Bathymetric Side-scan Backscatter Map Restoration based on Data Fusion

E. Coiras*, Y. Petillot, D. M. Lane

Ocean Systems Laboratory
Department of Electrical, Electronic & Computer Engineering
School of Engineering and Physical Sciences
Heriot-Watt University,
EH14 4AS Edinburgh. UK
*corresponding author: E.Coiras@hw.ac.uk

Abstract: Backscatter image mosaics frequently suffer from noticeable visual artifacts and missing regions, which reduce the performance of subsequent feature extraction or classification processes. In this paper, a new fusion-based method for the restoration of bathymetric side-scan backscatter images is presented. First, the values of missing backscatter pixels are estimated using contextual information provided by the underlying bathymetry mesh. As a final step, a histogram regularization process is performed to remove intensity artifacts caused by inaccurate compensation of sensor gain. Restoration results are presented for Simrad EM12S images.

I. INTRODUCTION

This paper briefly describes two techniques we have developed for the restoration of bathymetric side-scan images in project AMASON [1]: fusion-based missing ping estimation and histogram-based intensity compensation. The restoration process is mostly oriented to improve the performance of subsequent feature extraction or classification processes. As a byproduct, the subjective visual quality of the backscatter map is also improved, thus easing the interpretation of the observed scene by a human operator.

II. MISSING PING ESTIMATION

Missing ping data in bathymetric side-scan images is usually compensated by linear interpolation between nearest in-line neighboring values. This technique is easy to implement and can be conveniently applied at the moment of reading each data line from the source file. Performance is, however, generally poor, and usually results in highly noticeable image patches. As an example, top part of Fig. 1 shows a Simrad EM12S bathymetric side-scan image, displaying missing regions corresponding to rejected pings, and the result of restoration by linear interpolation. The image was rendered from raw backscatter return measures following the procedure described in N. C. Mitchell [2].

To improve results, some indirect knowledge about the missing values can be gathered from related contextual information. Assuming a good navigation correction is provided, support values can be found in adjacent swath lines. Still, even incorporating this new information, the application of bilinear or bicubic interpolation for the estimation of missing ping values is a poor solution, resulting in patchy side-scan maps. This is directly related to the noisy nature of side-scan images, in which smooth and gradual tint variations are rarely present. A smooth interpolation between any two known pixel values will undoubtedly result in a smooth patch of consistently deviated values from the local mean gray-level.

A more adequate approach can be found in synthesizing texture structure for the missing regions. This method achieves better results, replicating the noisy character of the backscatter images in the restored areas. Among the many image inpainting techniques proposed in the literature, we chose the method described in A. A. Efros and T. K. Leung [3] for its ease of implementation. In their paper, textures are modeled as Markov Random Fields, assuming the probability distribution for the value of a pixel given the values of its spatial neighborhood is independent of the rest of the image. Yet, to avoid the complications of constructing such a complex model, a simple non-parametric sampling technique is used to replicate its statistical properties. This approach produces fairly good results in the backscatter images (see Fig. 1), successfully improving their visual appearance and the results of subsequent processing routines.

Some minor problems—replication artifacts, long processing times—had yet to be solved, nevertheless. And, more importantly, some justification for the whole restoration procedure was required. A proper physical approach to the problem would result in a more specific and direct method for the particular type of images being considered, possibly improving performance along the way.

0-933957-31-9

Two particular techniques were developed to this end: noise synthesis and support-region selection based on local slope values.

