3-D Mapping for the Characterization of the Seafloor using an Underwater Robot within the MUMAP Project

Rafael Garcia*, Nuno Gracias*, Ricard Campos*, Ricard Prados*, Laszlo Neumann*,
Tudor Nicosevici*, Javier Escartin**
***University of Girona, Spain,** rafael.garcia@udg.edu, ngracias@eia.udg.edu, rcampos@eia.udg.edu, rprados@silver.udg.edu, lneumann@eia.udg.edu, tudor@eia.udg.edu,
**** Institut de Physique du Globe de Paris, France,** escartin@ipgp.fr

Abstract

The study of marine habitats is a key component for understanding and managing ocean resources. In addition to its intrinsic scientific interest, underwater studies can be used for assessing sites that are important from a conservation point of view, determining areas sensitive to disturbance and pollution, development of coastal and marine protected areas management plans or evaluation of fundamental environmental issues relating to the exploitation of deep-sea mineral and energy resources. In consequence, such studies provide valuable indicators for environment impact assessment and monitoring of environmental changes.

Recent technical advances in the field of robotics provide scientists with the right tools to acquire, process, and merge vast amounts of ocean data, making possible the generation of quantitative and qualitative information on marine habitats. However, while several processing tools are available in the market for analysis and interactive 3D visualization of multi-beam sonar data, limitations in image processing techniques and data ingestion to date have greatly restricted the exploitation of optical data to its full potential. Typical imaging surveys gather tens of thousands of images and hours of video that are only inspected visually by end-users but are not put in a geographic context.

This paper presents the results of the MUMAP project, aiming at the development of high-resolution seafloor mapping systems for optically monitoring relatively large areas of the ocean floor using an underwater robot. These high-resolution maps have direct applications to environmental studies, oceanography, geology and biology, as well as for the offshore industry.

Keywords: underwater robots, seafloor characterization, mapping, offshore.

1. Introduction

We are nowadays witnessing an increasing interest in the study of the oceans. The urge to find alternative food and energy sources and to understand climate changes, as well as biological and geological phenomena has determined scientists to multiply their efforts toward understanding the complex underwater environment. An example of such is the rapid decline of some habitats (for example coral reef communities worldwide), which is stimulating the development of innovative assessment tools to rapidly and effectively document the distribution and condition of endangered benthic organisms (Solan et al., 2003; Fisher et al., 2007).

However, direct interaction between scientists and the underwater environment is generally not feasible, due to its inaccessibility and hostility. An emerging alternative is remote data acquisition using underwater robots and the use of advanced processing techniques to digest the huge amount of data collected.

With the development of better and more affordable photography and videography equipment, the use of digital imagery in benthic monitoring has increased dramatically in the last decade. Video surveys are now routinely conducted as complements to diver-based measurements (Ninio et al., 2003). Moreover, several large-scale monitoring programs are now based almost exclusively on the analysis of video imagery. One such example is the Coral Reef Monitoring Program of the Florida Reef Tract where permanent belt transects are surveyed annually and video frames are sub-sampled to obtained estimates of coral cover and condition (Porter et al., 2002).

This paper describes the outcomes of the MUMAP project, which proposed the development of a new seafloor mapping system for large areas of the ocean floor, with direct applications to oceanography, geology, biology, engineering, underwater archaeology, and environmental studies, in addition to applicability for offshore industry (oil, gas, CO_2 offshore storage). Consider the following scenario, where we aim to monitor the evolution of a given area of the ocean floor, with the aid of an Autonomous Underwater Vehicle (AUV). As a first step, the AUV is deployed from the mother ship to perform a near-bottom, high-resolution geophysical survey (e.g., bathymetry, imagery, magnetics), while simultaneously gathering additional data from the seafloor and water column (turbidity, temperature, salinity, etc.). At the same time, an acoustic positioning system linking the ship and the AUV (Ultra Short Base Line – USBL) provides geo-referenced navigation data as well as communication with the end user, albeit with a resolution in position well below that of the geophysical data acquired. During this phase of the mission, the AUV will follow programmed paths while tracking the terrain to maintain optimal navigation, and to optimize survey coverage and data quality (Ridao et al., 2013). The ship's navigation will be coordinated with that of the AUV for precise positioning and reliable acoustic communications. Once the survey is completed, the AUV is recovered and the acquired data is retrieved. The data is then used to build a map of the interest area. It should be noted that near-bottom surveys provide unprecedented resolution of the seafloor structure, either from multi-beam acoustic systems or from imagery (video and electronic still images), at resolutions varying from tens of centimeters for acoustic systems to a few millimeters for optical systems (Gracias et al., 2013).

Although several multi-beam processing commercial tools for analysis and interactive 3D visualization are currently available (Fledermaus, iView4D, CUBE, etc.), limitations in image processing techniques and data ingestion have greatly restricted the exploitation of optical data to its full potential. Typical surveys will gather tens of thousands of images and hours of video. Currently, such data is largely inspected manually, resulting in a great deal of time being spent by scientist in laborious tasks related to image by image analysis for identifying regions of interest, visual and structural differences for evolution assessment, etc. Furthermore, data acquired using different sensors such as sonars, vision, or water property measurement devices, has to be either separately analyzed or manually aligned. However, modern multi-disciplinary, complex underwater studies clearly lead to the necessity of automated, multi-modal spatial representations for data analysis.

