

Variational Distance-Dependent Image Restoration

Ran Kaftory, Yoav Y. Schechner and Yehoshua Y. Zeevi
Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa, Israel, 32000

kaftoryr@tx.technion.ac.il, yoav@ee.technion.ac.il, zeevi@ee.technion.ac.il

Abstract

There is a need to restore color images that suffer from distance-dependent degradation during acquisition. This occurs, for example, when imaging through scattering media. There, signal attenuation worsens with the distance of an object from the camera. A 'naive' restoration may attempt to restore the image by amplifying the signal in each pixel according to the distance of its corresponding object. This, however, would amplify the noise in a non-uniform manner. Moreover, standard space-invariant denoising over-blurs close by objects (which have low noise), or insufficiently smooths distant objects (which are very noisy). We present a variational method to overcome this problem. It uses a regularization operator which is distance dependent, in addition to being edge-preserving and color-channel coupled. Minimizing this functional results in a scheme of reconstruction-while-denoising. It preserves important features, such as the texture of close by objects and edges of distant ones. A restoration algorithm is presented for reconstructing color images taken through haze. The algorithm also restores the path radiance, which is equivalent to the distance map. We demonstrate the approach experimentally.

1. Introduction

Attention is increasingly being directed towards imaging through scattering and attenuating media, such as haze [5, 13, 15, 18, 23], water [9, 12, 14, 16, 21, 22, 27] and tissue [6, 7]. It is important to overcome image degradation problems associated with such media. In the open air, this is needed for human activities such as driving and flight [18], as well as for computer vision algorithms such as recognition, tracking, and remote sensing. Underwater imaging is used for research in marine biology, archaeology and mapping, as well as for various industrialized applications [2, 17, 26]. Needless to say, imaging through tissue is

important for physicians.

While we mention very different imaging media and modalities, they do share common challenging aspects. In such media, the signal from a scene point is significantly altered due to scattering and absorption. One of the main effects is the decay of this signal with distance. Since objects in the field of view are at different distances from the imager, image degradation is spatially varying. Several methods have been proposed to restore image taken through scattering and attenuating media [9, 12, 14, 15, 16, 27], some of which account for part of the spatially varying degradation effects.

Although significant progress has been achieved, a major problem that still needs to be addressed is noise. When attempting to amplify attenuated signals, noise is also amplified. Since the degradation of the image is spatially varying, noise amplification is spatially varying as well. Therefore, standard denoising of the image might over-blur close by objects (which have low noise), or insufficiently smooth distant objects (which are very noisy). As a result, noise suppression should be adaptive to the distance map of the objects. This distance map may be estimated using *prior* knowledge, range finders or from the image itself. Since the *estimation of the distance map might be noisy*, caution should be taken in the adaptive denoising process. In the case of color images, special care should be taken to preserve color fidelity and edges.

In this paper, recovery is posed as an optimization problem, in which *both* the restored image and the path radiance (which is a function of the distance map), are regularized. The regularization operator is spatially weighted in order to achieve dependency on object distances and edges. The regularization operator itself, when used on color images, exhibits special coupling between the color channels and preserves the gradient directions. Restoration is achieved by minimizing a functional. It results in a reconstruction-while-denoising scheme.

We demonstrate our method by reconstructing images taken through haze. They are based on frames taken us-

ing a polarizer. The restoration process preserves important features, such as texture of close by objects, edges of distant ones and gradient directions of the color channels.

2. Distance-Dependent Degradation

A major way in which object distance affects image degradation stems from the falloff of radiation reaching the object at the distance, or the object radiance falloff due to a disturbing medium. Here we describe a few examples. Consider the use of an active illuminator, such as flash in photography. Then, the object irradiance is approximately inversely proportional to r^2 , where r is the distance of the illuminator from the object. This relation applies to photography in an open, clear environment. However, if the scene is embedded in an attenuating medium, the problem is exacerbated. The medium between the illuminator and the object has a transmissivity of $T^{\text{lighting}} \approx \exp(-cr)$, where c is the medium's extinction coefficient [5]. This effect, as the $1/r^2$ falloff mentioned above, affects the object irradiance, hence the radiance.

Even if the object irradiance is uniform, the object radiance may fall off with the distance, if the object is viewed through an attenuating medium. The radiation emanated from the object towards the camera is attenuated by $T = \exp(-cz)$, where z is the distance between the object and the imager. Any of these effects reduces the wanted signal as the object becomes more distant. Thus, reconstruction of objects is liable to be more noisy as distances increase. Let R be the ideal object radiance, as measured if no distance-dependent disturbances existed. Due to a distance-dependent falloff characterized by T , the detected object radiance is $D = RT$. Generally, the distance varies with the image coordinates (x, y) . Furthermore, the medium properties as well as the object radiance depend on the wavelength band (or color), indexed by λ . Hence, we write

$$D_\lambda(x, y) = R_\lambda(x, y)T_\lambda(x, y). \quad (1)$$

If the medium is scattering, it induces an additional disturbance. Some radiation *unrelated* to the object is scattered towards the camera. Its origin is the surrounding lighting, part of which is intercepted by particles in the medium and scattered. This creates an additive disturbance $A_\lambda(x, y)$, which is termed *airlight* [15] or path radiance when imaging in haze. A similar effect exists in other media, e.g., underwater, where it is called *backscatter* [9, 21, 22, 27]. Then, the overall sensed radiance is

$$I_\lambda(x, y) = R_\lambda(x, y)T_\lambda(x, y) + A_\lambda(x, y). \quad (2)$$

The path radiance and T are interdependent. In case of vision under natural lighting [15], $A_\lambda = a_\lambda^\infty(1 - T_\lambda)$, where a_λ^∞ is an environmental parameter, expressing the path radiance when $T_\lambda \rightarrow 0$, i.e., when the object is very distant.

