Efficient autonomous image mosaicing with applications to coral reef monitoring

Armagan Elibol,
Dept. of Mathematical Engineering
Yildiz Technical University, Istanbul Turkey
aelibol@yildiz.edu.tr

Nuno Gracias, Rafael Garcia
Computer Vision and Robotics Group
University of Girona, Girona Spain
{ngracias, rafa}@eia.udg.edu

Art Gleason, Brooke Gintert,
Diego Lirman and R. Pam Reid
University of Miami, Miami USA
{agleason,bgintert,dlirman,preid}@rsmas.miami.edu

***Abstract*—Over the past decade, several image mosaicing methods have been proposed in robotic mapping and remote sensing applications. Owing to the rapid developments on obtaining optical data from areas beyond human reach, there is a high demand from different scientists for creating large-area image mosaics often using images as the only source of information. One of the most important steps in the mosaicing process is the motion estimation between overlapping images to obtain the topology, i.e, the spatial relationships between images. In this paper, we propose a generic framework for feature-based image mosaicing capable of obtaining the topology with a reduced number of matching attempts and to get the best possible trajectory estimation. Innovative aspects include the use of a fast image similarity criterion combined with a minimum spanning tree (MST) solution, to obtain a tentative topology. This topology is improved by attempting image matching over the pairs of higher overlap evidence. Unlike previous approaches for large-area mosaicing, our framework is able to naturally deal with the cases where time consecutive images cannot be matched successfully, such as completely unordered sets. We conclude this paper by presenting an environmental application of this mosaicing approach for monitoring coral reefs.**

I. Introduction

Image mosaicing methods have been widely used for panoramic imaging [1] and mapping [2]. Aerial and satellite imagery of the Earth's surface, merged with high-resolution topographic models, have proven a key tool to understand physical processes of our planet (geological, hydrological, biological, etc.), to monitor environmental changes, whether man-induced or natural, resource management, development and infrastructure planning (public works, remediation plans, land use, etc.). Robots are becoming more and more important in gathering optical data from places where human cannot reach. In robot mapping(i.e., aerial and/or underwater), when a robot is surveying a large area using only a down-looking camera, it is of interest to obtain a global view of the area. To have a wide-area visual representation of the scene, it is necessary to create large-area maps (mosaics). Mosaics enable different applications such as geological [3], [4] and archaeological surveys [5], ecology studies [6], [7], [8], environmental damage assesment [9], [10] and detection of temporal changes in [11]. Therefore, there is a high demand from different science communities for creating optical maps of areas where human cannot reach.

When a robot is mapping an area in a scattering media [4], illumination effects, noise, lack of image contrast and blurring are phenomena that make image registration difficult. This leads to inaccuracies in image registration that cause misalignment when images are mapped onto the mosaic (global) frame. To compose images into a mosaic form, several steps are needed and one of the most important steps is image registration. In the absence of any other information, most of the existing methods try the exhaustive strategy of matching all images against all. However, this approach is only feasible for a small number of images. Since large-area surveys might comprise of several hundreds to tens of thousands of images [4], [5], the all-against-all strategy becomes impractical.

To overcome this problem, we propose in this paper a generic image mosaicing scheme aiming to get the complete topology with minimum number of image matching attempts while simultaneously obtaining a globally coherent mosaic. The algorithm takes as input a set of images that has been previously acquired. The time order is not taken into account; therefore, the image set can be totally unordered. Our technique first infers image similarity information using a fast method based on the Euclidean distance between feature descriptors. Then, this similarity is used to create a tentative topology with associated uncertainty. The estimated uncertainty, although large at the beginning, is gradually reduced by successive iterations of image matching and bundle adjustment.

II. Related Work

Quality constraints of image mosaics are usually very strict, especially for mapping purposes, as the mosaic might be used for global navigation [12], localization of interest areas [4] and detection of temporal changes [11]. Several image mosaicing approaches for creating underwater mosaics have been proposed over the last decade [13], [14], [15], [16]. Pizarro et al. [13] proposed a mosaicing system that exploited navigation and attitude information for bundle adjustment. Madjidi and Negahdaripour [14], addressed the global alignment problem for a submersible equipped with stereo cameras, using a mixed adjustment model to recursively determine the pose of the vehicle. Rzhanov et al. described in [15] a methodology that exploited navigation data to build geo-referenced photomosaics of the mid-ocean ridges at the East Pacific Rise. Ferrer et al. [16] proposed a global alignment method for creating