2. Acquiring images of the seafloor using an underwater robot

A large number of underwater studies rely on visual inspection, frequently using optical mosaics when the area to be explored is larger than what can be covered by a single image. Here, images of the region of interest are acquired by underwater cameras deployed by either remotely-operated vehicles (ROVs) or autonomous underwater robots (AUVs). Due to the high absorption rate of the underwater medium that limits the amount of light present, high sensitivity cameras are required for such surveys. Moreover, as a result of the light scattering determined by the particles suspended in water, dehazing methods are required to increase the effective range of the cameras (Gracias et al. 2013). In cases where the surveys take place at depths where natural light is not present or not sufficient, artificial lighting systems need to be employed.

As a rule of thumb, visual surveys should be conducted at altitudes not exceeding 8-10m, while each still image or video frame should have an overlap of at least 40% with the previous one in order to ensure consistent mapping results. For this, it is crucial to have a close coordination between the imaging system and the navigation of the robot.

3. High-resolution 2D Mapping

The acquired images are later processed and stitched to form high-resolution, detailed two-dimensional (2D) maps, also known as mosaics. Image mosaicing deals with the process of combining the information from multiple images of the same area, to create a single representation with extended field of view (Elibol et al., 2014a). In order to construct ocean floor photo-mosaics, the individual images forming the mosaic are usually obtained by setting a camera on an underwater vehicle (Gracias et al., 2013; Elibol et al., 2013) – usually in a down-looking orientation.

During the mosaicing process, the systems looks for a series of salient points, which are then matched in corresponding (adjacent) images. The tracking of such interest points across multiple images allows motion estimation (Elibol et al., 2010). Using the motion estimations, new non-consecutive overlapping images (or *loops*) can be detected (Nicosevici and Garcia, 2012), then the image topology is computed, and next the images are globally aligned and merged, forming composite view of the surveyed area (Elibol et al., 2011; Elibol et al., 2012).

However, the underwater medium adds particular challenges to the image acquisition task, which are not present in land imagery (Garcia and Gracias, 2011a). Phenomena such as light attenuation and scattering, force the image acquisition to be performed as close to the seabed as possible, using artificial light sources with limited power (often due to vehicle autonomy reasons in the case of autonomous robots). These phenomena cause uneven illumination across each image and the presence of high-frequency noise. Furthermore, the images acquired are used to build a two-dimensional representation of the seafloor, while this may contain salient tridimensional structures. Thus, when several frames are stitched together to build a photo-mosaic, the seams along image boundaries are often noticeable, due to both photometrical and geometrical registration inaccuracies (Prados et al., 2011). Image blending is the merging step in which those artifacts are addressed (see Figure 1). In a first step, lighting artifacts such as refracted sunlight (Shihavuddin et al., 2012) or non-uniform illumination of the images are compensated. Next, the optimal seam over the overlapping image regions is found, minimizing the photometric differences along its path (Prados et al., 2012). Finally, the overlapping image information around the seams is fused using gradient domain techniques. This pipeline ensures a continuous a consistent appearance of the built mosaic while reducing visual artifacts (Prados et al., 2014).