This parameter can be extracted based on the raw image data, and can be assumed to be constant across the field of view. Note that $0 \leq A_\lambda(x, y) \leq a_\lambda^\infty$, since $1 \geq T_\lambda \geq 0$. This stems from the fact that $T_\lambda = \exp(-c_\lambda z)$ and $z > 0$. This relation between path radiance and T_λ yields a useful estimation [23] of T_λ :

$$\hat{T}_\lambda(x, y) = 1 - \frac{A_\lambda(x, y)}{a_\lambda^\infty}. \quad (3)$$

Coupled Unknowns

Eq. (3) means that if we have an estimate of $A_\lambda(x, y)$, then we can derive T_λ at each pixel. This facilitates inversion of Eq. (1), hence restoring the object radiance $R_\lambda(x, y)$. In vision in haze or underwater [21, 22, 23, 24], an estimate of $A_\lambda(x, y)$ can be obtained by mounting a polarizing filter on the camera. Then, two images are taken with the filter, in each of which the polarizer is oriented in a different direction. One image is dimmer than the other and is termed I_λ^{min} . The brighter image is termed I_λ^{max} . The physical model of these images is¹

$$I_\lambda^{\text{min}}(x, y) = R_\lambda(x, y)T_\lambda(x, y) + A_\lambda(x, y)(1 - p_\lambda), \quad (4)$$

$$I_\lambda^{\text{max}}(x, y) = R_\lambda(x, y)T_\lambda(x, y) + A_\lambda(x, y)(1 + p_\lambda). \quad (5)$$

Here $p_\lambda \in [0, 1]$ is an environmental parameter, which can be automatically extracted [24] from these raw images.² The unpolarized image is defined as the sum of the polarized images:

$$I_\lambda = I_\lambda^{\text{min}} + I_\lambda^{\text{max}}. \quad (6)$$

For illustration, such image data is simulated in Fig. 1. The increase of degradation with distance is evident.

Eqs. (3,4,5) provide sufficient constraints to make the joint estimation of $R_\lambda(x, y)$, $A_\lambda(x, y)$ and $T_\lambda(x, y)$ well posed, though it may be noisy. However, such a joint estimation introduces *coupling* between these unknown. As we show later in this paper, this means that regularization should be jointly applied to all the unknown fields.

3. Noisy Restoration

Let us add noise $n_\lambda(x, y)$ to a sensed image expressed by Eq. (2). Then, suppose for the moment that both $T_\lambda(x, y)$ and $A_\lambda(x, y)$ are perfectly known, i.e., these variables are noiseless. In this case, a direct inversion of Eq. (2) yields

$$\begin{aligned} \hat{R}_\lambda(x, y) &= \frac{I_\lambda(x, y) - A_\lambda(x, y)}{T_\lambda(x, y)} \\ &= R_\lambda(x, y) + \frac{n_\lambda(x, y)}{T_\lambda(x, y)}. \end{aligned} \quad (7)$$

¹The filter itself attenuates the image globally, but this scalar multiplication is disregarded.

²It expresses the degree of polarization of the path radiance.

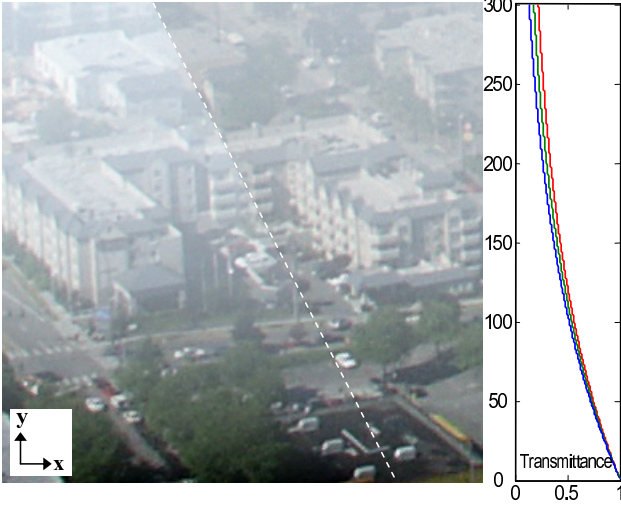


Figure 1. Simulated imaging through haze using a polarizer. There is random image noise (PSNR=40 dB). [Left] I^{\max} . [Right] I^{\min} . [Plot] The transmittance, T , as a function of y . It is plotted for the red, blue and green color channels with respective plot colors. The haze becomes dense where the transmittance is low.