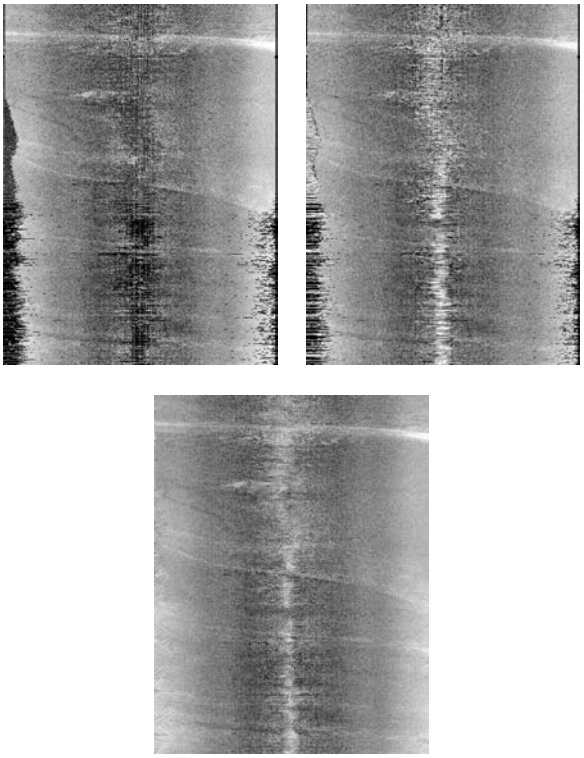

Fig. 1: Results of missing ping restoration procedures. Top left: bathymetric side-scan image; pixels where amplitude data is missing are shown in black. Top right: result of in-line interpolation. Bottom: backscatter map restoration by texture inpainting using the method proposed by A. A. Efros and T. K. Leung.

Support-Regions and Noise Synthesis

The study of the frequency components of the backscatter images proves they are very similar to noise in the scales relevant to our problem, so texture synthesis is probably more than what's actually needed. Using noise synthesis should suffice.

Although several complex and detailed models have been proposed to accurately describe backscatter signal characteristics, the restriction to suitable small neighborhoods with certain properties permits the application of a simple gaussian model to approximate the local gray-level distribution, as explained below.

For a given frequency, models for backscatter return functions depend mainly on two parameters: ensonified material and angle of ensonification [4]. The composition of the scene being observed is normally unknown, but the surface angle is actually being measured by the multi-beam sensor in the form of bathymetry information. Thus, after navigation correction and bathymetry mesh construction, local slope values can be computed for every pixel in the backscatter map.

Restricting the angle of ensonification of the pixels considered for computing the local gray-level distribution effectively means that differences in material composition are directly reflected in the values of the mean and standard deviation computed for the local region. Therefore, replication of the noise distribution presented by neighboring pixels with congruent slopes is an effective way of restoring the values of missing pixels, as ultimately the process is equivalent to assigning the pixel a gray-value corresponding to the type of material of their neighbors under the angle of incidence under which it is being observed.

As a consequence, the proposed reconstruction method preserves periodic or structured features that might be present in the ensonified scene. Of course, they will be propagated through the backscatter map only if they ultimately emerge from corresponding features in the slope-map. Therefore, the underlying idea for the proposed fusion-based restoration is that the presence of edges in the observed scene is directly related to differences in shape and composition of the objects contained in the scene. This assumption has successfully been used before as a foundation for the development of registration techniques for different types of images [5].

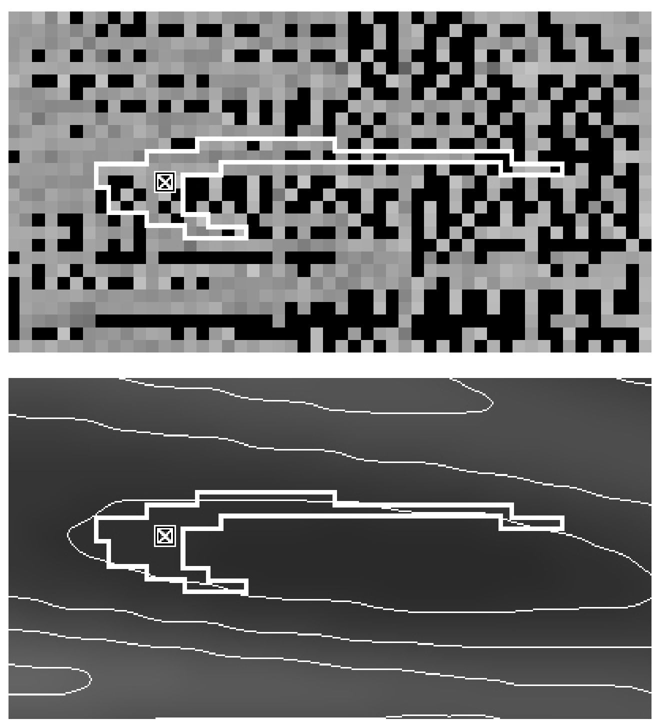

Fig. 2: Support region for a missing pixel of the backscatter image. Top: missing pixel to be restored (signaled with a cross) and its support region, which is determined by region growth on the corresponding slope map. Bottom: contour-lined map of slope values showing the corresponding location of the support region.

