Towards a model-free denoising of underwater optical images

Andreas Arnold-Bos, Jean-Philippe Malkasse, Gilles Kervern

TOWARDS A MODEL-FREE DENOISING OF UNDERWATER OPTICAL IMAGES

Andreas Arnold-Bos* , Jean-Philippe Malkasse†, Gilles Kervern †

*ENSIETA / E^{3}I^2
2, rue François Verny, 29806 Brest cedex 9, France
Email: arnoldan@ensieta.fr

†Thales Underwater Systems
Route de S^{te} Anne du Portzic, 29238 Brest cedex 3, France
Email: firstname.lastname@fr.thalesgroup.com

***Abstract*— A major hindrance to underwater operations using cameras comes from the light absorption and scattering by the marine environment, which limits the visibility distance up to a few meters in coastal waters when using low-end cameras. We propose a complete preprocessing framework able to handle the entire spectrum of noises present in underwater images. We show that most, if not all of this preprocessing can be done with very generic methods that do need any knowledge of the scene or of the turbidity characteristics of the water, while still remaining coherent with the underwater images formation model.**

I. INTRODUCTION

The increasing interest in Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs) for underwater operations has called for the development of efficient, widely available sensors. Optical cameras meet such requirements and have the additional benefit of having an excellent resolution. However, the major obstacle to their use use is that light, unlike sound, is poorly propagated in the water. The effective range of visibility is limited to about twenty meters in clear water and less than three meters in turbid, coastal waters.

These poor performances are explained by the peculiar propagation properties of light in the aquatic medium [8], [9], [14]. First, a ray of light is exponentially attenuated as it travels in the water so the background of the scene will be poorly contrasted and hazy. The visibility range may indeed be augmented with artificial lighting. Unfortunately, water will reflect a significant fraction of the light power towards the camera before it actually reaches the objects in the scene. This process, known as backward scattering, causes a characteristic glowing veil that superimposes itself on the image and hides the scene. Finally, forward scattering, i.e. randomly deviated light on its way from an object to the camera, causes blurring of the image features. One could also consider macroscopic floating particles ("marine snow") as being unwanted signal, although they belong to the scene. In orders of magnitude, backscattering and marine snow are the greatest degradation factors, attenuation comes second and forward scattering follows closely. Figure 1 is an example of a fairly typical underwater image taken in daylight conditions.

When specialized hardware such as range gated light systems [5] or polarized cameras [12] are not available, image quality must be improved via software processing. These algorithms deal either via deconvolution or via generic enhancement methods (GEMs) such as contrast enhancement that do not rely on any physical model. Both approaches have their advantages and flaws. Deconvolution is rigorous but hard to perform in a real situation because the parameters of the model are unknown. In controlled situations, deconvolution can be complete, but in natural environments, only deconvolution of forward scattering with restrictive assumptions on the point of view have been achieved [7], [10]. GEMs can be used without these limitations and provide fair results, which explains their popularity [3]; however they lack the theoretical backup that deconvolution provides. Therefore GEMs are often used empirically, focusing on the prominent aspects of noise (contrast inequalities) without really caring about the fine intricacies of the image formation process. All in all, the *real* effect of GEMs on underwater images is unknown, which sheds doubts on their relevance and performances.

Fig. 1. A typical underwater image (Picture by Anthony DeLullo).

We investigate here the possibility of addressing the whole range of noises present in underwater images by using a rational combination of deconvolution and generic enhancement methods. Our paper is structured as follows. First we present the image formation model (section II). In section III we describe a preprocessing algorithm designed to address the whole spectrum of noises present in underwater images acquired with standard optical cameras. Section IV is dedicated to the analysis of our results and a discussion on the limits of our approach. We exhibit promising results using both simulated and real images.

II. THE IMAGE FORMATION MODEL

First, let us introduce some notation conventions: the generic 3-D object point will have coordinates x, y, z and its image has coordinates (x', y') on the image formation plane. We call $Z(x', y')$ the *depth map*, that is, the distance of the camera to the object that would be projected in (x', y') if the propagating medium is of unit transmittance (ideal medium).

