi_Y = [\alpha_Y \mathbf{I} \ \beta_Y \mathbf{I} \ \gamma_Y \mathbf{I}], \quad \Phi_I = [\alpha_I \mathbf{I} \ \beta_I \mathbf{I} \ \gamma_I \mathbf{I}], \quad \Phi_Q = [\alpha_Q \mathbf{I} \ \beta_Q \mathbf{I} \ \gamma_Q \mathbf{I}], \quad (22)$$

while \mathbf{I} is the identity matrix. The α , β , and γ coefficients used in (22) are the ones used in (16), as defined in [5].

The gradient and Hessian of the fitting term are given by

$$\mathbf{g}_{\text{Fitting}} = 2\mathbf{T}_{\text{RGB}}^T (\mathbf{T}_{\text{RGB}}\mathbf{p} - \mathbf{s}_{\text{RGB}}) \quad \text{and} \quad \mathbf{H}_{\text{Fitting}} = 2\mathbf{T}_{\text{RGB}}^T \mathbf{T}_{\text{RGB}}, \quad (23)$$

respectively, where T denotes transposition. The gradient and Hessian of the regularization term are respectively given by

$$\mathbf{g}_{\text{Regularization}} = 2\lambda_Y \Phi_Y^T \mathcal{D}\mathbf{W}^T \mathbf{W}\mathcal{D}\Phi_Y\mathbf{p} + 2\lambda_C \Phi_I^T \mathcal{D}\mathbf{W}^T \mathbf{W}\mathcal{D}\Phi_I\mathbf{p} + 2\lambda_C \Phi_Q^T \mathcal{D}\mathbf{W}^T \mathbf{W}\mathcal{D}\Phi_Q\mathbf{p} \quad (24)$$

and

$$\mathbf{H}_{\text{Regularization}} = 2\lambda_Y \Phi_Y^T \mathcal{D}\mathbf{W}^T \mathbf{W}\mathcal{D}\Phi_Y + 2\lambda_C \Phi_I^T \mathcal{D}\mathbf{W}^T \mathbf{W}\mathcal{D}\Phi_I + 2\lambda_C \Phi_Q^T \mathcal{D}\mathbf{W}^T \mathbf{W}\mathcal{D}\Phi_Q. \quad (25)$$

Hence, the gradient of the cost function is $\mathbf{g}_{\varphi(\mathbf{p})} = \mathbf{g}_{\text{Fitting}} + \mathbf{g}_{\text{Regularization}}$, while its Hessian is $\mathbf{H}_{\varphi(\mathbf{p})} = \mathbf{H}_{\text{Fitting}} + \mathbf{H}_{\text{Regularization}}$. The seamless use of RGB fitting and YIQ regularization in the optimization stems from the linear relation (16) between these color spaces.

4 RESULTS

First, we describe simulations. They were based on hazy versions of the Lena picture, which was used to texture-map a depth-map made of a slanted plane, having a range of $1m - 10km$. The airlight has a slightly blue hue. The simulations varied the panchromatic attenuation coefficient β and the STD σ of noise added to the simulated hazy images, testing performance in various levels of noise and turbidity. The maximum pixel value was 255. A typical simulated frame is shown in Fig. 2. We then performed recovery. To avoid image tweaking, in all simulations and the subsequent real experiments, the *same* cost-function parameters were used: $\lambda_Y = 1/20$ and $\lambda_C = 1/2$. The simulated results are quantitatively compared to the ground-truth, Lena, using root-mean-squared (RMS) error, normalized by the mean intensity value

$$E = \frac{\sqrt{\frac{1}{3N} \sum_{\chi=R,G,B} \sum_{x,y} [\hat{L}_{\text{object}}^{\chi}(x,y) - L_{\text{object}}^{\chi}(x,y)]^2}}{\frac{1}{3N} \sum_{\chi=R,G,B} \sum_{x,y} L_{\text{object}}^{\chi}(x,y)}. \quad (26)$$

Here, N is the number of image pixels. Fig. 2 plots E as a function of the attenuation distance ($1/\beta$). In poor visibility (low attenuation distance), the output error is very high when using simple inversion and is much reduced by the regularization. Also in good visibility (high attenuation distance), regularization reduces



Fig. 4. (a) Raw frame. (b) Simple inversion. (c) Adaptive regularization attenuates noise at the background, while sparing consequent blur from foreground objects.

the recovery error. A similar conclusion is drawn from the plot (in Fig. 2) of E as a function of the raw image noise.