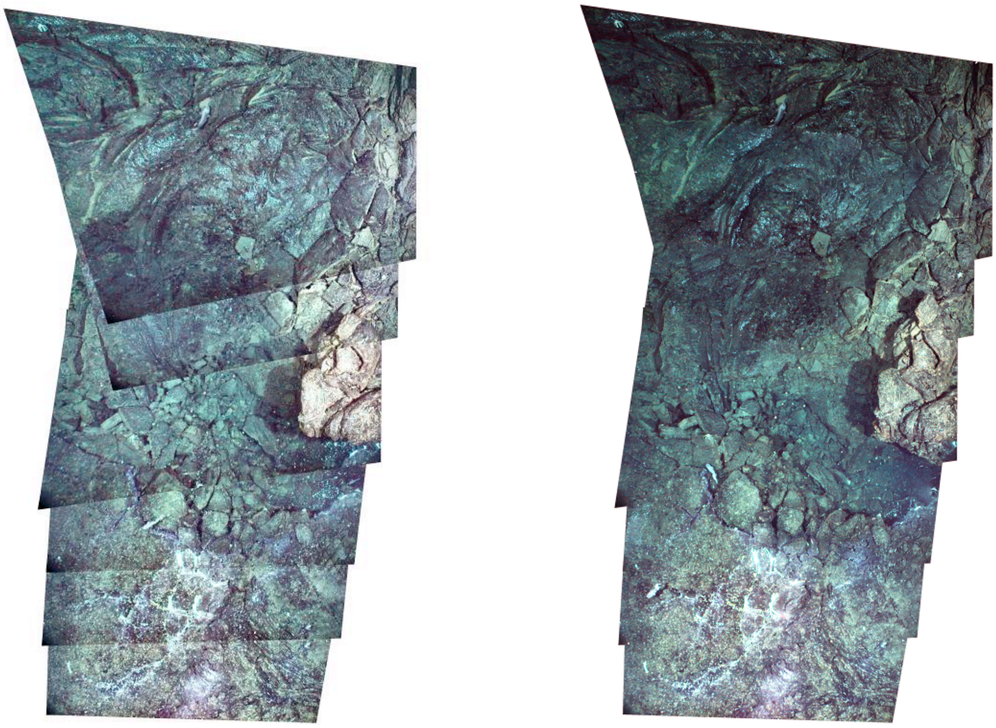

Fig. 1. Colour photo-mosaic of an interest region before (left) and after (right) the application of an image blending technique in order to reduce visual artefacts.

Figure 2 illustrates the end result of the mosaicing process using the techniques developed in the MUMAP project, representing a high-accuracy 2D map of a seafloor patch of nearly 1 square kilometer. This process is now largely automated, and can be systematically applied to optical surveys (in this case with a resolution of 10mm/pixel). Moreover, we have extended current mapping capabilities to combine a large 2D mosaic with a bathymetry of the area acquired with multibeam systems but at a coarser resolution (see Figure 3). In addition, the approach has been also extended to create high-resolution 3D representations of regions of interest (see Section 4), using either a structure-from-motion monocular camera or with a stereo optical system (Ferrer and Garcia 2011), and achieving a much higher resolution when compared to multibeam bathymetry. These georeferenced 3D models can be co-registered with the 2D map, allowing the visualization (and virtual navigation) of the 3D models as well as the accurate orthoprojection of those complex 3D sites onto the 2D mosaic (Nicosevici and Garcia 2013). The resulting hybrid maps provide better resolution than acoustic bathymetry, while improving the accuracy of the 2D mapping. The maps can be represented in a multimodal environment allowing new visualization and manipulation techniques as well as the integration of data provided by other sensors mounted on the survey vehicle, both from the seafloor and the water column (e.g. temperature, PH, salinity, water current speeds, turbidity, etc.).