The scene is represented by the *reflectance map* $R(x', y')$ which more or less represents its intrinsic color (we represent it as a gray level between 0, black, and 1, white). This is truly the information we wish to recover.

The perceived light intensity in the camera, when the medium is ideal, is influenced by the incident light power on the object point B (beam pattern), the material characteristics C_M (specularity, diffusivity, etc.) and the lens characteristics (focal length, aperture)

C_L. Thus the received light power is:

$$I(x', y') = R.B.C_M.C_L(x', y') = R.F(x, y)$$

F is called the illumination function.

When the medium is not ideal, such as under water, the semi-empirical image formation model commonly used [4], [8] states that the received signal I_R is the sum of three components, plus some noise n:

$$I_R = I_D + I_{FS} + I_{BS} + n \tag{1}$$

The *direct component* I_D is an attenuated version of I with distance due to the absorption of light by floating particles:

$$I_D = I.e^{-cZ} \tag{2}$$

where c is called the volume attenuation constant, in m^{-1}. Typical values for c^{-1} range from 3 m for coastal waters up to 20 m for clear, deep ocean waters.

However, an elementary volume of water and floating particles does not only absorb photons but does also deviate incident rays of light. This process is called scattering. The volume scattering function, which gives the probability for a ray of light to be deviated of an angle θ from its direction of propagation, has been determined experimentally [9].

The *forward scattered component* I_{FS} is an attenuated and blurred version of I due to the random deviation of light on its way to the camera by floating particles. The point spread function (PSF) is a Cauchy function, hence its Fourier transform is an inverse exponential:

$$I_{FS} = (e^{-g.Z} - e^{-c.Z})\mathcal{F}^{-1}\left(e^{-B.Z.i.f}\right) \star I \tag{3}$$

Here, $\mathcal{F}^{-1}$ stands for the inverse Fourier transform, $|g| < c$ and $B < c$ are empirical constants, i is a scaling constant to convert image plane units into meters, and f is the angular frequency in the image plane.

The last component is the *backscattered component*, which corresponds to light on their way from the source to the objects of the scene, that are reflected by the water and floating particles, then attenuated and forward scattered on their way to the camera. Backscattering acts like a glaring veil superimposed on the objects in the scene. Therefore, it is generally seen as the predominant limiting factor in underwater optics. An exact derivation of the backscattering may be found in [8], but we do not really need it here. We shall note that component I_{BS} and it suffices to know that it is nearly uncorrelated from I_D and I_{FS}, which are directly linked to the signal coming from the scene.

We illustrate the combined contribution of attenuation, forward scattering, and backscattering in a simulation with the same setting than the one proposed by McGlamery [8]. The image is planar and made of a series of stripes. The camera and the spotlight are three meters away from the scene, and the spotlight is five meters away from the camera. For the sake of simplicity, the image is unidimensional, that is, only the intensity along axis x is given (fig. 2). The resulting image clearly shows the predominance of backscattering in the final image (fig. 3). We will use this simulation as a companion tool in this paper to illustrate the effect of our processing.

III. Detailed algorithm

In this section, we describe our preprocessing algorithm to enhance images. The algorithm workflow is given by figure 4. First, we describe the use of GEMs to tackle problems caused by backscattering, attenuation and lighting discrepancies, and we justify our approach with respect to the image formation model. The following last

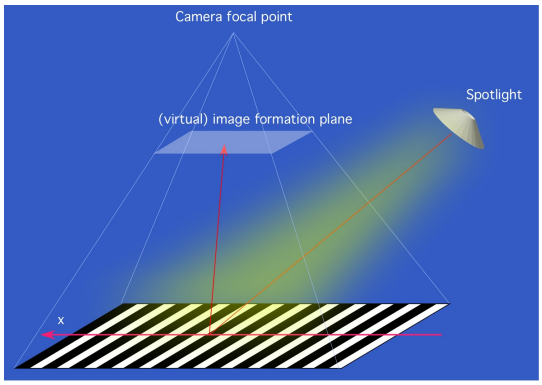

Fig. 2. Simulated configuration ($g = 0.195\ m^{-1}$, $g = 0.25\ m^{-1}$, $B = 0.375\ m^{-1}$)