A result of a real experiment is shown in Fig. 3 (acquisition is as described in [17]). Compare it to the inversion result and the raw image in Fig. 1. The visibility has improved significantly relative to the raw image, while the regularization results in less noise than that of simple inversion. Close objects are not blurred at all. Quantitatively, the contrast is $C = 0.49$ in the regularized area surrounding the front row of buildings. This value is very similar to the one yielded by simple inversion. In the distant ridge (on the upper left part of the image), the regularized result yields $C \approx 0.07$, which is $\approx 4.3 \times C(I^{\min})$. Yet, this value is more moderate than the output of simple inversion since the associated denoising blur somewhat decreases the image variance.

An additional set of experimental results is shown in Fig. 4. It zooms in on part of an outdoor scene in haze to better observe the adaptive suppression of noise. Here, the raw frames were acquired using a Nikon D100 camera. This camera was also used underwater (see [16]) in a watertight housing. The underwater results are white balanced prior to display to counter the strong blue hue of the illumination in this environment. The raw image and the image recovered by adaptive regularization are displayed in Fig. 5. The process improves visibility significantly. For example, in the rectangular areas marked in Fig. 5, which correspond to distant scene parts, we measure $C(\hat{L}_{\text{object}}) \approx 2.6 \times C(I^{\min})$.

5 DISCUSSION

The approach was demonstrated here in natural settings of haze and water. Nevertheless, we believe that it may apply to other media and scenarios, where noise is affected by object distance. It

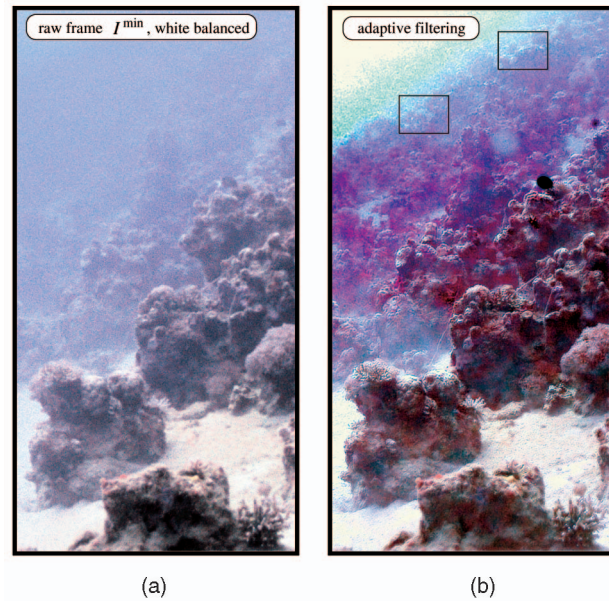


Fig. 5. (a) A white balanced raw frame acquired underwater. (b) Reconstruction with adaptive regularization.

may be worth adapting this approach to recovery based on images taken over long periods, in scenes of known distances [12]. An important extension to this work is to explore operators other than the ones we used, for weighting and for penalizing unsmooth results. In particular, it may be interesting to generalize median filtering to this spatially varying distance-dependent problem. Another extension is to account for random noise and bias in \hat{t} while estimating L_{object} , regularizing them jointly, e.g., using [15].

ACKNOWLEDGMENTS

The authors thank Einav Namer for image acquisition and Michael Elad for useful discussions. 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 the Magnetron program of the Israeli Ministry of Commerce. The work was done in a joint project with ELOp Ltd, and conducted in the Ollendorff Minerva Center in the Electrical Engineering Department at the Technion. Minerva is funded through the BMBF.