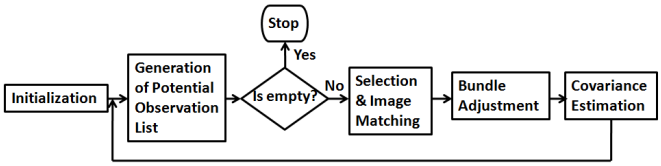

Figure 1. Pipeline of the algorithm.

large-scale underwater photo-mosaics that combines image registration information and 3D position estimates provided by navigation sensors available in deep water surveys. Bulow et al. [17] proposed an online mosaicing (image-to-mosaic registration) method for Unmanned Aerial Vehicles using an image registration method based on Fourier-Mellin transformation. As they have stated, the proposed method fails if there is not enough overlapping area between time consecutive images. None of the methods mentioned above have concentrated on the challenge of finding the non-time consecutive matches when only image information is available.

III. Topology Estimation

In this work, we assume that the optical axis of the camera is kept perpendicular to the scene, which is approximately flat [1]. Each image has an associated planar transformation [18] with 4 degrees of freedom, ${}^{M}\mathbf{H}_i$, detaied in Eq. (1), that relates the image frame i to a common mosaic frame M (i.e., absolute homography).

$$ {}^{M}\mathbf{H}_i = \begin{bmatrix} a_i & -b_i & c_i \\ b_i & a_i & d_i \\ 0 & 0 & 1 \end{bmatrix} \tag{1} $$

For simplicity, we consider the reference frame to be the coordinate system of the first image, so that ${}^{M}\mathbf{H}_1$ is equal to identity, thus it is not part of the parameter vector to be estimated.

Our scheme is composed of five different steps: (1) Initialization, (2) Generation of the list of potential overlapping image pairs, (3) Image selection and matching, (4) Minimisation of the reprojection error and (5) Covariance propagation. The pipeline of the proposed method is illustrated in Fig. 1

Initialization The initialization step aims at obtaining information on the similarity between images and to establish an initial link between them. This similarity information is intended to be computed in a fast and approximate way. First, Scale Invariant Features (SIFT) [19] are extracted. Then a small subset of feature descriptors (between 100 and 200) are randomly selected from each image, and compared against the subsets of all other images. This comparison is performed using the Euclidean distance between feature descriptors [19].

For a given pair of images, our similarity measure is proportional to the number of descriptors that are associated using the distance criterion. The computational cost of this similarity measure is comparatively low, since it mainly involves computing the angles between a small set of descriptor vectors. Our multi-threaded C implementation allows for computing the measure in 2.5 miliseconds on a standard desktop machine for a pair of images with 200 descriptors each. In order to establish the initial link between images, we use a MST where weights of the edges are the inverted initial similarity values. MST of a weighted graph is a subset of edges that form a tree whose sum of weights of edges is minimum[20]. MST provides a connected tree which is composed of the most similar image pairs according to the similarity information. The initial relative homographies between those image pairs that are in the MST are treated as identity mappings with very large uncertainty. Using these relative homographies, the absolute homographies are computed along with its uncertainty which is propagated using a first order approximation [21].

Finding Potential Overlapping Image Pairs This step aims to find the overlapping image pairs given an estimate of the absolute homographies and its uncertainty. We propose to use an approach which employs two successive different tests. The first test computes the distance between image centers by taking into account their uncertainties. If this distance is smaller than a selected threshold (such as, size of the image diagonal) then the second test is applied. The second test consists of generating several noisy instances of homographies using the propagated covariances and computing the overlapping area between images. If the normalized average overlapping area is above a given threshold (e.g., 30%), then the pair is added to the list of potential overlapping image pairs.

Image Selection and Matching This step starts by selecting a subset of image pairs from the potential overlapping pair list. The main reason for this selection is that it is not feasible to attempt to match the whole list since the list might contain many non-overlapping pairs due to the high uncertainty and drift on the current absolute homographies. We have used the estimated overlapping area between potential overlapping pairs as a ranking criterion. The size of the subset is determined by a simple Computational Time criterion, where the total matching time for the subset is approximately equal to the computational time spent on all other steps in the iteration (generation of list of potentially overlapping pairs, bundle adjustment and covariance propagation). For image matching, features are detected and matched using SIFT [19], followed by outlier rejection and motion estimation [18].