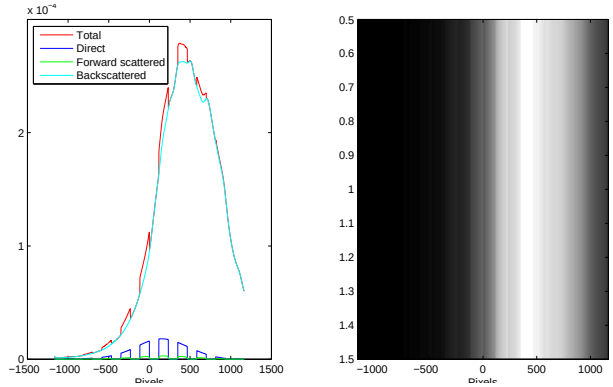

Fig. 3. Contribution of I_D, I_{FS} and I_{BS} in the final image

stage deals with the removal of remaining noises using an adaptive smoothing algorithm. Finally, an optional stage helps to remove the effects of forward scattering.

A. Elimination of attenuation, backscattering and other lighting discrepancies

The intrinsic illumination pattern of the spotlight, as well as attenuation with depth and backscattering are the cause to numerous discrepancies in contrast. Early attempts to enhance contrast involved histogram equalization; yet as the lighting in the image is unequal, global histogram equalization does not help much and local methods are needed. These methods are reviewed in depth by García, Nicosevici and Cufi [3] and benchmarked empirically (that is, visually) on a series of images. To enhance contrast in underwater images, local variations of the histogram equalization such as histogram equalization on a sliding window can be used, but these methods generally increase the noise level. Better methods rely on the fact that the shadows present in the images are spatially slowly varying. Thus the estimated level of lighting can be estimated using either a very loose spline approximation, or a series of morphological dilatations using a large structuring element, or simply using a low-pass filtering such as a Gaussian blur. Contrast can then be enhanced by subtracting the estimated shadows to the original image. Instead of subtraction, division can also be used: this comes essentially to estimating the reflectance of the image by assuming that the low-pass version is the illumination field. This last method is certainly reasonable and efficient to overcome nonuniform illumination caused by auxiliary lighting systems when water turbidity is low. Nothing, however, gives guarantees of optimality for this method when backscattering becomes important. Consequently, we dedicate the following paragraph to an analysis its behavior with respect the image formation process.

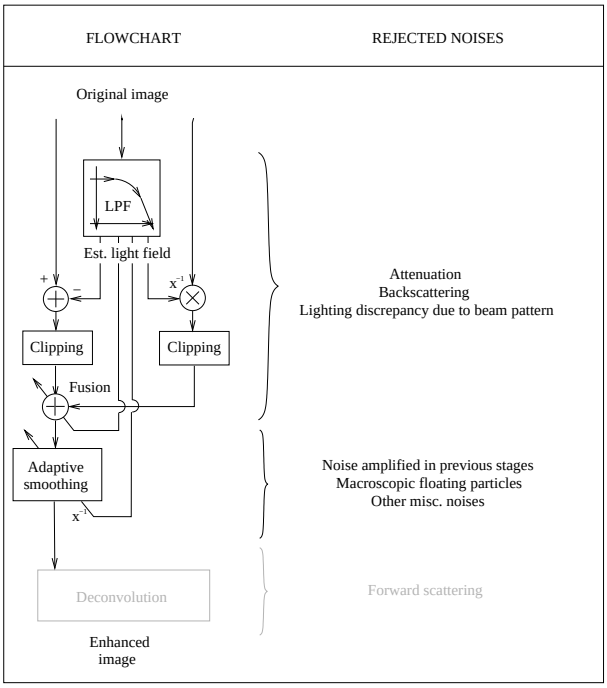

Fig. 4. Block diagram of our underwater denoising framework

We start by stating that the illumination function $F(x', y')$ and the backscattered component are spatially slowly varying in front of the reflectance $R(x', y')$. This is most often true, except in some ill-posed cases, where small objects are in front of a differently colored background (small fishes over a sandy bottom, for example). Also, I_{BS} and F are slowly varying in front of the extent of the PSF of the forward scattered component, so a first approximation of the PSF of the forward scattered component is a Dirac. Consequently, the forward scattered component can be approximated by:

$$I_{FS} \approx (e^{-g.Z} - e^{-c.Z}).I \tag{4}$$

Injecting (2) and the above into (1) yields:

$$I_R = e^{-g.Z}.I + I_{BS} + n \tag{5}$$

When computing a low-pass approximation of the received image by convolution with a Gaussian function $G_\sigma(x', y')$ of spatial extent σ very large in front of the spatial variations of R, we obtain:

$$I_{LP} \approx I_{BS} + F.e^{-g.Z}.\bar{R} = I_{BS} + F'.\bar{R} \tag{6}$$

In this expression, the term $F' = F.e^{-g.Z}$ represents an effective illumination function and $\bar{R}(x', y') = G_\sigma \star R(x', y')$ is the local mean of R. Now from the above equations, it appears that the natural method to reject backscattering would consist in subtracting I_{LP} to the received image I_R (subtraction method) :

$$I_{\text{sub}} = I_R - I_{LP} \approx F'.(R - \bar{R}) + n \tag{7}$$

The result of the difference effectively eliminates backscattering; however, the effective illumination factor remains. Consequently the reconstruction is only good at close distance to the camera, and in zones where the scene is well lit by the sources.

Another method is the local normalization, also known as the illumination- reflectance model. It consists in dividing the received image by its low-pass version[1]:

[1]In this paper, multiplication and division used for matrices are element-wise and not traditional matrix multiplication or division.

$$I_{\text{div}} = I_R/I_{LP} = \frac{I_{BS} + F'.R + n}{I_{BS} + F'.\bar{R}} \tag{8}$$

By introducing $V = I_{BS}/(F'.\bar{R})$, which we call the *veiling ratio*, the above equation can also be written:

$$I_{\text{div}} = \frac{V}{V+1} + \frac{1}{V+1}\frac{R}{\bar{R}} + \frac{n}{I_{LP}} \tag{9}$$

When $V \to 0$, *i.e.* when no backscattering is present in the image, then:

$$I_{\text{div}} = \frac{R}{\bar{R}} + \frac{n}{F'\bar{R}} \tag{10}$$

which is the reflectance, normalized by the local mean reflectance of the image, hence the name of the method. On the contrary, when high backscattering is present, $V \to \infty$, constant $V/(V+1)$ tends towards one and $1/(V+1)$ tends towards zero. As a result, the intensity histogram of the divided image tends to be centered around one with a small dispersion (especially for poorly contrasted images) except for a few outliers. This is the explanation why contrast equalization by division is always followed by histogram clipping to reject these outliers:

$$I_{\text{clip}}(i,j) = \begin{cases} i_1 & \text{if } I_{\text{div}}(i,j) < i_1, \\ i_2 & \text{if } I_{\text{div}}(i,j) > i_2, \\ I_{\text{div}}(i,j) & \text{else.} \end{cases} \tag{11}$$

where $[i_1..i_2]$ is typically the 95% percentile interval of the intensity distribution of I_{div}. Then the dynamic of the image may be expanded so that all available intensity levels $[0..i_{\max}]$ are occupied:

$$I_{\text{clip}}(i,j) \leftarrow \frac{I_{\text{clip}}(i,j) - i_1}{i_2 - i_1}.i_{\max} \tag{12}$$

As a crude approximation, histogram clipping amounts to reject the constant term $V/(V+1)$ in equation 9.