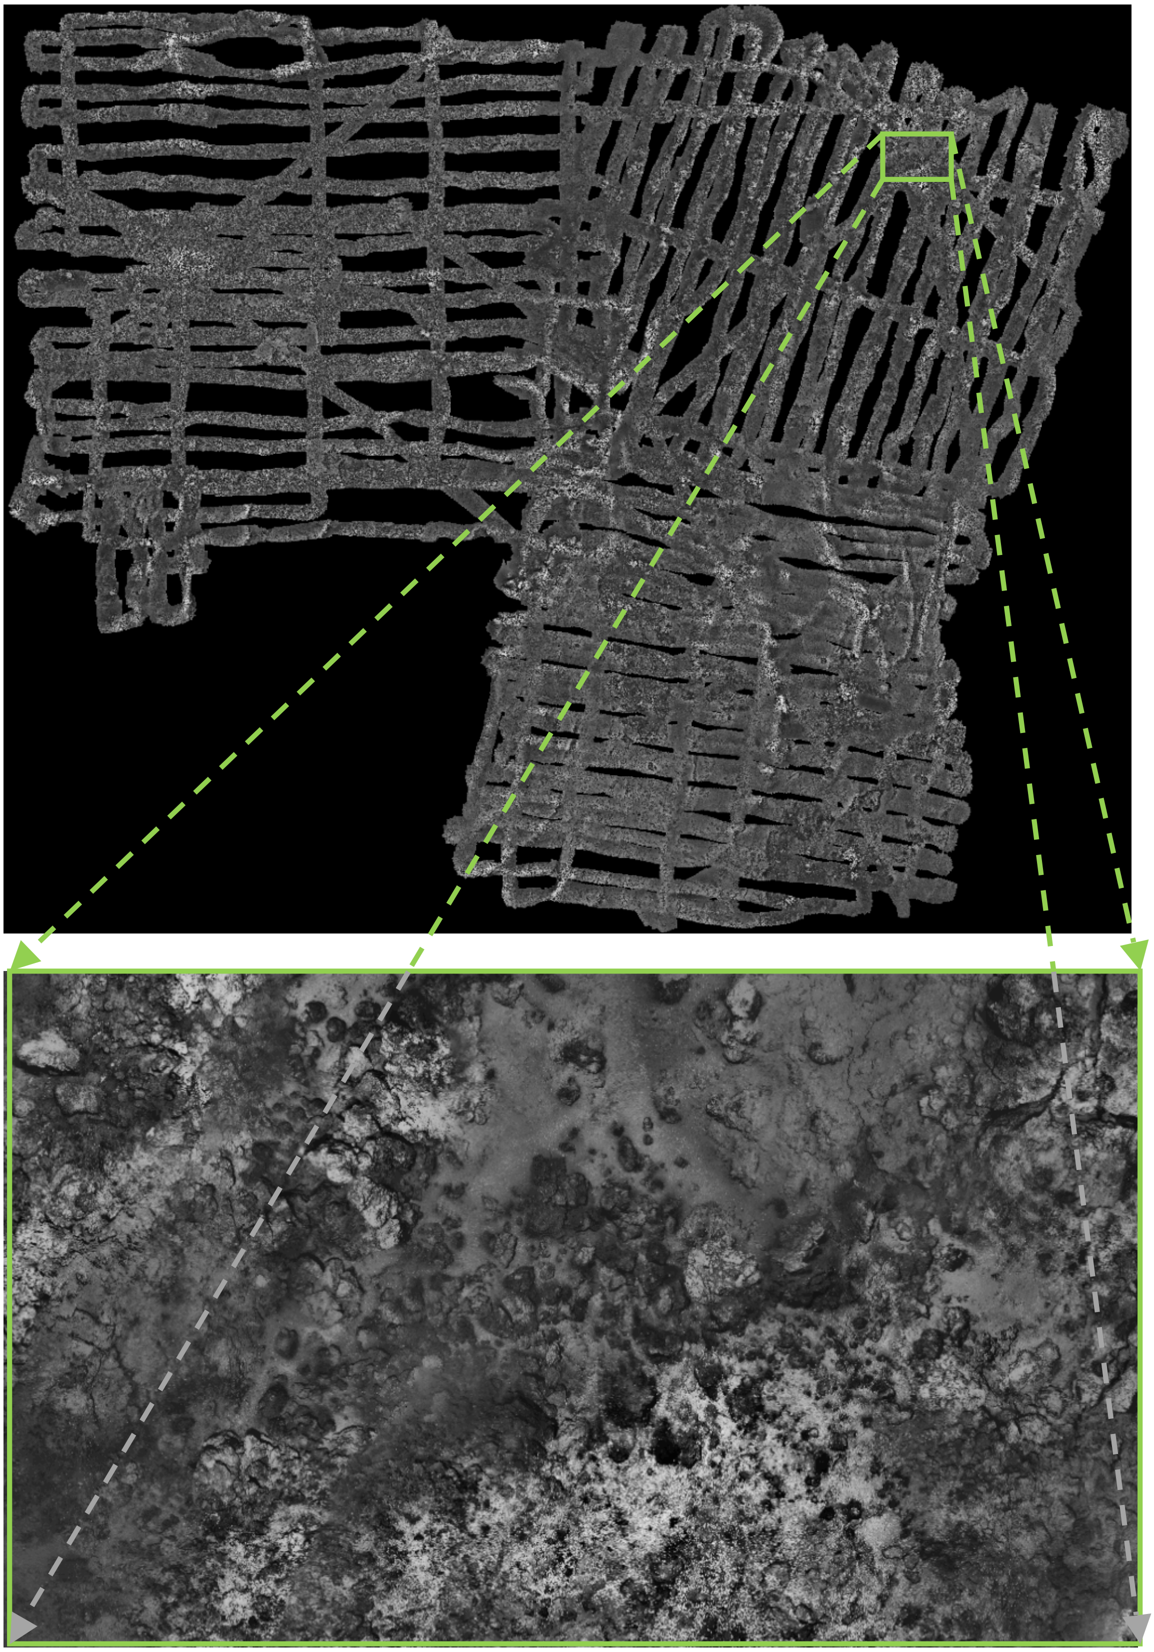

Fig. 2. 2D mosaic of the MOMAR08 cruise in the mid-Atlantic Ridge. The mosaic represents 3 full days of survey with Remotely Operated Vehicle VICTOR6000 and covers an area of nearly one square kilometer.

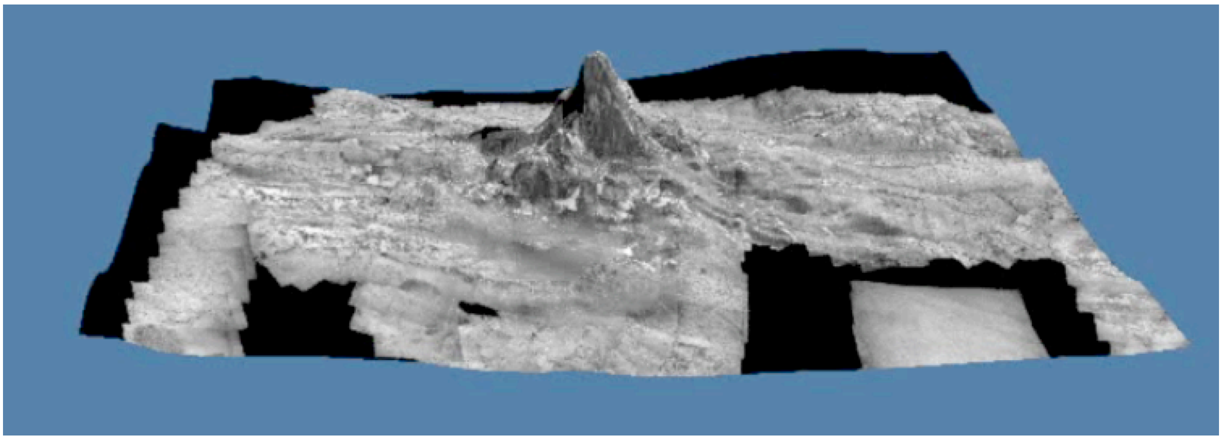

Fig. 3. Combination of 2.5D bathymetric relief and texture from a photomosaic.