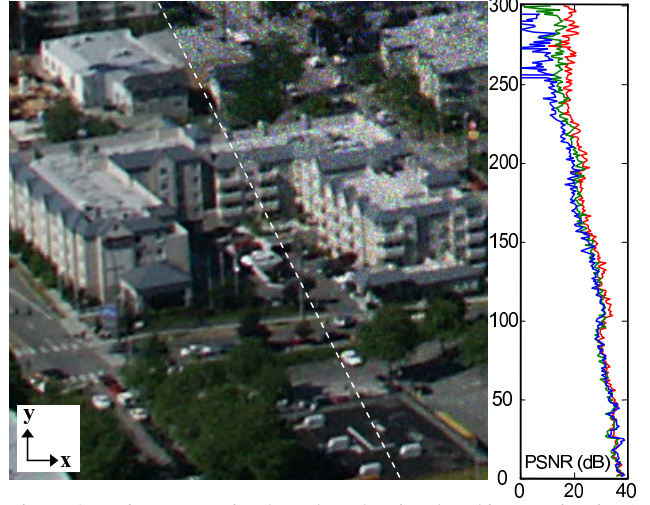


Figure 2. Noisy restoration based on the simulated images in Fig. 1 using Eqs. (7,8) and (9). [Left] Original ground truth image. [Right] Restored image. The restoration is noisy in upper rows, which correspond to low transmittance (see the plot in Fig. 1). [Plot] The PSNR of the reconstructed image as a function of y , for each color channel.

The noise is amplified since generally $T_\lambda(x, y) < 1$. Moreover, noise amplification is *spatially varying* through its distance dependency. Different pixels generally correspond to different object distances z , as described in Sec. 2. This affects $1/T_\lambda(x, y)$, hence the noise amplification. At very distant objects, where $T_\lambda(x, y) \rightarrow 0$, the restoration becomes highly ill-conditioned.

The problem of noise is exacerbated by the fact that $T_\lambda(x, y)$ itself is typically unknown, and should be estimated somehow. Such an estimate is noisy as well, and this eventually increases the error and noise in the estimated radiance $\hat{R}_\lambda(x, y)$. We can see this specifically when looking at polarization-based vision through haze, modelled in Sec. 2. There, an unregularized estimation of the path radiance is obtained from Eqs. (4,5) as

$$\begin{aligned} \hat{A}_\lambda(x, y) &= \frac{I_\lambda^{\max}(x, y) - I_\lambda^{\min}(x, y)}{p_\lambda} \\ &= A_\lambda(x, y) + n_\lambda^{\text{path}}(x, y). \end{aligned} \quad (8)$$

Here $n_\lambda^{\text{path}}(x, y)$ is the noise in the path radiance estimate. Its variance is $\sigma_{\text{path}}^2 = 2(\sigma/p_\lambda)^2$, where σ^2 is the noise variance in each of the acquired images I^{\min} and I^{\max} . Here, the raw images have independent noise. Based on Eq. (3), the unregularized estimation of T is

$$\begin{aligned} \hat{T}_\lambda(x, y) &= 1 - \hat{A}_\lambda(x, y)/a_\lambda^\infty \\ &= T_\lambda(x, y) + [n_\lambda^{\text{path}}(x, y)/a_\lambda^\infty]. \end{aligned} \quad (9)$$

When this noisy estimate of T is used in Eq. (7), it affects $\hat{R}_\lambda(x, y)$.

An example for the effect of noise in an unregularized dehazed image is shown in Fig. 2. Clearly, the noise intensity varies spatially with the scene distance.

4. A Variational Approach

Noise can be suppressed by regularization [20] which can be based on a variational framework [8, 28]. This framework employs the action of partial differential equations (PDEs) on images, serving the function of smoothing operators. By a proper choice of these operators, this smoothing can filter out the noise and still be edge preserving. Our approach has several interesting aspects. The PDE formulation *adapts to the spatially varying medium transmittance* T (implicitly, to z), to counter the intrinsic problem of noise amplification in the restoration process. In addition, the formulation affects *coupled fields* which have different natures: $R_\lambda(x, y)$, $A_\lambda(x, y)$ and $T_\lambda(x, y)$. Finally, this formulation applies to color (or spectral) images, hence smoothing should be consistent in different channels λ . In the following, we describe this framework in the context of vision under hazy conditions.

Define $\{\hat{R}\} \equiv \{\hat{R}_\lambda\}_{\text{all } \lambda}$ as the set of estimated fields corresponding to the object radiance in all the wavelength bands. Specifically, for color images, $\{\hat{R}\} \equiv \{\hat{R}_\lambda\}_{r,g,b}$, where $\lambda = r, g, b$ denote the red, green and blue color channels, respectively. Similarly, define $\{\hat{A}\} \equiv \{\hat{A}_\lambda\}_{\text{all } \lambda}$ as the estimated path radiance distribution in these bands. These are the fields that we seek to optimize. The estimated transmittance field $\hat{T}_\lambda(x, y)$ is not explicitly included. Rather, it is derived directly from $\hat{A}_\lambda(x, y)$ based on Eq. (3). Con-

sider a cost functional over these fields

$$\mathbb{J} = \iint_{\Omega} \mathbb{F}(\{\hat{R}\}, \{\hat{A}\}) + \alpha \mathbb{R}(\{\hat{R}\}, \{\hat{A}\}) dx dy, \quad (10)$$

where Ω is the image domain. Here $\mathbb{F}(\{\hat{R}\}, \{\hat{A}\})$ is a *fidelity* term, which measures how the observed data fits the physical model given in Eqs. (3,4,5). The term $\mathbb{R}(\{\hat{R}\}, \{\hat{A}\})$ expresses regularization of the sought fields. We focus on this term in Sec. 4.2. The parameter α weights the regularization term relative to the fidelity term.