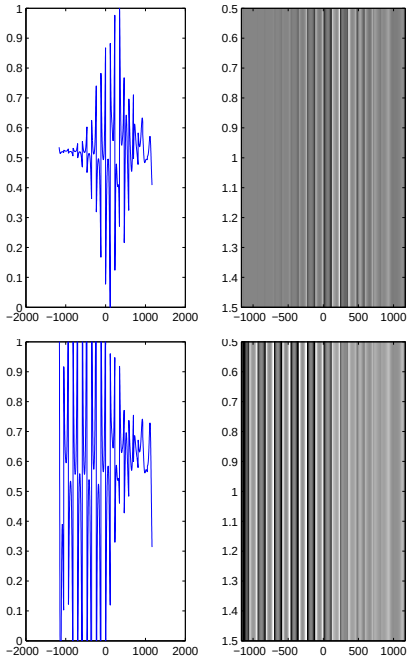

Fig. 5. I_{sub} (above) and I_{div} (below) for simulated image (after clipping and scaling)

To decide whether dividing or substracting is the best, we can use two criteria: *i)* the proximity of the recovered image to the ideal signal and *ii)* the local contrast, which is directly linked to the magnitude of the gradient of the image. The comparison can be both made on

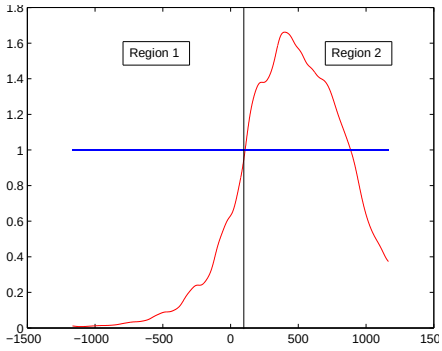

Fig. 6. Normalized gradient ratio S

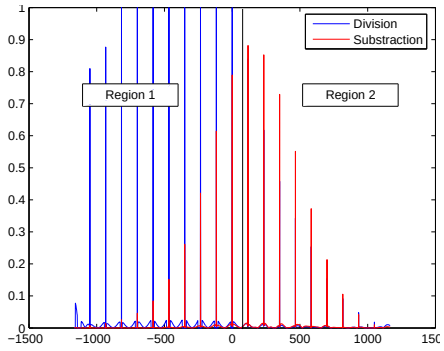

Fig. 7. Gradient of I_{sub} and I_{div}

a theoretical basis and an experimental basis, using the simulator, so as to have a ground truth.

Contrast provides an essential means to evaluate the quality of the restored image. The contrast is directly linked to the gradient of the image. The greater the gradient, the greater the contrast. Assuming zero noise, the gradient of the image recovered by the difference method is:

$$\overrightarrow{\nabla} I_{\text{sub}} = \overrightarrow{\nabla} F'.R + F'.\overrightarrow{\nabla} R \approx F'.\overrightarrow{\nabla} R \tag{13}$$

while for the localy normalized image, it is:

$$\overrightarrow{\nabla} I_{\text{div}} = \frac{F'.\overrightarrow{\nabla} R}{I_{BS} + F'\bar{R}} = \frac{F'.\overrightarrow{\nabla} R}{I_{LP}} \tag{14}$$

The normalized ratio of the two gradients is:

$$S = \overrightarrow{\nabla} I_{\text{sub}} / \overrightarrow{\nabla} I_{\text{div}} = F'\bar{R} + I_{BS} = F'.\bar{R}(1+V) \tag{15}$$

Often, gradients are not normalized, so one must multiply it by factor $K_{\text{div}}/K_{\text{sub}}$, where K_{div} and K_{sub} are respectively the amplitude of the values range of the images estimated by division and subtraction. When the normalized ratio of the two gradients is greater than one, the image estimated after subtraction has a greater local gradient than the image estimated after division, so the subtraction method must be preferred. The opposite is true when the normalized ratio is inferior to one. We observe that the contribution of backscattering is better estimated by an additive illumination field than a multiplicative contribution in well-lighted zones of the image. Figures 6 and 7 show respectively the fusion ratio S and the gradient of I_{sub} and I_{div} for the simulated image (after clipping and normalization); in region 1, $S < 1$ and indeed, $\overrightarrow{\nabla} I_{\text{div}} > \overrightarrow{\nabla} I_{\text{sub}}$. The inverse is true in region 2.