In addition, to fully enable temporal studies, critical in the long-term tracking of natural and human processes, it is also possible to develop a software module that is tailored to incorporate change detection tools, to monitor and quantify active processes (Escartin et al., 2013). These change detection tools can be either user-supervised (Barreyre et al., 2012) or automatic (Shihavuddin et al., 2013a; Shihavuddin et al., 2013b) using machine learning techniques. Such temporal analysis is critical in the assessment of seabed morphology over time, with a representative example being the monitoring of coral reef regression as a direct consequence of global warming (Delaunoy et al., 2008).

4. Mapping in 3D

Many of the regions of interest for the scientific community are located in areas with 3-dimensional relief, and optical images can also help in the modeling and monitoring of these areas (De Filippo et al., 2013). For this reason, research and development of new algorithms is a must to enable the creation of 3D representations of the seafloor over areas of a few hundreds of square meters, while having enough ground resolution to allow the identification of colonies of a few centimeters size. Thus, advanced solutions are required to optimally process, combine and visualize the collected data, so that it can be exploited by the end user.

Ongoing work is addressing the streamlining of the processing of 2D and 3D data so that it can be merged into georeferenced, multi-modal maps of the areas of interest. Each map consists of a 2D photo-mosaic of the area, a lower resolution bathymetry, 3D models of specific regions of interest and a set of data from other sensors carried by the AUV (e.g. temperature, PH, salinity, water current speeds, turbidity, etc.).

It should be mentioned here that one of the important aspects of the project is the combined use of *orthomosaicing* and very detailed 3D models of the ocean floor. This novel mapping methodology addresses the limitations of classic mosaicing techniques, by enabling high resolution, accurate 2D mapping of large areas, even in the presence of 3D relief.

A quick review of the state of the art on underwater 3D modelling revels that most of the proposals retrieve the surface in a 2.5D representation (see Figure 3). Furthermore, given the straightforwardness of changing from depth readings to 2D representation, the creation of a triangulated elevation map is usually considered a side result (Singh et al., 2007; Nicosevici and Garcia 2013). However, novel 3D reconstruction processes are built to include a more general scenario, where the sensor (such as a camera) can be mounted in a general configuration; that is, located anywhere on the robot and with any orientation. With this new configuration, objects can be observed from new viewpoints, so that the retrieved ranges are no longer suitable for projection onto a plane (and hence, a 2.5D map cannot be built). Viewing the object from arbitrary positions allows a better understanding of the global shape of the object since its features can be observed from angles that are more suitable to their exploration.

In this direction, Structure-from-motion (SFM) approaches use the corresponding points between pairs of images in the sequence to reconstruct both the motion of the camera during acquisition, and the structure of the scene in the form of a sparse point cloud, as illustrated in Figure 4 (Nicosevici and Garcia, 2013).

Once the camera poses are estimated, dense point cloud retrieving techniques are used to obtain a 3D point cloud description of the scene with a larger number of samples. These methods are based on relaxing the matching criteria in order to produce as many point matches between images as possible. These techniques may follow a close to brute-force approach (Yang and Pollefeys, 2003) or an iterative growing procedure (Furukawa and Ponce, 2010). Obviously, this intensive matching is very likely to generate noisy point clouds containing a huge number of outlier points, given that wrong or poor matches translate into wrong 3D locations that corrupt the final point set (see Figure 5 (a)).

The lack of continuity in the most likely corrupted dense point set retrieved by the above presented techniques burdens further computations and proper visualization to be performed using this data. Thus, the next step aims at obtaining a smooth continuous surface representation in the form of a triangle mesh. We observed that there is a lack of application of these techniques in underwater applications (Campos et al., 2011). In order to deal with the previously mentioned problem of corruption in the dense point sets, we developed techniques allowing the input data to be corrupted with both noise and outliers (Campos et al., 2013; Campos et al., 2014). These methods rely on providing an approximation of the shape of the object in the form of small surfaces of local support generated from the input points. Then, by means of an adaptation of the restricted Delaunay triangulation paradigm (Boissonnat and Oudot, 2005), these approximations can be meshed following a user-desired set of criteria regarding the quality of the output triangles. In Figure 5 (b) we find a sample of a retrieved surface. While most well-known methods in the state-of-the art require the additional information of normals at input points (Kazhdan and Hoppe, 2013), the meshing methods developed within the MUMAP project can be used with raw point sets, without the need of knowing the normals (Campos et al., 2013; Campos et al., 2014). Moreover, our methods allow the recovery of bounded surfaces that usually appear when exploring an underwater scenario.