REFERENCES

- [1] F. Cozman and E. Kroktov, "Depth from Scattering," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, pp. 801-806, 1997.
- [2] H. Farid and E.H. Adelson, "Separating Reflections from Images by Use of Independent Component Analysis," *J. Optical Soc. Am. A*, vol. 16, pp. 2136-2145, 1999.
- [3] E. Hecht, *Optics*, pp. 562-628. fourth ed., Addison Wesley, 2002.
- [4] W.S. Jagger and W.R.A. Muntz, "Aquatic Vision and the Modulation Transfer Properties of Unlighted and Diffusely Lighted Natural Waters," *Vision Research*, vol. 33, pp. 1755-1763, 1993.
- [5] A.K. Jain, *Fundamentals of Digital Image Processing*, pp. 67-71, 553-555. Prentice Hall, 1989.
- [6] D.M. Kocak and F.M. Caimi, "The Current Art of Underwater Imaging with a Glimpses of the Past," *MTS J.*, vol. 39, pp. 5-26, 2005.
- [7] N.S. Kopeika, *A System Engineering Approach to Imaging*, pp. 446-452. SPIE Press, 1998.
- [8] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas, "Synthetic Aperture Confocal Imaging," *ACM Trans. Graphics*, vol. 23, pp. 825-34, 2004.
- [9] D. Miyazaki, M. Kagesawa, and K. Ikeuchi, "Transparent Surface Modeling from a Pair of Polarization Images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, pp. 73-82, 2004.
- [10] S.G. Narasimhan and S.K. Nayar, "Vision and the Atmosphere," *Int'l J. Computer Vision*, vol. 48, pp. 233-254, 2002.

- [11] S.G. Narasimhan and S.K. Nayar, "Interactive Deweathering of an Image Using Physical Models," *Proc. IEEE Workshop Color & Photometric Methods in Computer Vision*, 2003.
- [12] S.G. Narasimhan, C. Wang, and S.K. Nayar, "All the Images of an Outdoor Scene," *Proc. European Conf. Computer Vision*, pp. 148-162, 2002.
- [13] S. Negahdaripour and H. Madjidi, "Stereovision Imaging on Submersible Platforms for 3D Mapping of Benthic Habitats and Sea Floor Structures," *IEEE J. Oceanic Eng.*, vol. 28, pp. 625-650, 2003.
- [14] J.P. Oakley and B.L. Satherley, "Improving Image Quality in Poor Visibility Conditions Using a Physical Model for Contrast Degradation," *IEEE Trans. Image Processing*, vol. 7, pp. 167-179, 1998.
- [15] J. Polzehl and V.G. Spokoiny, "Adaptive Weights Smoothing with Applications to Image Restoration," *J. Royal Statistical Soc., Series B*, vol. 62, pp. 335-354, 2000.
- [16] Y.Y. Schechner and N. Karpel, "Recovery of Underwater Visibility and Structure by Polarization Analysis," *IEEE J. Oceanic Eng.*, vol. 30, pp. 570-587, 2005.
- [17] Y.Y. Schechner, S.G. Narasimhan, and S.K. Nayar, "Polarization-Based Vision Through Haze," *Applied Optics*, vol. 42, pp. 511-525, 2003.
- [18] Y.Y. Schechner, J. Shamir, and N. Kiryati, "Polarization and Statistical Analysis of Scenes Containing a Semi-Reflector," *J. Optical Soc. Am. A*, vol. 17, pp. 276-284, 2000.
- [19] S. Shwartz, E. Namer, and Y.Y. Schechner, "Blind Haze Separation," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 1984-1991, 2006.
- [20] J. Sun, J. Jia, C.K. Tang, and H.Y. Shum, "Poisson Matting," *ACM Trans. Graphics*, vol. 23, pp. 315-321, 2004.
- [21] T. Treibitz and Y.Y. Schechner, "Instant 3Descatter," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 1861-1868, 2006.
- [22] L.B. Wolff, "Polarization-Based Material Classification from Specular Reflection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, pp. 1059-1071, 1999.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.