Minimizing \mathbb{J} is accomplished by reaching $\partial\mathbb{J}/\partial\hat{R}_\lambda = 0$ and $\partial\mathbb{J}/\partial\hat{A}_\lambda = 0$, or equivalently, solving its Euler-Lagrange equations

$$\frac{\partial\mathbb{F}(\{\hat{R}\}, \{\hat{A}\})}{\partial\hat{R}_\lambda} + \alpha \frac{\partial\mathbb{R}(\{\hat{R}\}, \{\hat{A}\})}{\partial\hat{R}_\lambda} = 0 \quad (11)$$

$$\frac{\partial\mathbb{F}(\{\hat{R}\}, \{\hat{A}\})}{\partial\hat{A}_\lambda} + \alpha \frac{\partial\mathbb{R}(\{\hat{R}\}, \{\hat{A}\})}{\partial\hat{A}_\lambda} = 0 \quad (12)$$

for every λ . A common practice of PDE-based methods is to apply the Euler-Lagrange equations in iterative steps. Similarly, here we apply Eqs. (11,12) iteratively, one field of \hat{A}_λ or \hat{R}_λ at a time.

4.1. The Fidelity Operator

The fidelity term is derived from the physical model of image formation, given in Eqs. (4) and (5). If we use the mean-squared error criterion, then³

$$\begin{aligned} \mathbb{F}(\{\hat{R}\}, \{\hat{A}\}) &= \\ &= \frac{1}{2} \sum_{\lambda} \left[I_{\lambda}^{\min} - \hat{R}_{\lambda}(1 - \hat{A}_{\lambda}/a_{\lambda}^{\infty}) - \hat{A}_{\lambda}(1 - p_{\lambda}) \right]^2 \\ &\quad + \frac{1}{2} \sum_{\lambda} \left[I_{\lambda}^{\max} - \hat{R}_{\lambda}(1 - \hat{A}_{\lambda}/a_{\lambda}^{\infty}) - \hat{A}_{\lambda}(1 + p_{\lambda}) \right]^2. \end{aligned} \quad (13)$$

The Euler-Lagrange equations require partial derivatives. These are

$$\begin{aligned} \frac{\partial\mathbb{F}(\{\hat{R}\}, \{\hat{A}\})}{\partial\hat{R}_\lambda} &= (1 - \hat{A}_\lambda/a_\lambda^\infty) \cdot \\ &\quad \cdot \left[2\hat{A}_\lambda - I_\lambda^{\min} - I_\lambda^{\max} + 2\hat{R}_\lambda(1 - \hat{A}_\lambda/a_\lambda^\infty) \right] \end{aligned} \quad (14)$$

and

$$\begin{aligned} \frac{\partial\mathbb{F}(\{\hat{R}\}, \{\hat{A}\})}{\partial\hat{A}_\lambda} &= p_\lambda(I_\lambda^{\min} - I_\lambda^{\max} + 2\hat{A}_\lambda p_\lambda) + \\ &+ \left(1 - \frac{\hat{R}_\lambda}{a_\lambda^\infty} \right) \left[2\hat{A}_\lambda - I_\lambda^{\min} - I_\lambda^{\max} + 2\hat{R}_\lambda \left(1 - \frac{\hat{A}_\lambda}{a_\lambda^\infty} \right) \right]. \end{aligned} \quad (15)$$

³For simpler notation, we omitted the (x, y) notation in these terms. Yet, the underlying (x, y) dependency is maintained.

4.2. The Regularization Operator

The regularization term is based on Beltrami-flow, but it is *adapted* to the spatially-varying nature of the noise. Ref. [25] interprets an image as a manifold (surface) embedded in a high dimensional space, where x, y are two coordinates of this space, and the intensity at each of the λ channels is represented as an additional dimension. Hence, a color image is a manifold embedded in a 5-dimensional space. From this viewpoint, image denoising can be interpreted as a process that minimizes the *surface area* of this manifold.

Here we deal with two major fields. One is the object radiance, which we wish to obtain per channel λ and location x, y . We denote its corresponding manifold as R . The other is the path radiance, having a corresponding manifold A . The surface area of R is measured by the Polyakov action [19], which⁴ is

$$\iint_{\Omega} \mathbb{S}(R) dx dy = \iint_{\Omega} \sqrt{\det(\mathbf{G}_R)} dx dy. \quad (16)$$

Here \mathbf{G}_R is a 2×2 matrix. Each of its elements depends on the spatial location x, y . The element values are given [11] by

$$\begin{aligned} g_{11}^R &= 1 + \beta^2 \sum_{\lambda} (\partial\hat{R}_\lambda/\partial x)^2 \\ g_{12}^R = g_{21}^R &= \beta^2 \sum_{\lambda} \frac{\partial\hat{R}_\lambda}{\partial x} \frac{\partial\hat{R}_\lambda}{\partial y} \\ g_{22}^R &= 1 + \beta^2 \sum_{\lambda} (\partial\hat{R}_\lambda/\partial y)^2, \end{aligned} \quad (17)$$

where β is a parameter that sets the scale of intensity dimensions. Following Ref. [10], the use of Eq. (16) in the regularization forces the color channels to spatially align. Specifically, consider places where an edge exist at multiple color channels at the same position x, y in the raw images. Then, its location will be maintained in $\{\hat{R}\}$, in all the relevant channels. This cross-channel coupling suppresses color distortions.

