Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Improving Automatic Target Recognition with Forward Looking Sonar Mosaics

F. Ferreira * **V. Djapic** ** **M. Micheli** *** **M. Caccia** ****

* *CNR-IEIIT, Via De Marini, 6, 16149 Genova, Italy, (e-mail: fausto.ferreira@ieiit.cnr.it). Work performed while at STO-CMRE.*
** *SPAWARSYSCEN Pacific, 53560 Hull Street, San Diego, CA 92152-5001 (e-mail: vladimir.djapic@navy.mil). Work performed while at STO-CMRE.*
*** *STO-CMRE, Viale San Bartolomeo, 400, 16126 La Spezia, Italy, (e-mail: micheli@cmre.nato.int).*
**** *CNR-ISSIA , Via De Marini, 6, 16149 Genova, Italy, (e-mail: max@ge.issia.cnr.it).*

Abstract:
Automatic Target Recognition (ATR) is a key element needed to make Mine Countermeasure missions using robots entirely autonomous. While there has been much progress in applying ATR algorithms on high-resolution Synthetic Aperture Sonar (SAS) and sidescan sonar data, performing ATR with a low cost Forward Looking Sonar (FLS) is much more challenging. An algorithm for the detection of underwater man-made objects in FLS previously developed can work in real-time although it suffers considerably from typical noise in sonar images and false alarms. The work presented here shows that ATR algorithms can be exercised on sonar mosaics built also in real-time instead of raw data coming from the FLS. The use of mosaics can help the detection of the targets by reducing some noise (including harmonics from other acoustic devices mounted on the robot) and giving a better contrast to the images to be processed. Moreover, mosaic images can be useful for post-processing and data analysis. The mosaicking algorithm also runs in real-time to maintain the performance of the system and to be useful in real missions. It was tested both on data previously collected and in real experiments with different set-ups and with different sonars. The wide range of results obtained with different surface vehicles and in different situations demonstrate the usefulness of the method.

Keywords: mosaicking, forward looking sonar, automatic target detection, acoustic image processing.

1. INTRODUCTION

Over the last years, sonar imagery has been used as an alternative to optical imagery due to its higher range and working conditions. In fact, sonars can be very useful under conditions of turbidity and lack of illumination, scenarios where optical cameras fail to provide good images. This is especially important either in vehicles that lack artificial light or in surface vehicles as the light attenuation in the water gives a very limited range to optical cameras.

While Side-Scan Sonars (SSS), Synthetic Aperture Sonars (SAS), and some Forward Looking Sonars (FLS) provide high resolution data, Forward Looking Sonars (FLS) with lower resolution are often used due to their satisfactory range resolution and lower cost. Moreover, they are very useful as their dimensions and power requirements allow them to be mounted on Remotely Operated Vehicles (ROVs), Autonoumous Underwater Vehicles (AUVs) and Autonoumous Surface Vehicles (ASVs) of medium size. SAS and SSS data can be used in large area surveys to give a first assessment of the area. After that initial survey, FLS can be used for a closer inspection to confirm the results and/or perform some kind of intervention task.

Concerning the high resolution data, since the work by (Hayes and Gough, 1992), several methods have been developed to detect man-made objects in SSS and SAS data, for example, (Reed et al., 2004; Beaujean et al., 2011). In particular, the target detection algorithm used in conjunction with the mosaicking algorithm presented herein is inspired by (Williams and Groen, 2011). All these methods detect a target in unknown, wide survey areas. The challenge of detecting targets while analyzing FLS data in real-time remains. Although there has been work on detecting obstacles using FLS data like (Martin et al., 2000; Petillot et al., 2001) and on registering FLS images Aykin and Negahdaripour (2013), most of the research is

⋆ STO-CMRE is the NATO Science and Technology Organization, Centre for Maritime Research and Experimentation. This work was performed under the 2013 CMRE Programme of Work, with funding from NATO Allied Command Transformation, and supported by the Centre's Visiting Researcher Programme.
⋆⋆Research supported in part by the Fundação para a Ciência e Tecnologia (FCT), Portugal with the PhD Grant SFRH/BD/72024/2010 and by the project MORPH, Contract Number: EU-FP7-ICT-288704 coordinated by ATLAS ELEKTRONIK GMBH.