We propose the following fusion process to have the best compromise between these two methods. First, images I_{sub} and I_{div} are computed, along with the normalized gradient ratio S, which only depends on the low-pass version of the image I_{LP} and the ranges K_{sub} and K_{div}. Next, I_{sub} and I_{div} are histogram-clipped so that their values ranges is about 95% of the confidence interval of their initial values. The images are then scaled between 0 and 1. The "optimal" contrast-equalized image is then:

$$I_{\text{eq}} = \frac{S.I_{\text{sub}} + \frac{1}{S} I_{\text{div}}}{S + \frac{1}{S}} \tag{16}$$

The clipping and the scaling will ensure that no color disparities exist between the two fused images.

B. Rejecting remaining noises

Generally, contrast equalization will raise the noise level in poorly contrasted areas of the original image. Signal to noise ratio would stay constant after equalization were the images not quantified; but the fixed color quantization step induces strong errors in dark zones. These errors are amplified by histogram clipping, because the color dynamic is so small that clipping will be equivalent to crude binarisation. Moreover, macroscopic floating particles which have images of roughly the size of a pixel induce a strong clutter after equalization. Although being part of the scene, these particles are unwanted and can be seen as a local noise. Compared to the substraction method, the illumination-reflectance model has a slightly more important effect on noise in dark zones. Indeed, the local gain of the contrast equalization process is:

$$G(i,j) = \frac{1}{I_{\text{LP}}(i,j)} . \frac{i_{\max}}{i_2 - i_1} \tag{17}$$

The standard deviation of the noise at location (i,j) will also be multiplied by $G(i,j)$; consequently, the lower the intensity, the higher the noise level. Gain G affects all types of noises remaining in the image, as the additive Gaussian noise of the camera electronics, the quantification noise, which can both be considered as Gaussian white noise.

When automatic object recognition is the final goal, good detectability of edges is important and all the previously mentioned noises may cause false alarms. Although edge detectors such as the Canny-Deriche algorithm [2] use a smoothing filter to eliminate noise, this filter is linear and does not adapt itself to the local noise intensity or the presence of edges. Thus the filtering level is overevaluated on some parts of the image. This can be prevented if noise is rejected using an adaptive image smoothing stage after contrast equalization. The principle behind adaptive smoothing is that the image is filtered less where edges are present, hence preserving their good visibility. In an earlier paper [1], we experimented several adaptive smoothing algorithms and showed that they could significantly improve edge detection and visibility in heavy backscattering conditions. We have selected a variation of a denoising algorithm presented by Kovesi [6], who proposed to use a multiscale denoising algorithm based on complex-valued log-Gabor wavelets. Compared to real-valued wavelets, the benefit of using complex wavelets is that we have an approximation of the amplitude and the phase of the signal at various scales. Then, only the amplitude of the coefficients is shrunk by the estimated noise contribution at that scale. As a result, the phase of the signal and thus good localization of edges is preserved. Most importantly, this algorithm may be used in a totally automatic way, because an elegant method exists to determine the amount of noise at each scale, assuming that the noise is white and Gaussian. Although this hypothesis does not hold in our case, it is of very negligible

importance in practice. However, unlike Kovesi's original algorithm, we do not use a constant threshold over space, but instead correct it by factor $G(i,j)$ so that heavier filtering is done where the initial image was dark. A great deal of other adaptive denoising algorithms exist that may be worth investigating in our context, for example curvelets [13] or anisotropic filtering [11]. However, when it comes to choosing the algorithm, we believe that speed and absence of manual tuning should be priviledged over performance, because performance is hard to quantify and no significant difference in denoising between the algorithms is to be expected anyway.

C. Forward scattering rejection

Forward scattering has a blurring effect which depends on the distance travelled by light. If the scene is assumed to be plane, as is the case in our simulation, and if the turbidity parameters are known, then inverse filtering can be performed. However, special care should be taken during this process, as the inverse filter has the form $f \mapsto e^{b \cdot f}$, which tends to amplify noise at high frequencies, even when Wiener filtering is used.