Minimization of Eq. (16) is aided by its partial derivatives.

$$\frac{\partial\mathbb{S}(R)}{\partial\hat{R}_\lambda} = \mathcal{D} \left[\sqrt{\det(\mathbf{G}_R)} (\mathbf{G}_R)^{-1} \nabla\hat{R}_\lambda \right], \quad (18)$$

where $\nabla\hat{R}_\lambda$ is the spatial gradient of R_λ and \mathcal{D} is the divergence operator.

Now, we would like to incorporate an adaptation to the medium transmittance. Hence, smoothing of R should not be performed in close by scene regions, but should increase

⁴This expression holds for Euclidian spaces.

as the physical attenuation in the medium becomes more significant. In addition, we would like to enhance the edge preservation, in places where smoothing does take place. We achieve this by using a *weighting* function

$$w(x, y) = w_{\text{transmittance}}(x, y)w_{\text{edge}}(x, y) \quad (19)$$

in the smoothing operator, as we describe next. Then the regularization derivative is set to be

$$\frac{\partial \mathbb{R}(\{\hat{R}\}, \{\hat{A}\})}{\partial \hat{R}_\lambda} = w(x, y) \frac{\partial \mathbb{S}(\mathbf{R})}{\partial \hat{R}_\lambda}. \quad (20)$$

Following Eq. (7), if no regularization is performed, the noise variance of \hat{R}_λ is

$$\sigma_R^2(x, y) \propto \frac{\sigma^2}{[T_\lambda(x, y)]^2} = \sigma^2 \left[\frac{a_\lambda^\infty}{a_\lambda^\infty - A_\lambda(x, y)} \right], \quad (21)$$

where Eq. (3) is used. Hence, we set

$$w_{\text{transmittance}}(x, y) = \left[\frac{\hat{A}_\lambda(x, y)}{a_\lambda^\infty} \right]^2, \quad (22)$$

where \hat{A}_λ is an estimate of the path radiance. Note that at image locations x, y where $A_\lambda(x, y) \rightarrow 0$, i.e., at close distances, $w_{\text{transmittance}}(x, y) = 0$. Hence, no regularization is performed in these locations, as desired. On the other hand, as $A_\lambda(x, y) \rightarrow a_\lambda^\infty$, i.e., at very large distances, $w_{\text{transmittance}}(x, y) = 1$, implying that regularization is fully employed there. Eq. (22) monotonically adapts to the estimated transmittance, between these two cases. A similar idea, of adapting a smoothing operator to a distance map by a weighting function was presented by [1] in the context of rendering effects of defocus.

To enhance the edge preservation in places where smoothing does take place (large object distances), we use

$$w_{\text{edge}}^{\mathbf{R}}(x, y) = [\det(\mathbf{G}_{\mathbf{R}})]^{-1/2}, \quad (23)$$

where $\mathbf{G}_{\mathbf{R}}$ depends on x, y as defined in Eq. (17). The term $\sqrt{\det(\mathbf{G})}$ is a multi-channel edge indicator. The use of Eq. (23) jointly with Eq. (20) (without the use of $w_{\text{transmittance}}$) is known as the *Beltrami* operator.⁵

Now, we address the estimated path radiance, \hat{A}_λ , having a manifold \mathbf{A} . We define $\mathbf{G}_{\mathbf{A}}$ in analogy to Eq. (17), where \hat{R}_λ is substituted by \hat{A}_λ . The surface area of this manifold is given similarly to Eq. (16), by

$$\iint_{\Omega} \mathbb{S}(\mathbf{A}) dx dy = \iint_{\Omega} \sqrt{\det(\mathbf{G}_{\mathbf{A}})} dx dy, \quad (24)$$

⁵This operator was also used for the processing of movies, color and textures in [11].

and its partial derivatives are

$$\frac{\partial \mathbb{S}(\mathbf{A})}{\partial \hat{A}_\lambda} = \mathcal{D} \left[\sqrt{\det(\mathbf{G}_{\mathbf{A}})} (\mathbf{G}_{\mathbf{A}})^{-1} \nabla \hat{A}_\lambda \right]. \quad (25)$$

The noise in \hat{A}_λ , is *not* amplified as a function of the transmittance, as can be seen in Eq. (8). Therefore, adaptation of the regularization to the transmittance is not needed. However, edge preservation is important, since it maintains distance discontinuities, such as mountain ridges. Similarly to \hat{R}_λ , the enhancement of edge preservation is obtained by using

$$w_{\text{edge}}^{\mathbf{A}}(x, y) = [\det(\mathbf{G}_{\mathbf{A}})]^{-1/2}, \quad (26)$$

and the regularization derivative is

$$\frac{\partial \mathbb{R}(\{\hat{R}\}, \{\hat{A}\})}{\partial \hat{A}_\lambda} = w_{\text{edge}}^{\mathbf{A}}(x, y) \frac{\partial \mathbb{S}(\mathbf{A})}{\partial \hat{A}_\lambda}. \quad (27)$$

A coupling between \hat{R}_λ and \hat{A}_λ is imposed by the Euler Lagrange equations. Particularly, optimizing \hat{R}_λ requires an estimate \hat{A}_λ . Therefore, we use an iterative scheme, in which optimization of \hat{R}_λ uses the estimate \hat{A}_λ from the previous iteration, and vice versa.