Implementation and performance optimization

Using this approach, region growth in the slope map is performed around every missing pixel in the backscatter image. Growth continues until enough region pixels have valid backscatter values (see Fig. 2, where a threshold of 63 pixels was selected). Then a gaussian is fitted to the gray-level distribution of the region and used for synthesizing the value of the missing pixel. Although overall backscatter return-values are better modeled by a Rayleigh function, a gaussian is enough to model the gray-level distributions of the small sets of samples contained in the support-regions.

Pixels are added to the support region according to how close their slope is to the average slope value of the pixels already in the region. To synthesize the value for the missing pixel, a random number is generated according to the parameters of the gaussian distribution computed for the region.

The method is further improved by considering that pixels adjacent to missing regions are actually noisier and brighter than the rest, as Fig. 3 shows. Rejecting those pixels by an initial 4-neighbor dilation of the missing areas yields much better results. A comparison of the noise-synthesis approach with and without this initial dilation stage is shown in Fig. 4.

To minimize processing times, all pixels with missing backscatter values in the connected support-region are synthesized at the same time. These would otherwise be synthesized afterwards by sampling values from a very similar support-region. Besides, experiments have shown this shortcut to have practically no influence on the final result (at least not more than the output variations of the gaussian noise generation routine), yet greatly reduce processing time.

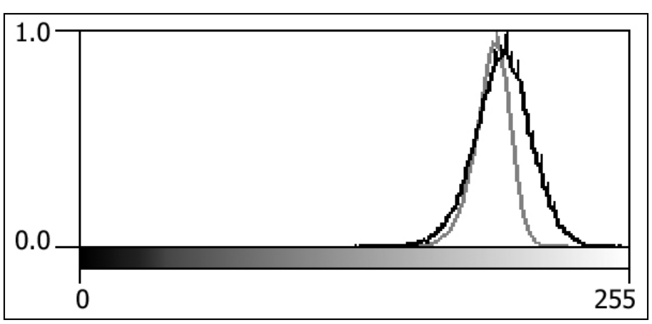

Fig. 3: Pixels closer to the missing regions are brighter and noisier than the rest. In grey, the histogram for the central half of the original image considering only those points more than 2 pixels far from missing regions (mean 191.73, standard deviation 9.27); in black, the histogram of the points less than 2 pixels far from missing regions (mean 196.74, standard deviation 14.13).

Other refinements that have been dropped after experiments showed them to be of little practical value are:

- Processing the other missing pixels in the grown region only if the standard deviation of the region's gray-level distribution is below a given threshold (proportional to the deviation of the initial backscatter map). This should have reduced overshooting in regions with abrupt changes in the underlying slope-map. Results show this mostly influences the backscatter image borders, which almost always have missing or very noisy slope values anyway.

- Processing the other missing pixels in the grown region only if they have at least one non-missing neighbor. This could be important to guarantee a gradual propagation of the gray-level distribution if:

  a) Holes in the backscatter are very big compared to the preset region size, which can be avoided by initially choosing a greater region size. Or,
  b) the ratio of valid pixels to total number of pixels in the region is very small, which can only happen if the initial region seed-point sits on a very steep or noisy slope region (because the seed is always adjacent to non-missing pixels). And again that almost every time corresponds to the image borders.