not focused on detecting and identifying man-made objects in real-time. Recently, an algorithm working in real-time was presented but it focus only on obstacle detection, not specifically target recognition (Karabchevsky, 2011). (Fallon et al., 2013) presents a SLAM-based algorithm to reacquire features (targets) but does not recognize them. The work presented in (Galceran et al., 2012) deals precisely with real time detection and identification of man-made objects. It does not need any previous training but it integrates knowledge about the target to be searched, particularly, the shape and dimensions of the object. Since satisfactory results were obtained with that work, the idea presented here is to use the mosaicked data instead of the raw frames to further improve the results.

Taking into account previous work on real-time mosaicking of optical camera data Ferreira et al. (2012), a novel algorithm for mosaicking of FLS data in real time is hereby presented. Not all the assumptions assumed for the optical camera data are valid in the acoustic domain and the task is much more challenging. However, as the results will demonstrate, the algorithm works satisfactory for the applications tested. It will be shown that the mosaics produced for the FLS data can improve target detection and recognition and diminish the false alarms rate without compromising the real-time capability. It can also be useful for post-processing and data analysis and examples will be presented in Section 4. This is not the only mosaicking algorithm that exists for the FLS data. In (Hurtos et al., 2012), a phase correlation-based mosaicking algorithm was applied to FLS data with the same kind of data that was used on this work but it was not concentrating on the real-time aspect neither in ATR. More recently in (Hurtos et al., 2013a) an evolution of it was applied to chain inspection with good results although in this case a high resolution FLS sonar (DIDSON) was used with a very limited range. In (Yong, 2011), mosaicking techniques working close to real-time were investigated for FLS data in a ship-hull inspection application but not applied to ATR. To the best of the authors' knowledge using mosaicked data for real-time ATR is new in the state of the art.

The remaining of this paper is organized as it follows. Section 2 presents the mosaicking algorithm. Section 3 discusses briefly the detection algorithm and the connection between the two algorithms. Results obtained with real data are presented in Section 4. Finally, conclusions and future work are described in Section 5.

2. MOSAICKING ALGORITHM

One of the main reasons to use mosaics instead of raw data is the reducing of noise. In real experiments, the authors verified that a regular pattern of noise was appearing in the raw data collected by the FLS. Taking into account that there is an echosounder working at 200KHz onboard the vehicle used in the experiments and that the FLS is not detecting only echos at its center frequency (either 450KHz or 900KHz) but instead has a bandwith ranging from 300 to 600Khz (or 600 to 1200Khz respectively), it is easy to conclude that that regular pattern comes from the 2nd harmonic of the echosounder. Figure 1 shows an example where this pattern appears. The brighter blobs are the echosounder interference and the target is the rectangular shape with a clear shadow in the middle of the image. When this pattern falls into the window of interest defined, then a false alarm in target recognition can occur since normally the echo coming from the echosounder is stronger (and thus brighter) than the one coming from the farther bottom target. A way of eliminating most of these false positives is to use mosaics as these ones are a collection of several raw frames and thus average the scene viewed resulting that this pattern will appear less often. With mosaics, not only the false positives are eliminated but also the natural changes in insonification from frame to frame and the noise coming from waves or water column particles reflections are also removed or diminished considerably.

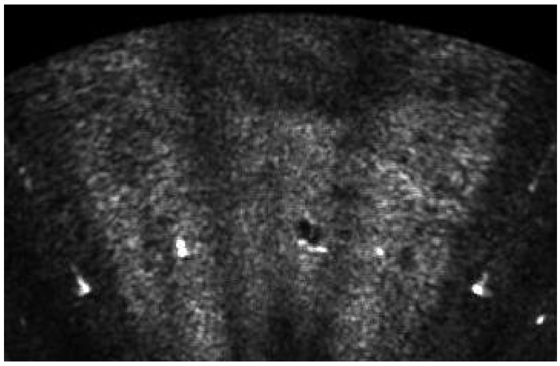

Fig. 1. Raw data depicting a target in the middle and echosounder interference to the left and right of the target.