In real-life conditions, the scene is rarely planar; secondly the turbidity parameters are often unknown. Experimentally, it appears that the first hypothesis is not very important as noticeable improvements to the sharpness of the edges can be obtained by assuming the scene to be planar even when it is not. The second problem is the most limiting, as water turbidity parameters are only scarcely known in tables. To this date, the best approach to solve this problem in a real application is perhaps the one proposed by Olmos in her PhD thesis [10]. The implementation is as follows. Typical *a priori* values for the depth and turbidity coefficient are used as starting values in an iterative process where one optimizes a cost function representing the contrast of the deconvolved image. Special care must be taken when designing the optimization process so as not to converge to local minima. Figure 8 shows the effect of inverse Wiener filtering on the simulated image where all parameters are known. Although a fairly noticeable enhancement can be observed when the medium is moderately turbid, this method does not remove backscattering, which explains why deconvolution cannot be used alone.

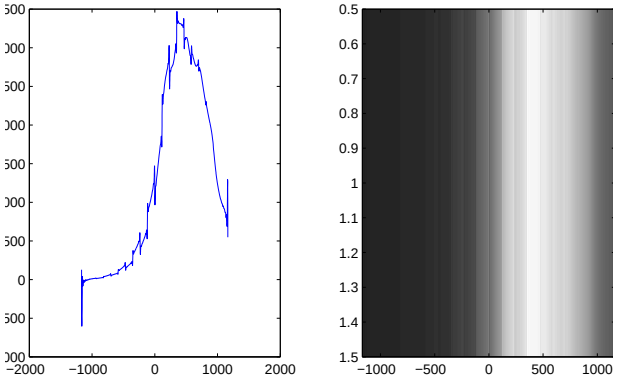

Fig. 8. Simulated image after deconvolution of forward scattering using true model parameters

IV. Experiments and discussion

Figures 9 to 14 present various stages of our denoising algorithm applied to the image shown in fig. 1. Only the luminance channel of the image has been processed. While the contrast is better in the final equalized image (fig. 11), noise remains; fig. 14 is the image (in color) after adaptive denoising and deconvolution; notice how the signal to noise ratio has improved.

The bulk of the work is performed by the contrast equalization step. Yet this stage can lead to severe problems when contrast equalization is pushed to its limits. Removing low-frequencies is good because the local signal mean is a good estimation of the lighting field. However this introduces bias, as uniform zones of typical size similar to the size of the low-pass filter kernel will be all normalized to one; consequently the color of object is lost and only the contrast of high frequency details is enhanced. Secondly, enhancing high frequency components is –again– desirable, because they define the edges. But this corrupts the perceptual relative importance of edges in the image. This means that slight details pertaining to the objects texture will become as visible, that is, as significant than objects contour, which definitely play a bigger role in human or machine object recognition. These problems are inherent to all contrast equalization techniques. Furthermore, the use of Gaussian blurring when estimating the illumination field may be the cause of halos around strongly contrasted objects of size similar to the Gaussian kernel (see, for example, the dark fish over the sandy bottom in our example image). To limit those effects, the width of the Gaussian kernel must be very large. Here, the standard deviation of the kernel was 20% of the image diagonal length. In a nutshell, a very general statement about contrast equalization would be that it improves the image at its first order and beyond (which explains why edge detection is better, as the gradient is higher), yet recovery of the image at the order zero is a hard thing to obtain.

Another interesting question is the importance of deconvolution in the restoration process and where to put it in the preprocessing pipeline. When used as the first stage, the noise level increase that comes with deconvolution should normally be counterbalanced by the adaptive denoising coming at the end of our algorithm. In fact, this is not true, as shown in fig. 13, to be compared with fig. 12 where no deconvolution is used. Deconvolution can also be put after contrast enhancement and denoising. Doing this is avoids noise being amplified by deconvolution as it was removed before. Yet results obtained by adding deconvolution at the end of our process (fig. 14) bring only marginally better benefits than without its use. For all practical purposes, we feel that deconvolution can very often be safely omitted without any noticeable loss in edge detection.