Minimizing the Reprojection Error The error terms resulting from image registration are measured in the image reference frames. We have employed a standard Bundle Adjustment (BA) approach [12] which minimizes the weighted reprojection error over homographies. Reprojection error is expressed as follows:

$$ \varepsilon = \sum_k \sum_t \sum_{j=1}^{n} \| {}^{k}\mathbf{p}_j - {}^{M}\mathbf{H}_k^{-1} \cdot {}^{M}\mathbf{H}_t \cdot {}^{t}\mathbf{p}_j \|_2 + \| {}^{t}\mathbf{p}_j - {}^{M}\mathbf{H}_t^{-1} \cdot {}^{M}\mathbf{H}_k \cdot {}^{k}\mathbf{p}_j \|_2 \tag{2} $$

where k and t are a pair of images that were successfully

[1] In this work, it is assumed that the navigation altitude of the vehicle is large with respect to the 3D relief of the scene

matched, n is the total number of correspondences between the overlapping image pairs, $({}^{M}\mathbf{H}_k, {}^{M}\mathbf{H}_t)$ are the absolute homographies for images k and t, respectively. ${}^{k}\mathbf{p}_j = ({}^{k}x_j, {}^{k}y_j, 1)$ encodes the coordinates of the j^{th} feature point in image k, while ${}^{t}\mathbf{p}_j$ are the coordinates of the same scene point in image t. The weight included cost function is given in Eq. (3), which is the L_2 norm of a stack of weighted residues. f is minimized over θ, which contains the parameters for all image homographies.

$$f = \mathbf{R}^T \cdot \mathbf{W} \cdot \mathbf{R} \tag{3}$$

where $\mathbf{R} = \left[\left| \begin{array}{l} {}^{i}\mathbf{r}_j^k = {}^{i}\mathbf{p}_k - {}^{i}\mathbf{H}_j \cdot {}^{j}\mathbf{p}_k \\ {}^{j}\mathbf{r}_i^k = {}^{j}\mathbf{p}_k - {}^{j}\mathbf{H}_i \cdot {}^{i}\mathbf{p}_k \end{array} \right|_{stack} \right]$ is a $4N_{pm} \times 1$ vector and $\mathbf{W}$ is a diagonal $4N_{pm} \times 4N_{pm}$ matrix of weights for each residue. N_{pm} is the total number of correspondences. Finally, ${}^{i}\mathbf{H}_j = {}^{i}\mathbf{H}_M \cdot {}^{M}\mathbf{H}_j$ and ${}^{j}\mathbf{H}_i = {}^{j}\mathbf{H}_M \cdot {}^{M}\mathbf{H}_i$. The minimisation of the cost function in Eq. 3 was carried out using the MATLAB™ *lsqnonlin* function for large-scale methods. The optimization algorithm requires the computation of the Jacobian matrix containing the derivatives of all residuals with respect to all parameters. The Jacobian matrix has a clearly defined block structure, and the sparsity pattern is constant [22], [23]. In our implementation, analytic expressions were derived and used for computing the Jacobian matrix.

Covariance Propagation We apply Haralick's method [21] to propagate the uncertainty of the resulting homography estimations of BA. f in Eq. 3 is a function of parameter vector θ and $\mathbf{x}$ containing all data affected by noise. After optimization, the first order approximation to the uncertainty in the parameters is given by [21]

$$\mathbf{\Sigma}_\theta = \left(\frac{\partial g}{\partial \theta}\right)^{-1} \cdot \frac{\partial g}{\partial \mathbf{x}} \cdot \mathbf{\Sigma}_\mathbf{x} \cdot \left(\frac{\partial g}{\partial \mathbf{x}}\right)^{T} \cdot \left(\frac{\partial g}{\partial \theta}\right)^{-1} \tag{4}$$

where $\mathbf{\Sigma}_\mathbf{x}$ is the covariance matrix of $\mathbf{x}$ and g is the jacobian of f with respect to θ.