The mosaicking algorithm is based on the one presented in Ferreira et al. (2012) and updated recently in Ferreira et al. (2013a). In this work, we apply our mosaicking algorithm to FLS data. While in the case of mosaicking optical data, the algorithm was purely based on optical data (with acoustic altimeters only used in case of failure), here, other sensors onboard the ASV can be used together with the acoustic data. Motion estimates are available from a Differential GPS (DGPS) system for the ASV. In fact, the vision-based motion estimation algorithm that was used in (Ferreira et al., 2013a) is redundant in this case. In the same way, instead of using a magnetic compass, heading information is available from the same DGPS system. The DGPS system is not mounted coincidentally with the sonar head but there is a compensation and the exact GPS position of the sonar head is available as an input to the mosaicking algorithm.

The sonar head is mounted in a pole that can go up and down and has a pan-and-tilt unit. While the pole only influences in the detection of the target and calculation of the distance to the bottom target, the pan-and-tilt unit has to be considered for mosaicking purposes. As in many works in the literature (Hurtos et al., 2012; Aykin and Negahdaripour, 2013), the 3D point is projected onto 2D using an orthographic projection as approximation. This approximation is valid as long as the scene relief is despicable when compared to the sonar range. This is true if the elevation angle (ϕ) is small as $cos(\phi) \approx 1$ and $sin(\phi) \approx 0$ for small angles. Typically, in our target recognition missions, the elevation angle is smaller than $5°$ thus, not violating that assumption. A study on the effect of this approximation in the images registration can be found in (Johannsson et al., 2010). Equation 1 defines the 3D coordinates in spherical coordinates (r, θ, ϕ) with r being the sonar's range and θ and ϕ the bearing and elevation angles respectively. Equation 2 presents the projected 2D point.

$$P = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} r\cos\theta\cos\phi \\ r\sin\theta\cos\phi \\ r\sin\phi \end{bmatrix} \tag{1}$$

$$\hat{p} = \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} r\cos\theta \\ r\sin\theta \end{bmatrix} \tag{2}$$

Given the orthographic projected point, an affine transformation can be used to relate two consecutive frames. The rotation matrix and translation parameters for the affine transformation can be directly obtained from the DGPS and the pan-and-tilt unit with the pan being treated as a pure rotation around the elevation axis. Skipping the details, the final relation between two frames can be described as in equation 3. For convenience, rewriting Equation 3 using homogeneous coordinates for the projected point $(u, v, 1)$ gives us the final formula in Equation 4 with $\mathbf{R}$ being the rotation matrix and $\mathbf{t}$ a vector with the translations in x and y directions.

$$\begin{bmatrix} u_2 \\ v_2 \end{bmatrix} = \mathbf{R} \begin{bmatrix} u_1 \\ v_1 \end{bmatrix} + \mathbf{t} \tag{3}$$

$$\begin{bmatrix} u_2 \\ v_2 \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ 0_{1x2} & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ v_1 \\ 1 \end{bmatrix} \tag{4}$$

As in (Ferreira et al., 2012), the motion estimates (now coming from DGPS instead of the camera) can be used as initial guess of the point where the the images should be stitched together to form the mosaic. Then, a correlation method can be applied in a neighborhood of that initial guess instead of the full frame to improve computational efficiency and take advantage of the DGPS information. This refinement step is particularly important when mosaicking large areas while for small areas such as the ones present in some of the Mine Countermeasures missions, this is not so critical as the robot moves slowly when compared to the frame rate obtained. For instance, in circular missions around a target or in missions where the robot keeps a fixed distance to the target, the DGPS estimates can be enough and actually work better as the environment is practically featureless (with the exception of the target). However, filtering of the DGPS estimates is needed to avoid possible jumps in the GPS position. Both maximum speed and acceleration of the vehicle are taken into account in this filter.