After the recovery of the surface in the form of a mesh of triangles, one can benefit from the photogrammetric origin of the models generated to use original images to provide texture to the model (Garcia et al., 2011b), as seen in Figure 5 (c). However, blending methods might be needed in these texture mapping techniques to alleviate the differences in illumination on parts of the texture obtained from different images, as well as to mitigate possible camera registration errors (Lempitsky and Ivanov, 2007).

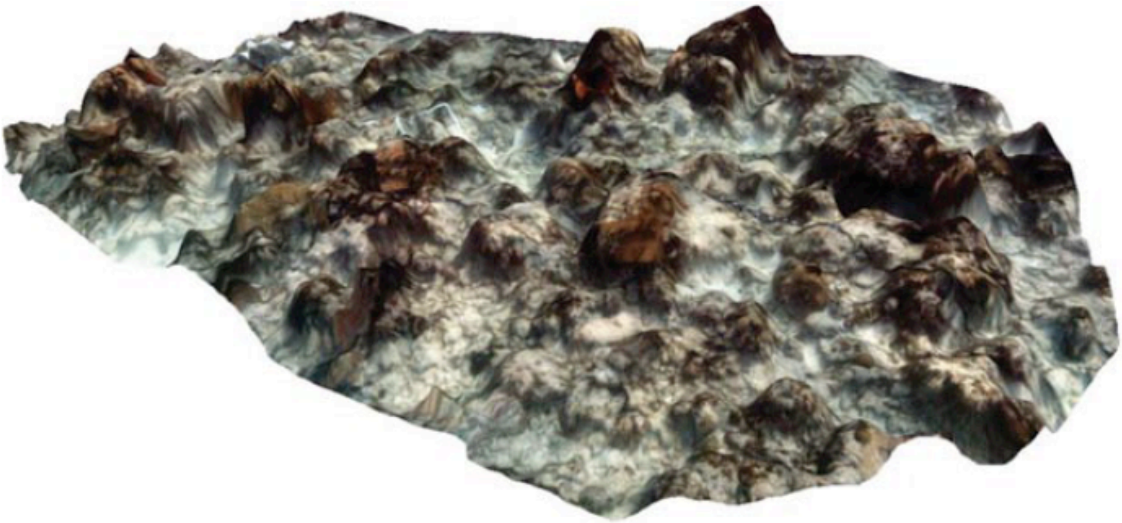

Fig. 4. Estimation of the 3D relief of the seafloor using a single calibrated camera and a structure-from-motion approach. The model has approximately 240,000 vertices and covers an area of 12 × 19 m.

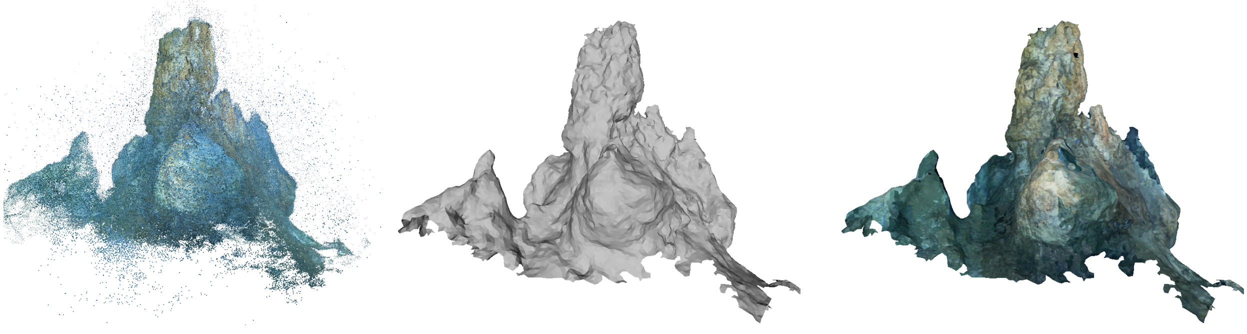

Fig. 5. Steps of the 3D reconstruction pipeline to recover the structure of a hydrothermal chimney. The dense point set representation of the object can be seen in (a), along with its reconstructed surface in (b) and its texture-mapped version in (c).

5. Conclusions

While the seafloor covers 70% of the Earth's surface, our understanding of processes operating there has been greatly hindered by the lack of imaging techniques similar to those available for land. Seafloor imaging requires obtaining numerous photos over large areas, and is therefore less efficient than the imaging of the Earth's surface from airborne and space systems, which can provide single images over very large areas, making the mapping of whole planets feasible. In recent times, however, underwater seafloor exploration and intervention techniques have advanced greatly. We are now able to image large areas of the seafloor with AUVs.

This paper describes the MUMAP project, which proposes the development of a new seafloor mapping system for large seafloor surfaces, with direct applications to oceanography, geology, biology, engineering, underwater archaeology, and environmental studies, and with applicability for offshore industry (oil, gas, CO_2 offshore storage).