4.3. Minimization Scheme

We prefer to use an Alternating Minimization scheme as was used by [4]. It iteratively propagates the derivatives to solve the Euler-Lagrange equations:

1. Perform a noisy restoration (Sec. 2). Denote the results as $\hat{A}_\lambda^{(0)}$ and $\hat{R}_\lambda^{(0)}$. Let them initialize the following steps.
2. Use $\hat{A}_\lambda^{(t)}$ and $\hat{R}_\lambda^{(t)}$ obtained in the previous iteration t .
3. For every λ , solve $\hat{R}_\lambda^{(t+1)}$ by:

$$\frac{\partial \mathbb{F}(\{\hat{R}^{(t+1)}\}, \{\hat{A}^{(t)}\})}{\partial \hat{R}_\lambda^{(t+1)}} + \alpha w_{\text{edge}}^{\mathbf{R}} \frac{\partial \mathbb{S}(\{\hat{R}^{(t+1)}\})}{\partial \hat{R}_\lambda^{(t+1)}} = 0. \quad (28)$$

4. For every λ , solve $\hat{A}_\lambda^{(t+1)}$ by:

$$\frac{\partial \mathbb{F}(\{\hat{R}^{(t+1)}\}, \{\hat{A}^{(t+1)}\})}{\partial \hat{A}_\lambda^{(t+1)}} + \alpha w_{\text{edge}}^{\mathbf{A}} \frac{\partial \mathbb{S}(\{\hat{A}^{(t+1)}\})}{\partial \hat{A}_\lambda^{(t+1)}} = 0. \quad (29)$$

5. Repeat steps 2,3,4 until convergence.

Eqs. (28,29) are nonlinear PDEs. Rather than solving them directly, we use the Fixed Point Lagged Diffusive method, which was used in [29]. More details can be found in the appendix.

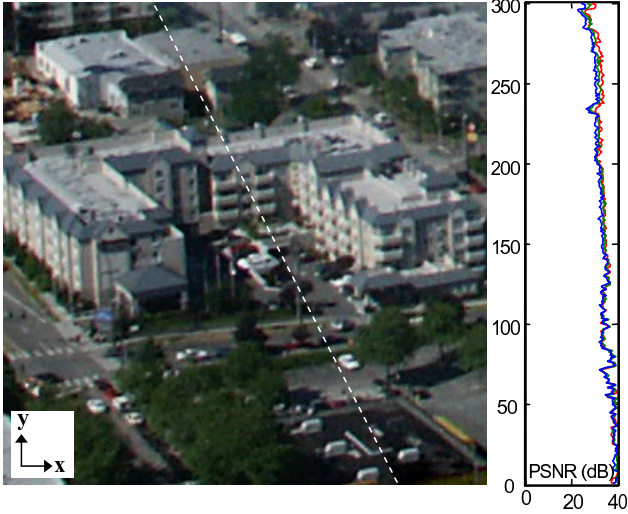


Figure 3. Variational-based restoration of the images simulated in Fig. 1. [Left] Original ground truth image. [Right] Restored image. The restoration has very low noise and the edges are sharp. [Plot] As a function of y , the PSNR of the reconstructed image is much higher than in Fig. 2.

5. Results

The minimization scheme presented in the previous section was applied to the simulated example of Fig. 1. The results are presented in Fig. 3. It can be seen that the reconstructed object radiance resembles the original and it has low noise. The edges are sharp and the colors are clear. The PSNR as a function of the image row number is almost constant and much better than that of the ‘naive’ reconstruction (Fig. 2).

In order to compare our method to a state-of-the-art denoising algorithm, the Non-Local Means algorithm [3] post-processed the ‘naive’ result (shown in Fig. 2).⁶ The postprocessing result is shown in Fig. 4. The restored image suffers from an artifact of blurred texture (e.g. the trees in the yellow bounding box). The PSNR is constant as a function of y but it is worse than that of our method.

Our method was tested on real images taken in a hazy day by a Nikon D100 having a linear radiometric response. A polarizer was mounted on the camera and two consecutive frames were taken with different polarization states. The captured images are depicted in Fig. 5. The ‘naive’ reconstruction is shown in Fig. 6. The reconstruction is noisy especially in distant pixels. Our variational-based reconstruction⁷ is shown in Fig. 7. It can be seen that the image has low noise. Important features, such as texture for the close fields and edges for the distant mountains are pre-

⁶Acknowledgements to G. Peyre for implementing the algorithm and making it available via the MATLAB File Exchange.

⁷The object radiance of pixels where $\hat{A}_\lambda \approx a_\lambda^\infty$ cannot be restored. Therefore, they were replaced by the noisy restoration of the path radiance.




Figure 4. Non-Local Means denoising of the ‘naive’ restoration of Fig. 2. [Left] Original ground truth image. [Right] Restored and denoised image. A major artifact: texture is blurred (e.g. the trees in the yellow bounding box). [Plot] As a function of y , the PSNR of the reconstructed image is worse than in our method .

served. The restored path radiance is shown in Fig. 8. It is equivalent to the distance map.

6. Discussion

Our proposed method estimates coupled unknowns: scene radiance and path radiance. Instead of a noisy estimate, we obtained a regularized solution. This is obtained using a minimization scheme, which adapts to important characteristic of the scene: edges and distances. We believe that this framework can deal with other non-uniform degradation models. This would require changing the fidelity term according to the physical model. Also the weights should change according to the distance dependency.

Note that apart from attenuation and scene radiance, the raw image is also blurred due to multiple scattering. We now explore removal of this blur using distance-adaptive deconvolution, within the same framework.