When the images are put together to form the mosaic special care has to be taken. Most of the pixels close to the origin of the sonar head should be discarded as they do not represent accurately the bottom but include instead reflections along the water column. Only a valid region of interest should be mosaicked. When the goal is to mosaic quasi-static setups or setups where a moving target is actually close to the sonar head, then the full area should be used. The blending strategy can not be the same as the one used for optical cameras. In that case, only new pixels were added to regions that were empty. For FLS data that approach does not apply. First, in quasi-static setups where there is a moving target, no update would occur and the history of the mission would be lost. Second, in circular missions, that would mean adding only pixels in the boundaries and not reducing any noise in the overlapped areas. In the lawn mowing patterns that strategy can be used but it is still suboptimal as does not improve areas that are seen in more than one frame. Due to the changes in insonification, an object can be clearly seen in one frame and then not anymore in the next one. If the blending technique would add only new pixels in a case where the object was not seen on the previous frame, then that object could not be seen at all. Hence, a different approach is here used. The overlapping area between the mosaic and the frame to be added is computed and then only on the overlapping area, an average of the mosaic and the new frame is copied to the mosaic. This ensures that some of the noise is filtered and gives more importance to the new frame than what it was done with the optical camera. This simple technique works well for the purpose devised. Other more complicated techniques like the ones described in (Yong, 2011) or recently in (Hurtos et al., 2013b) including different weights for each of the images or for each region of the image can be used in a non real-time context but in this case, the authors chose the simplest yet effective solution.

2.1 Implementation issues

The mosaicking algorithm described was implemented taking into account the real-time constraint and the interface with the detection algorithm. The MOOS framework (Newman, 2007) used onboard the CMRE's vehicles was chosen and the mosaicking code is fully integrated with the code running onboard the vehicles. It is highly configurable through the use of a simple configurable file without the need of recompiling the code. The need of a high degree of flexibility is explained by the different applications and missions in which the algorithm is tested. It can be used for at least two different kinds of Forward Looking Sonars with a different interface, for large scale areas, for a limited defined area (by starting and ending point), in circular missions with regular resetting and publishing of the mosaic and even for static set-ups. Examples of these different applications will be shown in the Results section.

3. TARGET RECOGNITION ALGORITHM

The target detection and recognition algorithm is the one described in (Galceran et al., 2012) with several improvements. It is a real-time detection algorithm for FLS raw data that makes the use of integral image representation to achieve the real-time capability. Targets are detected by comparing the echo map of a region of interest and the background map of the same area. Only pixels that have an echo a certain amount higher than the background (threshold configured) are considered as possible targets. The resulting blobs after thresholding are morphologically analysed taking advantage of prior knowledge about the kind of objects that are being looked for. Then, a minimum echo threshold is used to filter out lower intensity blobs. The target expected location is given to the target detection algorithm by the results of a survey using SAS. In the case of a cylinder, its orientation can be estimated from the SAS data and provided as an input to the detection algorithm. Only location and orientation are provided by the SAS survey not the full map. Having the orientation of the target, one can search for the cylinder only when its broadside is visible. This is especially important in the case of circular missions around the target as its shape changes with the relative bearing to the sonar. The maximum deviation from ideal case where the cylinder is perfectly horizontal on the sonar frame is parameter configurable and it ranged from 30° to 50° in the experiments. The final decision of choosing the target location from the several detections produced was also improved. The detections are grouped and the distance from the centroid of each group to the expected target location is used. In this way, the closer a cluster of detections is to the target, the most probable is to be chosen. The Euclidean distance between the shape of each detection and the ideal shape is calculated and averaged. Then, the score obtained for each cluster is weighted with the inverse of mean distance. In this way, a cluster that has less detections but whose detections are closer to the real object increases its chances of being chosen. Due to the modularity of the MOOS framework very few changes were needed to connect the output of the mosaicking algorithm with this algorithm. The target detection can be enabled by the mosaicking algorithm when this produces and publishes a new mosaic. After performing a detection, the target detection algorithm will wait for a new mosaic to analyse. The detections are all saved in a list and the final decision is taken after either a certain amount of time, number of mosaics analysed or by request. How often the mosaicking algorithm publishes a mosaic is configurable and has to do with the type of mission.