The development proposed in the MUMAP project provides techniques and tools to a broad scientific community (biologists, ecologists, geologists, physicists), and that will allow advancements that are possible with existing data and processing tools. Furthermore, these tools will open the door to numerous non-scientific applications ranging from industry (offshore operations, drilling, mining) to the public domain (environmental impact and remediation, infrastructure construction, public works, risk assessment, etc.).

In the near future we can foresee efficient seafloor mapping algorithms using simultaneously multiple AUVs, where the distributed mapping system should be able to identify overlapping image pairs in the trajectories carried out by the different robots during the topology estimation process (Elibol et al., 2014b), being this a cornerstone for efficiently mapping large areas of the seafloor.

Acknowledgements

This work was funded through the Spanish Ministry of Science and Innovation (MICINN) under grant CTM2010-15216.

References

Barreyre, T., Escartin, J., Garcia, R., Cannat, M., Mittelstaedt, E. and Prados, R. (2012). Structure, temporal evolution, and heat flux estimates from a deep-sea hydrothermal field derived from seafloor image mosaics, Geochemistry, Geophysics and Geosystems, vol. 13, no. 4, pp. 1-29.

Boissonnat, J.-D. and Oudot, S. (2005). Provably good sampling and meshing of surfaces. Graphical Models, 67: 405-451.

Campos, R., Garcia, R. and Nicosevici, T. (2011). Surface reconstruction methods for the recovery of 3D models from underwater sites, IEEE OCEANS Conference, pp. 1-10, Santander, Spain.

Campos, R., Garcia, R., Alliez, P. and Yvinec, M. (2013). Splats-based Surface Reconstruction from Defect-Laden Point Sets, Graphical Models. Vol. 75, no. 6, pp. 346-361.

Campos, R., Garcia, R., Alliez, P. and Yvinec, M. (2014). A Surface Reconstruction Method for In-Detail Underwater 3D Optical Mapping, International Journal of Robotics Research. (*in press*)

De Filippo, F., Gracias, N., Garcia, R., Ferrer, J. and Bruno, F. (2013). Incremental underwater mapping in 6DOF with stereo tracking, Instrumentation Viewpoint, no. 15, pp. 20-21.

Delaunoy, O., Gracias N. and Garcia, R. (2008). Towards detecting changes in underwater image sequences, MTS/IEEE Techno-Ocean Conference (Oceans'08) Kobe, Japan.

Elibol, A., Gracias, N. and Garcia, R. (2010). Augmented State-Extended Kalman Filter Combined Framework for Topology Estimation in Large Area Underwater Mapping, Journal of Field Robotics, vol. 27, no. 5, pp. 656-674.

Elibol, A., Garcia, R. and Gracias, N. (2011). "A New Global Alignment Approach for Underwater Optical Mapping," Ocean Engineering, vol. 38, no. 10, pp. 1207-1219.

Elibol, A., Gracias, N. and Garcia, R. (2012). Efficient Topology Estimation for Large Scale Optical Mapping, Springer. ISBN 978-3-642-30312-8

Elibol, A., Gracias, N. and Garcia, R. (2013). Fast Topology Estimation for Image Mosaicing using Adaptive Information Thresholding, Robotics and Autonomous Systems, vol. 61, no. 2, pp. 125–136.

Elibol, A., Gracias, N., Garcia, R. and Kim, J. (2014a). Graph Theory Approach for Match Reduction in Image Mosaicing, Journal of the Optical Society of America A. vol. 31, no. 4, pp. 773-782.

Elibol, A., Kim, J., Gracias, N. and Garcia, R. (2014b). Efficient Image Mosaicing for Multi-robot Visual Underwater Mapping, Pattern Recognition Letters. Vol. 46, pp. 20–26.

Escartin, J., Garcia, R., Barreyre, T., Cannat, M., Gracias, N., Shihavuddin, A. and Mittelstaedt, E. (2013). Optical methods to monitor temporal changes at the seafloor: The Lucky Strike deep-sea hydrothermal vent field (Mid-Atlantic Ridge), IEEE International Underwater Technology Symposium.

Ferrer, J. and Garcia, R. (2011). Bias Reduction for Stereo Triangulation, Electronics Letters, vol. 46, no. 25, pp. 1665-1666.

Fisher, W. S., Davis, W. P., Quarles, R. L., Patrick. J., Campbell, J. G., Harris, P. S., Hemmer, B. L. and Parsons, M. (2007). Characterizing coral condition using estimates of three-dimensional colony surface area. In Environmental Monitoring and Assessment, vol. 125, no. 1, pp. 347–360.

Furukawa, Y. and Ponce, J. (2010). Accurate, dense, and robust multiview stereopsis. IEEE Transactions on Pattern Analysis and Machine Intelligence, 32(8):1362-1376.

Garcia, R. and Gracias, N. (2011a). Detection of Interest Points in Turbid Underwater Images, IEEE OCEANS Conference, pp. 1-9, Santander, Spain.