Acknowledgments

We thank Einav Namer for the image acquisition. Ran Kaftory gratefully acknowledges the special doctoral fellowship awarded by HP. Yoav Schechner is a Landau Fellow - supported by the Taub Foundation, and an Alon Fellow. This research was supported by the Israeli Science Foundation (Grant 315/04) and the ‘Dvorah’ Fund of the Technion. It was conducted in the Ollendorff Minerva Center in the Elect. Eng. Dept. at the Technion. Minerva is funded through the BMBF.



Figure 5. Real polarization filtered images taken in a hazy day. [Left] I^{\max} . [Right] I^{\min} .



Figure 7. Variational-based restoration of the scene in Fig. 5. The restoration has low noise, without excessive blur.



Figure 6. Noisy restoration of the images in Fig. 5 using Eqs. (7,8) and (9). The restoration is noisy, especially in pixels corresponding to the distant mountain.



Figure 8. The estimated path radiance. Dark pixels indicate higher path radiance and a larger distance to the objects shown in Fig. 5.

Appendix

The minimization scheme in Sec. 4.3 involves a nonlinear PDE. This PDE can be solved by lagging the nonlinear term of the Beltrami operator one iteration behind. Then, for example, by knowing $\hat{R}_\lambda^{(t)}$ in some iteration t , $\hat{R}_\lambda^{(t+1)}$ can be found by solving a linear PDE.

In operator form, we give the following operator definitions:

$$\mathcal{L}_R(\hat{R}_\lambda^{(t)})\hat{R}_\lambda^{(t+1)} \equiv 2 \left(1 - \frac{\hat{A}_\lambda^{(t)}}{a_\lambda^\infty} \right)^2 \hat{R}_\lambda^{(t+1)} \quad (30)$$

$$- \frac{\alpha(\hat{A}_\lambda^{(t)})^2 \mathcal{D} \left[\sqrt{\det(\mathbf{G}_{R^{(t)}})} (\mathbf{G}_{R^{(t)}})^{-1} \nabla \hat{R}_\lambda^{(t+1)} \right]}{(a_\lambda^\infty)^2 \sqrt{\det(\mathbf{G}_{R^{(t)}})}},$$

$$\mathcal{L}_A(\hat{A}_\lambda^{(t)})\hat{A}_\lambda^{(t+1)} \equiv 2\hat{A}_\lambda^{(t+1)} \left[\left(\frac{\hat{R}_\lambda^{(t+1)}}{a_\lambda^\infty} - 1 \right)^2 + p_\lambda^2 \right] \quad (31)$$

$$- \frac{\alpha \mathcal{D} \left[\sqrt{\det(\mathbf{G}_{A^{(t)}})} (\mathbf{G}_{A^{(t)}})^{-1} \nabla \hat{A}_\lambda^{(t+1)} \right]}{\sqrt{\det(\mathbf{G}_{A^{(t)}})}},$$

$$\mathcal{K}_R \equiv \left(1 - \hat{A}_\lambda^{(t)}/a_\lambda^\infty \right) \left(I_\lambda^{\min} + I_\lambda^{\max} - 2\hat{A}_\lambda^{(t)} \right) \quad (32)$$

$$\mathcal{K}_A \equiv \left(\hat{R}_\lambda^{(t+1)}/a_\lambda^\infty - 1 \right) \left(2\hat{R}_\lambda^{(t+1)} - I_\lambda^{\max} - I_\lambda^{\min} \right) +$$

$$+ p_\lambda \left(I_\lambda^{\max} - I_\lambda^{\min} \right). \quad (33)$$

Eqs. (28) and (29) can be written as

$$\mathcal{L}_R(\hat{R}_\lambda^{(t)})\hat{R}_\lambda^{(t+1)} - \mathcal{K}_R = 0 \quad (34)$$

and

$$\mathcal{L}_A(\hat{A}_\lambda^{(t)})\hat{A}_\lambda^{(t+1)} - \mathcal{K}_A = 0, \quad (35)$$

respectively. A direct solution for these equations, as required by each minimization step, would be

$$\hat{R}_\lambda^{(t+1)} = \mathcal{L}_R(\hat{R}_\lambda^{(t)})^{-1}\mathcal{K}_R = 0 \quad (36)$$

and

$$\hat{A}_\lambda^{(t+1)} = \mathcal{L}_A(\hat{A}_\lambda^{(t)})^{-1}\mathcal{K}_A = 0. \quad (37)$$

Since it is difficult to find and invert $\mathcal{L}_R(\hat{R}_\lambda^{(t)})$ and $\mathcal{L}_A(\hat{A}_\lambda^{(t)})$, the steps $d\hat{R}_\lambda^{(t)}$ and $d\hat{A}_\lambda^{(t)}$ are introduced

$$\hat{R}_\lambda^{(t+1)} = \hat{R}_\lambda^{(t)} + d\hat{R}_\lambda^{(t)}, \quad \hat{A}_\lambda^{(t+1)} = \hat{A}_\lambda^{(t)} + d\hat{A}_\lambda^{(t)}. \quad (38)$$

These steps can be found by

$$\mathcal{L}_R(\hat{R}_\lambda^{(t)})d\hat{R}_\lambda^{(t)} = \mathcal{K}_R - \mathcal{L}_R(\hat{R}_\lambda^{(t)})\hat{R}_\lambda^{(t)}, \quad (39)$$

and

$$\mathcal{L}_A(\hat{A}_\lambda^{(t)})d\hat{A}_\lambda^{(t)} = \mathcal{K}_A - \mathcal{L}_A(\hat{A}_\lambda^{(t)})\hat{A}_\lambda^{(t)}. \quad (40)$$

Eqs. (39) and (40) can be solved using the conjugate gradient algorithm.