4. EXPERIMENTAL RESULTS

The results here presented are a small subset of the full results obtained. The range of results selected to be shown try to represent the flexibility of the algorithm, its possible applications and usefulness. For a more complete insight on the full results including mosaics obtained with another kind of sonar (Reson Seabat) and a ROV without access to DPGS, please refer to (Ferreira et al., 2013b). Therefore, the results show not only the improvement in ATR but also some other capabilities of the mosaicking algorithm. All the results here shown were obtained with BlueView sonars mounted either on a fixed pole or onboard the ASVs Gemellina and Gulliver from CMRE. The algorithm can run at a frame rate between 1Hz and 4Hz depending only on the number of mosaics saved for post-processing and data analysis.

4.1 Improving target contrast

Before going through the results for the target recognition it is useful to show a comparison between the kind of raw data one can expect from the sensor and the mosaic produced after the collection of some frames. The data here presented was obtained with a BlueView P450-130 with a Field of View (FOV) of 130° and working at 450KHz and the sonar was mounted onboard the ASV Gemellina. The data was collected in Marciana Marina, a marina in the Island of Elba, during the ANT'11 trial conducted by NATO Undersea Research Centre (NURC), currently CMRE. Figure 2 shows partial views of both the raw data and the mosaic. As it can be easily seen, the mosaic improves the contrast between a possible target and the background and gives a better definition of the shadow. It also allows the user to see details hardly seen on the raw data such as the long mooring rope on the right of the possible target. While in the raw data it is hard to understand that there is a mooring rope present on the scene due to the several interruptions in the image, in the mosaic it is very clear and thus identifiable.

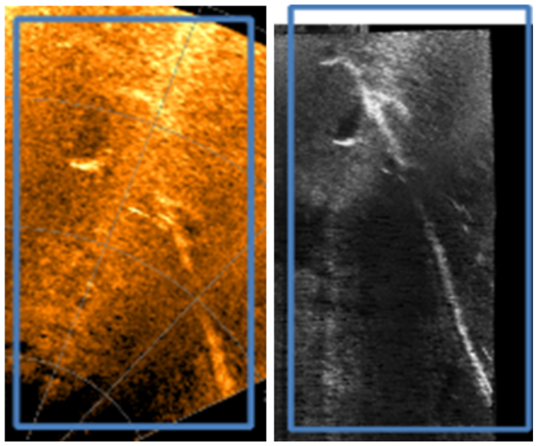

Fig. 2. On the left, a detail of a raw frame. On the right, the same area represented on the mosaic.

4.2 Improving target recognition

The primarily use of the mosaicking algorithm is the automatic target recognition. Therefore, preliminary results from Multinational AutoNomous Experiment (MANEX'13) are hereby presented as that was the best and most recent example of application of the algorithm. MANEX'13 was conducted in late October 2013 off the coast of the Island of Elba. The sonar is the same as in the previous section but the ASV is the new Gulliver, a catamaran also from CMRE which is slightly bigger than Gemellina. Typical missions executed for target reacquisition are: circling around the bottom target or keeping a distance to that target. An initial information of the position of the target is given by an AUV with an onboard SAS system as in a typical mission scenario. In this case, results are shown for a circular mission. Figure 3 represents a good detection of the target even in hard conditions such as the presence of the echosounder interference. There are two visible echosounder blobs, one on the right close to the target and one farther on the left. Comparing with Figure 1, the improvement is clear as most of the blobs are filtered out in the mosaic. It is not possible to show the same detection for the raw data because the algorithm runs either with one or the other. To assess the improvement that mosaicking can bring to the target detection and recognition algorithm, quantitative results were obtained in terms of false positives and true positives detection rates both for mosaic and raw data and are presented in Table 1. Percentages of true and false detections are calculated with respect to the total number of frames analysed. In this case, the mosaicking algorithm was publishing a mosaic to the detection algorithm each 20 frames which means that the number of frames analysed is 20 times less than in the case of raw data being used. There is a tradeoff in how much one is able to wait for a detection and the area covered by the mosaic. For instance, another criteria used for publishing mosaics is a maximum heading span of the mosaic. If a mosaic is published each 30° of maximum heading difference between the last frame and the initial frame, the number of mosaics produced is 3 times smaller the one obtained with a 10° of maximum heading difference. In the case of circular missions, the object shadow and echo change with the sonar's angle of incidence. Therefore, it is mandatory to produce a mosaic after a defined maximum heading difference or a maximum number of frames in order not to have considerable changes of the object shape or else artifacts can deteriorate the performance.