V. Conclusion

We presented a complete preprocessing framework for underwater images. First, a contrast equalization system was proposed to reject backscattering, attenuation, and lighting inequalities caused by the beam pattern. This system does not need any knowledge of the medium characteristics yet its action on the images is clearly understood and coherent with the image formation process. Then, remaining noises corresponding to sensor noise, floating particles images, and miscellaneous quantification errors, are suppressed using a very generic self-tuning wavelet-based algorithm. Finally, we showed that forward scattering *could* be addressed as well using deconvolution but that the gains are negligible compared to the effect of previous stages.

Acknowledgment

Andreas Arnold-Bos would like to thank Thales Underwater Systems for having welcomed him as a student intern during his *Diplôme d'Etudes Approfondies* in Signal and Image processing (Supaero, Toulouse) final project.

References

[1] A. Arnold-Bos, J.-P. Malkasse, and G. Kervern. A preprocessing framework for automatic underwater images denoising. In *Proceedings of the Fifth European Conference on Propagation and Systems*, Brest, France, Mar. 2005.

[2] R. Deriche. Using Canny's criteria to derive a recursively implemented optimal edge detector. *The International Journal Of Computer Vision*, 1(2):167–187, May 1987.

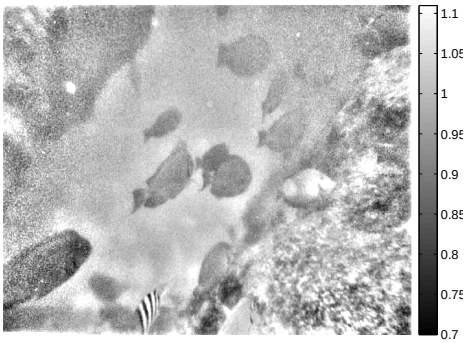

Fig. 9. Real image: I_{div}

Fig. 10. Real image: I_{sub}

[3] R. Garcia, T. Nicosevici, and X. Cufi. On the way to solve lighting problems in underwater imaging. In *Proceedings of the IEEE Conference on Oceans 2002*, pages 1018–1024, Oct. 2002.

[4] J. S. Jaffe. Computer modelling and the design of optimal underwater imaging systems. *IEEE Journal of Oceanic Engineering*, 15(2), Apr. 1990.

[5] J. S. Jaffe, J. McLean, M. P. Strand, and K. D. Moore. Underwater optical imaging: Status and prospects. *Oceanography*, 14:66–76, 2001.

[6] P. Kovesi. Phase preserving denoising of images. In *Proceedings of The Australian Pattern Recognition Society Conference DICTA'99*, Dec. 1999.

[7] Z. Liu, Y. Yu, K. Zhang, and H. Huang. Underwater image transmission and blurred image restoration. *SPIE Journal of Optical Engineering*, 40(6):1125–1131, June 2001.

[8] B. L. McGlamery. A computer model for underwater camera systems. *SPIE Ocean Optics VI*, 28, 1979.

[9] A. Morel. Diffusion de la lumière par les eaux de mer; résultats expérimentaux et approche théorique. *AGARD Lecture Series*, pages 3.1.1.–3.1.76, 1973.

[10] A. Olmos. *Detecting underwater man-made objects in unconstrained video images*. PhD thesis, Heriot-Watt University, Department of Computing and Electrical Engineering, Oct. 2002.

[11] P. Perona and J. Malik. Scale-space and edge detection using ansotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990.

[12] Y. Y. Schechner and N. Karpel. Clear underwater vision. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 536–543, 2004.

[13] J.-L. Starck, E. J. Candès, and D. L. Donoho. The curvelet transform for image denoising. *IEEE Transactions on Image Processing*, 11(2):670–684, June 2002.

[14] W. H. Wells. Loss of resolution in water as a result of multiple small-angle scattering. *Journal of the Optical Society of America*, 39(6), 1969.

Fig. 11. Fusion of I_{div} and I_{sub}

Fig. 12. Final image after adaptive smoothing (no forward scattering deconvolution)

Fig. 13. Final image after adaptive smoothing (with forward scattering deconvolution as first processing stage)

Fig. 14. Final image after adaptive smoothing (with forward scattering deconvolution as last processing stage