Garcia, R., Campos, R. and Escartín, J. (2011b). High-resolution 3D reconstruction of the seafloor for environmental monitoring and modelling". in IROS 2011, Workshop on Robotics for Environmental Monitoring.

Gracias, N., Ridao, P., Garcia, R., Escartin, J., L'Hour, M., Cibecchini, F., Campos, R., Carreras, M., Ribas, D., Palomeras, N., Magi, Ll., Palomer, A., Nicosevici, T., Prados, R., Hegedus, R., Neumann, L., De Filippo, F. and Mallios, A. (2013). Mapping the Moon: Using a lightweight AUV to survey the site of the 17th Century ship 'La Lune'. MTS/IEEE OCEANS Conference, Bergen, Norway.

Kazhdan, M. and Hoppe, H. (2013). Screened Poisson surface reconstruction. ACM Transactions on Graphics, 32(3): 29:1-29:13.

Lempitsky, V. and Ivanov, D. (2007). Seamless mosaicing of image-based texture maps. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1-6.

Neumann, L., Garcia, R., Basa, J. and Hegedus, R. (2013). Acquisition and Visualization Techniques for Narrow Spectral Color Imaging, Journal of the Optical Society of America A. Vol. 30, no. 6, pp. 1039–1052.

Nicosevici, T. and Garcia, R. (2012). Automatic Visual Bag-of-Words for Online Robot Navigation and Mapping, IEEE Transactions on Robotics, vol. 28, no. 4, pp. 886-898.

Nicosevici, T. and Garcia, R. (2013). Efficient 3D Scene Modeling and Mosaicing, Springer. ISBN 978-3-642-36417-4

Ninio, R., Delean, S., Osborne, K. and Sweatman, H. (2003). Estimating cover of benthic organisms from underwater video images: variability associated with multiple observers, Mar. Ecol. Progr. Ser. 265, pp. 107-116.

Porter, J.W., Kosmynin, V., Patterson, K.L., Jaap, W.C., Wheaton, J.L., Hackett, K., Lybolt, M., Tsokos, C.P., Yanev, G., Marcinek, D.M., Dotten, J., Eaken, D., Patterson, M., Meier, O.W., Brill, M. and Dustan, P. (2002). Detection of coral reef change in the Florida Keys Coral Reef Monitoring Project. Porter, J.W. and K.G. Porter (eds) The Everglades, Florida Bay, and Coral Reefs of the Florida Keys: An Ecosystem Sourcebook, CRC Press, pp. 749-769, Boca Raton, Florida.

Prados, R., Garcia, R. and Neumann, L. (2014). Image Blending Techniques and their Application in Underwater Mosaicing, Springer. ISBN 978-3-319-05557-2

Prados, R., Garcia, R., Escartin, J. and Neumann, L. (2011). Challenges of Close-Range Underwater Optical Mapping, IEEE OCEANS Conference, pp. 1-10, Santander, Spain.

Prados, R., Garcia, R., Gracias, N., Escartín, J. and Neumann, L. (2012). A Novel Blending Technique for Underwater Giga-Mosaicing, IEEE Journal of Oceanic Engineering, vol. 37, no. 4, pp. 626-644.

Ridao, P., Ribas, D., Palomeras, N., Carreras, M., Mallios, A., Hurtós, N., Gracias, N., Magí, Ll., Garcia, R., Campos, R., Prados, R. and Escartin, J. (2013). Operational validation of Girona500 AUV, 5th International Workshop on Marine Technology.

Shihavuddin, A., Gracias, N. and Garcia, R. (2012). Online Sunflicker Removal using Dynamic Texture Prediction, International Conference on Computer Vision Theory and Applications (VISSAP'2012), pp. 161-167, Rome (Italy).

Shihavuddin, A., Gracias, N., Garcia, R., Escartin, J. and Pedersen, R. (2013a). Automated classification and thematic mapping of bacterial mats in the North Sea, MTS/IEEE OCEANS Conference, Bergen, Norway.

Shihavuddin, A., Gracias, N., Garcia, R., Gleason, A.C.R. and Gintert, B. (2013b). Image-Based Coral Reef Classification and Thematic Mapping, Remote Sensing, 5(4), pp. 1809-1841.

Singh, H., Roman, C., Pizarro, O., Eustice, R. and Can, A. (2007). Towards high-resolution imaging from underwater vehicles. The International Journal of Robotics Research, 26(1):55{74, 2007.

Solan, M., Germano, J.D., Rhoads, D.C., Smith, C., Michaud, E., Parry, D., Wenzhofer, F., Kennedy, B., Henriques, C., Battle, E., Carey, D., Iocco, L., Valente, R., Watson, J. and Rosenberg, R. (2003). Towards a greater understanding of pattern, scale and process in marine benthic systems: a picture is worth a thousand worms. In J. Exp. Mar. Biol. Ecol., 285, pp. 313-338.

Yang, R. and Pollefeys, M. (2003). Multi-resolution real-time stereo on commodity graphics hardware. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), volume 1, pages 211-217 vol.1.