IV. Experimental Results

The generic scheme described in the previous section was tested on a general setup for image surveys using an underwater robot equipped with a down-looking camera. We have tested our scheme on a real data set in which some time-consecutive images do not have overlapping areas. The dataset was extracted from an underwater image sequence acquired by a Phantom 500 ROV during a survey in Andros, the Bahamas [7]. This data set is composed of two horizontal and three vertical transects. The total number of images is 112. In addition, we have changed the order of the images to have more non-overlapping consecutive pairs between ordered images. The initial similarity matrix is depicted in Fig. 2. The resulting final trajectory and uncertainty can be seen in Fig. 3 while the resulting mosaic is depicted in Fig. 4 and Table I summarises the results. The first column of the table corresponds to the tested method. The second column shows the total number of successfully matched image pairs that have least 20 inliers. The third column contains the total number of image pairs that were not successfully matched (*unsuccessful observations*). The percentage of the total number of image matching attempts with respect to all-against-all image matching attempts is given in the fourth column. The last column corresponds to the average reprojection error calculated using all the correspondences with the resulting set of homographies for each tested strategy.

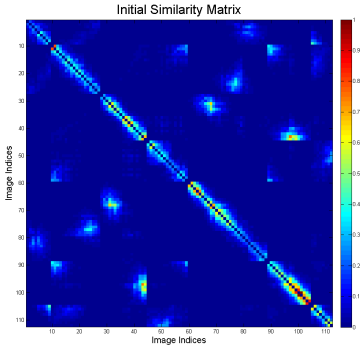

Figure 2. Initial Similarity Matrix of the dataset. This matrix was computed using a maximum of 200 feature points. Values are scaled to $[0, 1]$.

As there are some broken links between the time-consecutive images, the traditional iterative topology estimation method proposed in [12] cannot be applied. It can be observed in Table I that our scheme was able to get 94% of the total overlapping pairs with a considerably smaller number of image matching attempts. The second line shows the results for matching all the pairs for which the similarity matrix provides at least 20 descriptor associations to attempt RANSAC [18]. The third line is for the all-against-all strategy. Initial similarity matrix almost suggests all-against-all matching.

Table I
Summary of results.

Strategy	Successful Obs.	Unsuccessful Obs.	% of attempts wrt all-against-all	Avg. Error in pixels	Std. Dev. in pixels
1. Proposed Scheme	278	1,201	23.79	5.12	3.67
2. Similarity Matrix	294	5,900	99.65	4.86	3.61
3. All Against All	294	5,922	100.00	4.86	3.61

In order to show that the proposed scheme is not dependent on the image order in the dataset. We have also tested our scheme on a small dataset that is composed of approximately two transects having a few overlapping pairs between them. We have changed the image order fully in random manner only keeping the first image same. This is mainly to represent the topology in common global frame. Then, we reorganize the initial similarity matrix by taking into account this randomly generated new image order. The initial similarity matrices are depicted in Fig. 5. We have run our scheme on both original captured order and the randomly generated order of images. Results are summarized in Table II.

From the results, it can be seen that our scheme can work with fully unordered datasets as it uses similarity matrix obtained from images. Final trajectory and uncertainties on image centers are given in Fig.6 Final trajectory for the randomly ordered images is given in Fig.7.