References

- [1] M. Bertalmio, P. Fort, and D. Sanchez-Crespo. Real-time, accurate depth of field using anisotropic diffusion and programmable graphics cards. In *Proc. IEEE 3DPVT'04*, pages 767–773, 2004.
- [2] T. Boulton. Dove: Dolphin omni-directional video equipment. In *Proc. IASTED Int. Conf. on Robotics and Automation*, pages 214–220, 2000.
- [3] A. Buades, B. Coll, and J. Morel. A review of image denoising algorithms, with a new one. *Multiscale Modeling and Simulation (SIAM)*, 4:490–530, 2005.
- [4] T. F. Chan and C. Wong. Total variation blind deconvolution. *IEEE Trans. Image Processing*, 7:370–375, 1998.
- [5] F. Cozman and E. Kroktov. Depth from scattering. In *Proc. IEEE CVPR*, pages 801–806, 1997.
- [6] S. G. Demos and R. R. Alfano. Optical polarization imaging. *App. Opt.*, 36:150–155, 1997.
- [7] Y. Erez, Y. Y. Schechner, and D. Adam. Ultrasound image denoising by spatially varying frequency compounding. In *Proc. DAGM Symposium*, pages 1–10, 2006.
- [8] P. Favaro and S. Soatto. Seeing beyond occlusions (and other marvels of a finite lens aperture). In *Proc. IEEE CVPR*, volume 2, pages 579–586, 2003.
- [9] J. S. Jaffe. Computer modelling and the design of optimal underwater imaging systems. *IEEE J. Oceanic Eng.*, 15:101–111, 1990.
- [10] R. Kaftory, N. Sochen, and Y. Zeevi. Variational blind deconvolution of multi-channel images. *Int. J. Imaging Science and Technology*, 15:56–63, 2005.
- [11] R. Kimmel, R. Malladi, and N. Sochen. Images as embedded maps and minimal surfaces: Movies, color, texture and volumetric medical images. *Int. J. Comp. Vision*, 39:111–129, 2000.
- [12] D. M. Kocak and F. M. Caimi. *Computer vision in ocean engineering*, chapter 4. CRC Press, 2001.
- [13] N. S. Kopeika. *A System Engineering Approach to Imaging*, pages 446–452. Bellingham, Wash: SPIE, 1998.
- [14] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas. Synthetic aperture confocal imaging. *ACM TOG*, 23:825–834, 2004.
- [15] S. G. Narasimhan and S. K. Nayar. Interactive deweathering of an image using physical models. In *Proc. IEEE Workshop on Color and Photometric Methods in Computer Vision*, 2003.
- [16] S. G. Narasimhan, S. K. Nayar, B. Sun, and S. J. Koppal. Structured light in scattering media. In *Proc. IEEE ICCV*, pages 420–427, 2005.
- [17] S. Negahdaripour and H. Madjidi. Stereovision imaging on submersible platforms for 3d mapping of benthic habitats and sea floor structures. *IEEE J. Oceanic Eng.*, 28:625–650, 2003.
- [18] J. P. Oakley and B. L. Satherley. Improving image quality in poor visibility conditions using a physical model for contrast degradation. *IEEE Trans. IP*, 78:167–179, 1998.
- [19] A. M. Polyakov. Physics letters. 103B:207, 1981.
- [20] Y. Y. Schechner and Y. Averbuch. Regularized image recovery in scattering media. *To be published in IEEE Trans. PAMI*, 2007.
- [21] Y. Y. Schechner and N. Karpel. Clear underwater vision. In *Proc. IEEE CVPR*, volume 1, pages I–536–543, 2004.
- [22] Y. Y. Schechner and N. Karpel. Recovery of underwater visibility and structure by polarization analysis. *IEEE J. Oceanic Eng.*, 30:570–587, 2005.
- [23] Y. Y. Schechner, S. G. Narasimhan, and S. K. Nayar. Polarization-based vision through haze. *Applied Optics*, 42:511–525, 2003.
- [24] S. Shwartz, E. Namer, and Y. Y. Schechner. Blind haze separation. In *Proc. IEEE CVPR*, pages 1984–1991, 2006.
- [25] N. Sochen, R. Kimmel, and R. Malladi. A general framework for low level vision. *IEEE Trans. IP*, 7:310–318, 1998.
- [26] J. S. Talyor and L. B. Wolff. Partial polarization signature results from the field testing of the shallow water real-time imaging polarimeter (shrimp). In *Proc. MTS/IEEE OCEANS*, volume 1, pages 107–116, 2001.
- [27] T. Treibitz and Y. Y. Schechner. Instant 3Descatter. In *Proc. IEEE CVPR*, volume 2, pages 1861–1868, 2006.
- [28] D. Tschumperle and R. Deriche. Vector-valued image regularization with pde's : A common framework for different applications. *IEEE Tran. PAMI*, 27:506–517, 2005.
- [29] C. Vogel and M. Oman. Iterative methods for total variation denoising. *SIAM J. Scientific Computing*, 7:227–238, 1996.