The results of Table 1 show that using mosaicking one can get a higher ratio of true detections over false positives. The percentage of false positives is slightly smaller for the mosaic data but the false positives ratio decreases considerably due to the elimination of noise coming from echosounder harmonics mainly. The number of detections is small as mosaics are published only each 20 frames.

Table 1. Correct and False Detections [%], their Ratio and Total correct detections

	Correct [%]	False [%]	Ratio	Total correct
raw data	40.2%	23.4%	1.71	255
mosaic data	38.4%	15.3%	2.5	15

Therefore, an analysis of the influence of the publishing rate parameter was performed. Table 2 shows the different true detections and false positive rates and the ratio between them for different publishing rates, namely publishing a mosaic each 5, 10 and 20 frames, and the raw data results for the same dataset.

Table 2. Correct and False Detections [%], their Ratio and Total correct detections for different publishing rates

	Correct [%]	False [%]	Ratio	Total correct
raw data	40.2%	23.4%	1.71	255
mosaic data (20)	38.4%	15.3%	2.5	15
mosaic data (10)	29.5%	11.2%	2.63	21
mosaic data (5)	39.7%	14.6%	2.71	60

As it can be seen, the ratio between correct detections and false positives increases with the rate of publishing mosaics. The number of detections increases also to 60 which is as much as one can get with the raw data proportionally. The best case for the true detections is when a mosaic is published each 5 frames and although that does not correspond to the minimum false positives, it still beats the false positives percentage in the raw data (14.6% against 23.4%). As expected, the mosaicking algorithm can improve the ATR by decreasing the noise level and eliminating several false positives without compromising the true positives detection rate. The algorithm can run between around $1Hz$ and $4Hz$ depending on the quantity of data (mosaics) saved for post-processing and data analysis. The real-time capability is important because the results from the ATR are used in real-time and in the loop to start the next phase of the mission, namely, sending a small Unmanned Underwater Vehicle (UUV) to the identified target.

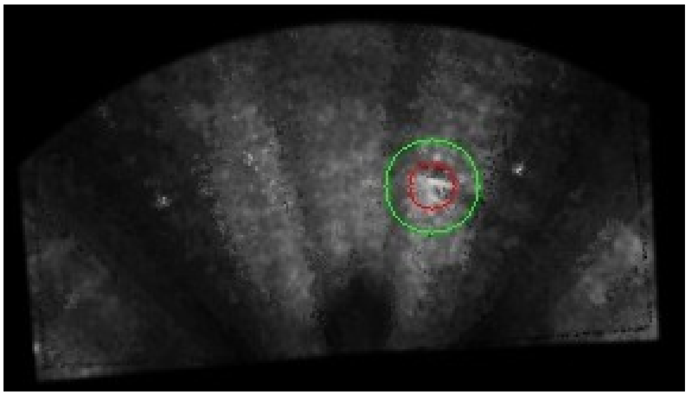

Fig. 3. An example of a correct detection in a mosaic even in the presence of noise.