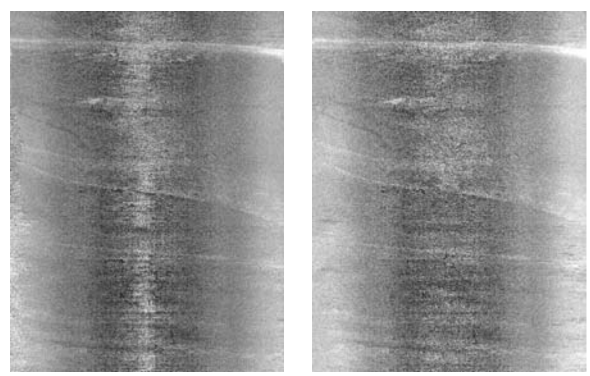

Fig. 4: Results of noise synthesis and effects of initial dilation stage. Left: without dilation; a brighter central band is still visible. Right: with initial dilation, which discards the pixels directly adjacent to missing regions.

III. HISTOGRAM-BASED INTENSITY COMPENSATION

Side-scan images frequently suffer from inadequate compensation of amplitude return intensity. As a consequence, sonar image mosaics usually present noticeable artifacts, which reduce the performance of subsequent processing and make their visual interpretation difficult.

The principal cause for the intensity variations is the attenuation of backscatter returns as distance and angle of incidence increase, which can be compensated for if the geometry of the acquisition process is known or can be recovered from navigation and bathymetry information.

A second source for intensity artifacts can be found in the sensors themselves, as they frequently apply an automatic compensation to the signal attenuation to minimize its effects. This correction is usually inaccurate, and—if the applied transformation is known— it's frequently reverted before any other processing is performed in the images. In most occasions, however, detailed information on the working parameters of the sensor is not available for the post-processing of legacy datasets, as is the case with the Simrad EM12s bathymetric side-scan sonar data used for testing the techniques described in this paper.

Although several methods have been proposed in the literature for the correction of intensity artifacts in backscatter maps (see for instance [6, 7]), they are mostly targeted to standard (non-bathymetric) side-scan sonar images, require some preliminary segmentation of the seabed or are sensitive to acquisition conditions. To overcome these limitations, we have developed a new histogram-based method for the removal of intensity variations in bathymetric side-scan images.

Our technique is based on the hypothesis that the gray-level distribution of the image should be the same for every detector in the sensor array—assuming the sensor configuration does not vary during acquisition. This, in turn, implies that the intensity correction is going to be the same for all the samples in a particular column of the backscatter image.

For the implementation, the average histogram of the whole image is first computed. Then the histograms of every column of the amplitude image (which roughly correspond to the returns sampled by one particular detector of the array) are fitted to the average histogram [8], thus removing the longitudinal banding artifacts present in the original image.

Fig. 5 shows an example of application of the proposed histogram correction to the results of the missing-ping restoration procedure shown in Fig. 4. The histogram correction can nevertheless be applied independently of the ping restoration, although care has to be taken to remove the influence of those columns that might have little or none valid pixels.

A final comparison of the results of the standard linear interpolation to those of the proposed restoration method, both texture-mapped over the bathymetry mesh, is presented in Fig. 6. The mapping of the original backscatter map with the missing pixels is also shown.

Fig. 5: Final result of the whole restoration process, after histogram-based intensity compensation to the output of the noise synthesis procedure with initial dilation (Fig. 4, right).

IV. CONCLUSIONS AND FUTURE WORK

The bathymetry mesh is mostly derived from ping phase information, while the backscatter image collects the amplitudes of ping returns. Standard interpolation methods can be used to patch holes in the mesh, but when applied to the backscatter map they generate poor results. The restoration method proposed in this article, however, integrates all the information provided by these two sources of data (phase and amplitude) to restore missing backscatter measures. It is therefore an effective way of using one source of data (bathymetry mesh) to complement a related one (backscatter map), and as such constitutes a proper fusion approach.

In the scope of project AMASON, we are currently reworking the techniques described in this article for the restoration of standard (non-bathymetric) side-scan systems, where bathymetry information is not directly available. Apart from better classification and feature extraction, the restored side-scan maps will improve the performance of the registration procedures used in AMASON's Concurrent Mapping and Localization methods for navigation enhancement [9].

ACKNOWLEDGMENTS

The authors would like to thank the Marine Geosciences Research Group of the University of Barcelona for providing the source EM12S survey data used in the elaboration of this paper.

This work has been partially supported by the 5th Framework Program of research of the European Community through the project AMASON (EVK3-CT-2001-00059).

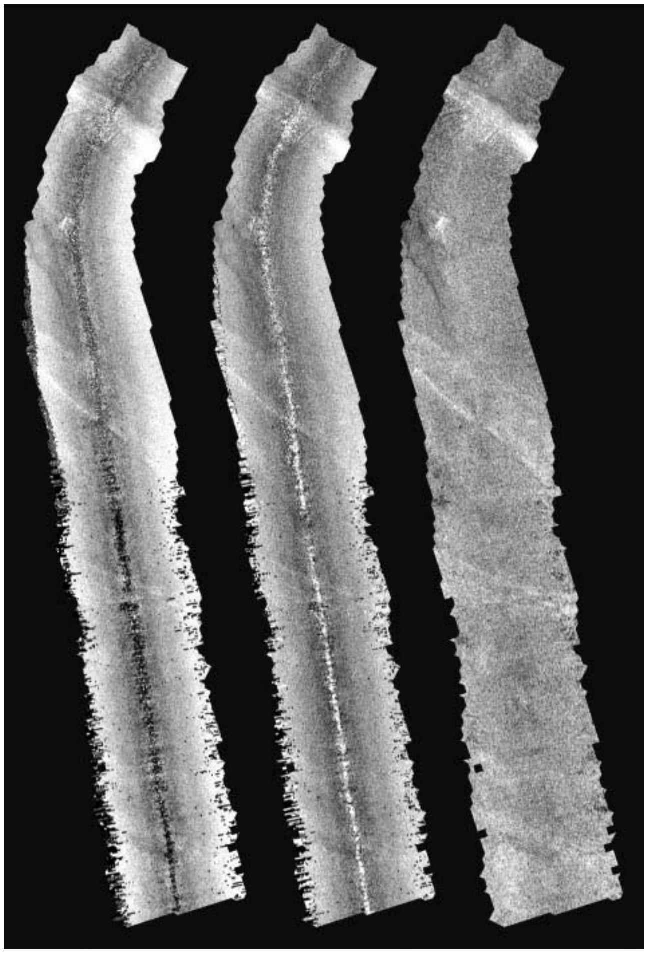

Fig. 6: Top view render of bathymetry mesh with different texture maps. From left to right: unprocessed backscatter image, restored backscatter by standard linear interpolation and restoration using the proposed method.

REFERENCES

[1] http://www.ece.eps.hw.ac.uk/~amason/

[2] N. C. Mitchell, "Processing and Analysis of Simrad Multi-beam Sonar Data," Marine Geophysical Researches 18 (pp. 729-739), 1996.

[3] A. A. Efros, T. K. Leung, "Texture synthesis by non-parametric sampling," IEEE International Conference on Computer Vision, Corfu, Greece, pp. 1033-1038, 1999.

[4] APL, "APL-UW High-Frequency Ocean Environmental Acoustic Models Handbook," Applied Physics Laboratory Technical Report, APL-UW TR9407, 1994.

[5] E. Coiras, J. Santamaria, C. Miravet, "A Segment-based Registration Technique for Visual-IR Images," Optical Engineering 39, no. 1, pp. 282-289, 2000.

[6] G. Shippey, A. Bolinder, R. Finndin, "Shade Correction of Side-scan Sonar Imagery by Histogram Transformation," IEEE OCEANS'94 Proceedings, Volume 2, 1994.

[7] Stuart Anstee, "Removal of Range-dependent Artifacts from Sidescan Sonar Imagery," DSTO-TN-0354, Aeronautical and Maritime Research Laboratory, Defence Science and Technology Organisation, Australia, 2001.

[8] R. C. Gonzalez, R. E. Woods, "Digital Image Processing," Addison-Wesley Publishing, 1992.

[9] I. Tena Ruiz, Y. Petillot, D. M. Lane, "Improved AUV Navigation using Side-scan Sonar," Proceedings of the OCEANS'03 Conference, San Diego, 2003.