4.3 Post-processing and data analysis

Finally, it is also important to show the usefulness of the mosaics in the post-processing and data analysis phase. The data here presented was collected using a BlueView P450-45 FLS with a 45° FOV and working at 450KHz during the Breaking the Surface 2014, the 5th International Interdisciplinary Field Training of Marine Robotics and Applications. The sonar was mounted on a pole fixed to the pier but the same results apply to the case where the sonar is mounted in a ASV. In this particular setup, besides a target on the sea bottom, there is a small UUV moving towards the bottom target while the ASV is stopped. Mosaicking this quasi-static scene (only the UUV is moving) can give the history of the mission in a very simple way without the need of replaying the whole sonar file. Details that can not be seen in the raw data (or are hardly noticeable), are easily detected looking to the produced mosaics. Namely, the case where the UUV misses the target and goes forward over the shadow of the bottom target. In the raw data, this is a very difficult situation to detect because the shadow prevails while in the mosaic it is very clear where is the UUV during the whole mission. Figure 4 represents this situation where the UUV missed the target and ended up his mission only after it in the shadow region of the bottom target.

5. CONCLUSIONS

The wide range of results show the effectiveness of the method. Several conclusions can be drawn. at best 2.71/1.71 often more than 10% First, mosaicking of FLS data can be done in real-time with satisfactory results. The noise present in the raw data is filtered, the false positives percentage is decreased (about 10%) and the ratio between true positives and false positives is increased at best 58%. Second, in featureless environments, image registration does not work properly and either GPS data should be used or other techniques to register the images, namely in the frequency domain should be investigated. Relying on the navigation info from the autonomous robot and using only positional data to estimate the displacement between frames can give useful mosaics for the purpose of Automatic Target Recognition. As it was shown, mosaics can be used not only in real-time operation but also for post-processing and data analysis in a much easier way

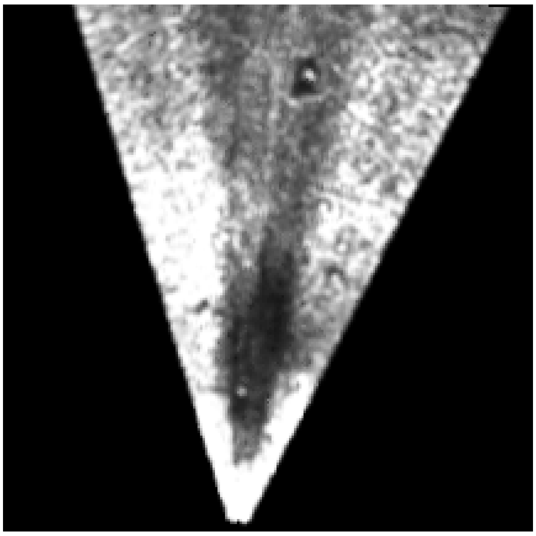

Fig. 4. Mosaic of a quasi-static setup with a bottom target and an UUV in the shadow of the bottom target.

than replaying the whole sonar file and looking at all the frames.

As future work, a method of registering the images in featureless environments should be found without compromising the real time capability. On the side of the automatic target recognition algorithm, the morphologic filter should be improved to detect targets with a large range of orientations. Currently, using the ratio of the principal axis versus the smaller axis did not work properly due to the difficulty of estimating the smaller axis in irregular blobs such as the ones detected by the BlueView sonar. It is worthwhile to test it first on data obtained by a higher resolution sonar where the theoretical assumptions about the shape of the object are closer to the real detections and then try to transpose the results to the BlueView data.

ACKNOWLEDGEMENTS

The authors wish to thank Professor Shahriar Negahdaripour (from University of Miami) for the fruitful discussions and Alberto Grati and Dr. Stefano Fioravanti (from CMRE) for their effort during the testing phase. In addition, the authors thank Dr. Warren Fox, ANMCM Programme Manager at CMRE, for his comments during the internal paper review process.

REFERENCES

Aykin, M.D. and Negahdaripour, S. (2013). On feature matching and image registration for two-dimensional forward-scan sonar imaging. *J. Field Robotics*, 30(4), 602–623.

Beaujean, P.P.J., Brisson, L.N., and Negahdaripour, S. (2011). High-resolution imaging sonar and video technologies for detection and classification of underwater munitions. *Marine Technology Society Journal*, 45(6), 62–74. doi:doi:10.4031/MTSJ.45.6.6.

Fallon, M., Folkesson, J., McClelland, H., and Leonard, J. (2013). Relocating underwater features autonomously using sonar-based slam. *Oceanic Engineering, IEEE Journal of*, 38(3), 500–513. doi: 10.1109/JOE.2012.2235664.

Ferreira, F., Veruggio, G., Caccia, M., Zereik, E., and Bruzzone, G. (2013a). A real-time mosaicking algorithm using binary features for ROVs. In *Control Automation (MED), 2013 21st Mediterranean Conference on*, 1267–1273. doi:10.1109/MED.2013.6608882.

Ferreira, F., Djapic, V., and Caccia, M. (2013b). Real time Mosaicking of Forward Looking Sonar Imagery. In *Journal of Field Robotics*. Submitted.

Ferreira, F., Veruggio, G., Caccia, M., and Bruzzone, G. (2012). Real-time optical SLAM-based mosaicking for unmanned underwater vehicles. *Intelligent Service Robotics*, 5(1), 55–71. URL http://dx.doi.org/10.1007/s11370-011-0103-x.

Galceran, E., Djapic, V., Carreras, M., and Williams, D.P. (2012). A real-time underwater object detection algorithm for multi-beam forward looking sonar. In F.L. Pereira (ed.), *Navigation, Guidance and Control of Underwater Vehicles (NGCUV) 2012*, volume 3, 306–311. Porto.

Hayes, M. and Gough, P. (1992). Broad-band synthetic aperture sonar. *IEEE Journal of Oceanic Engineering*, 17(1), 80–94.

Hurtos, N., Cufi, X., Petillot, Y., and Salvi, J. (2012). Fourier-based registrations for two-dimensional forward-looking sonar image mosaicing. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, 5298–5305. doi:10.1109/IROS.2012.6385813.

Hurtos, N., Palomeras, N., and Salvi, J. (2013a). Autonomous detection and following of a mooring chain using a forward-looking sonar. In *OCEANS - Bergen, 2013 MTS/IEEE*.

Hurtos, N., Cufi, X., and Salvi, J. (2013b). A novel blending technique for two-dimensional forward-looking sonar. In *OCEANS - San Diego, 2013 MTS/IEEE*.

Johannsson, H., Kaess, M., Englot, B., Hover, F., and Leonard, J.J. (2010). Imaging sonar-aided navigation for autonomous underwater harbor surveillance. In *IROS*, 4396–4403. IEEE.

Karabchevsky, S. (2011). *Real-Time Underwater Obstacle Detection using Forward Looking Sonar and FPGA*. Master's thesis, Ben-Gurion University of the Negev.

Martin, A., An, E., Nelson, K., , and Smith, S. (2000). Obstacle detection by a forward looking sonar integrated in an autonomous underwater vehicle. In *OCEANS 2000 MTS/IEEE Conference*, volume 1, 337–341. doi: doi:10.1109/OCEANS.2000.881281.

Newman, P. (2007). Introduction to programming with MOOS.

Petillot, Y., Ruiz, I., and Lane, D. (2001). Underwater vehicle obstacle avoidance and path planning using a muli-beam forward looking sonar. *IEEE Journal of Oceanic Engineering*, 26.

Reed, S., Petillot, Y., and Bell, J. (2004). Automated approach to classification of mine-like objects in side-scan sonar using highlight and shadow information. In *IEEE Radar, Sonar and Navigation*, volume 151, 4856.

Williams, D.P. and Groen, J. (2011). A fast physics-based, environmentally adaptive underwater object detection algorithm. In *OCEANS - Spain, 2011*, 1–7.

Yong, E.W. (2011). *Investigation of Mosaicing Techniques for Forward Looking Sonar*. Master's thesis, Heriot-Watt University,University of Girona,Universite de Bourgogne.