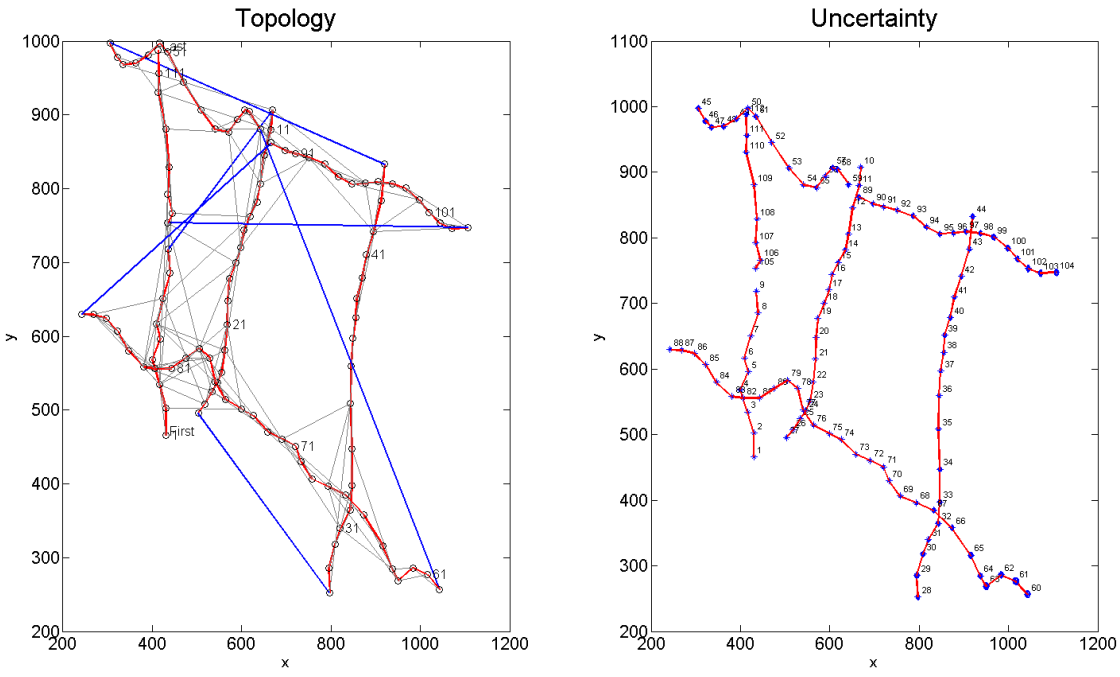

(a) (b)

Figure 3. Axes are in pixels and approximately 200 pixels per meter. (a) Final trajectory obtained by the proposed method. Red lines are links between time consecutive overlapping images while the black ones are between the non-time consecutive. Blue lines show the non-overlapping time consecutive images. (b) Uncertainty on the final trajectory. Uncertainty ellipses are drawn with 95% confidence level.

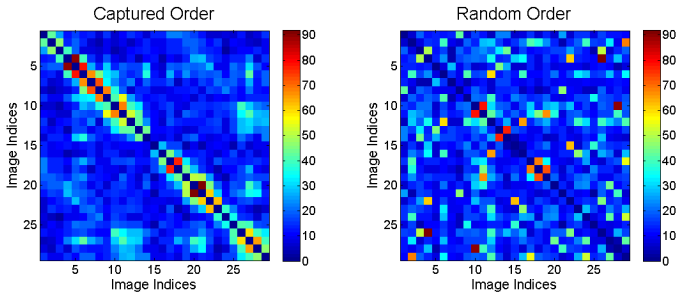

Figure 5. Initial Similarity Matrices of the second dataset. These matrices were computed using a maximum of 100 feature points.

Table II
SUMMARY OF RESULTS.

Strategy	Successful Obs.	Unsuccessful Obs.	% of attempts wrt all-against-all	Avg. Error in pixels	Std. Dev. in pixels
1. Captured Order of images	62	30	22.66	7.05	4.20
1. Random Order of images	62	31	22.90	7.03	4.20
2. Similarity Matrix	64	342	100.00	6.63	4.17
3. All Against All	64	342	100.00	6.63	4.17

V. ENVIRONMENTAL APPLICATION OF MOSAICS ON CORAL REEFS

Recent declines in coral reefs across the globe underscore the need for new scientific tools to better understand ecological patterns and rates of change. Given that multiple factors are typically responsible for changes within reef ecosystems, the monitoring of reef health must be carried out at multiple spatial and temporal scales, rather than relying on measuring only a few parameters. Comprehensive assessment of coral reef resources demands a hierarchical mapping strategy involving micro-scale to macroscale measurements. Image-based mosaics of the seabed enable observations on a mesoscale of 10's to 100's of m, with mm-scale resolution.

Underwater image-based mosaics address several limitations of traditional, diver-based, coral reef monitoring techniques. First, mosaics provide a landscape view of coral reefs that has previously been unobtainable [7]. Second, mosaics are efficient tools for tracking patterns of change over time [25]. Third, mosaics have high spatial accuracy at both the scale of an individual coral colony [7] and at the scale of the entire mosaic [10].

The potential use of mesoscale, or "landscape," mosaics has been investigated for several coral reef-related applications, including: documenting hurricane damage at both the colony and reef-framework scale [9], mapping mesophotic [26], [27] and deep-water [28] coral ecosystems, quantifying the area damaged by a ship that had run aground [10], and tracking individual colonies through time [9], [25]. Of these, the ship grounding and individual monitoring take particular advantage of the new scale of observation enabled by landscape mosaics.

Accurately documenting patterns of physical damage (and subsequent recovery patterns) to benthic habitats can be especially challenging when the spatial extent of injuries exceeds tens of square meters. Such injuries are often too large and

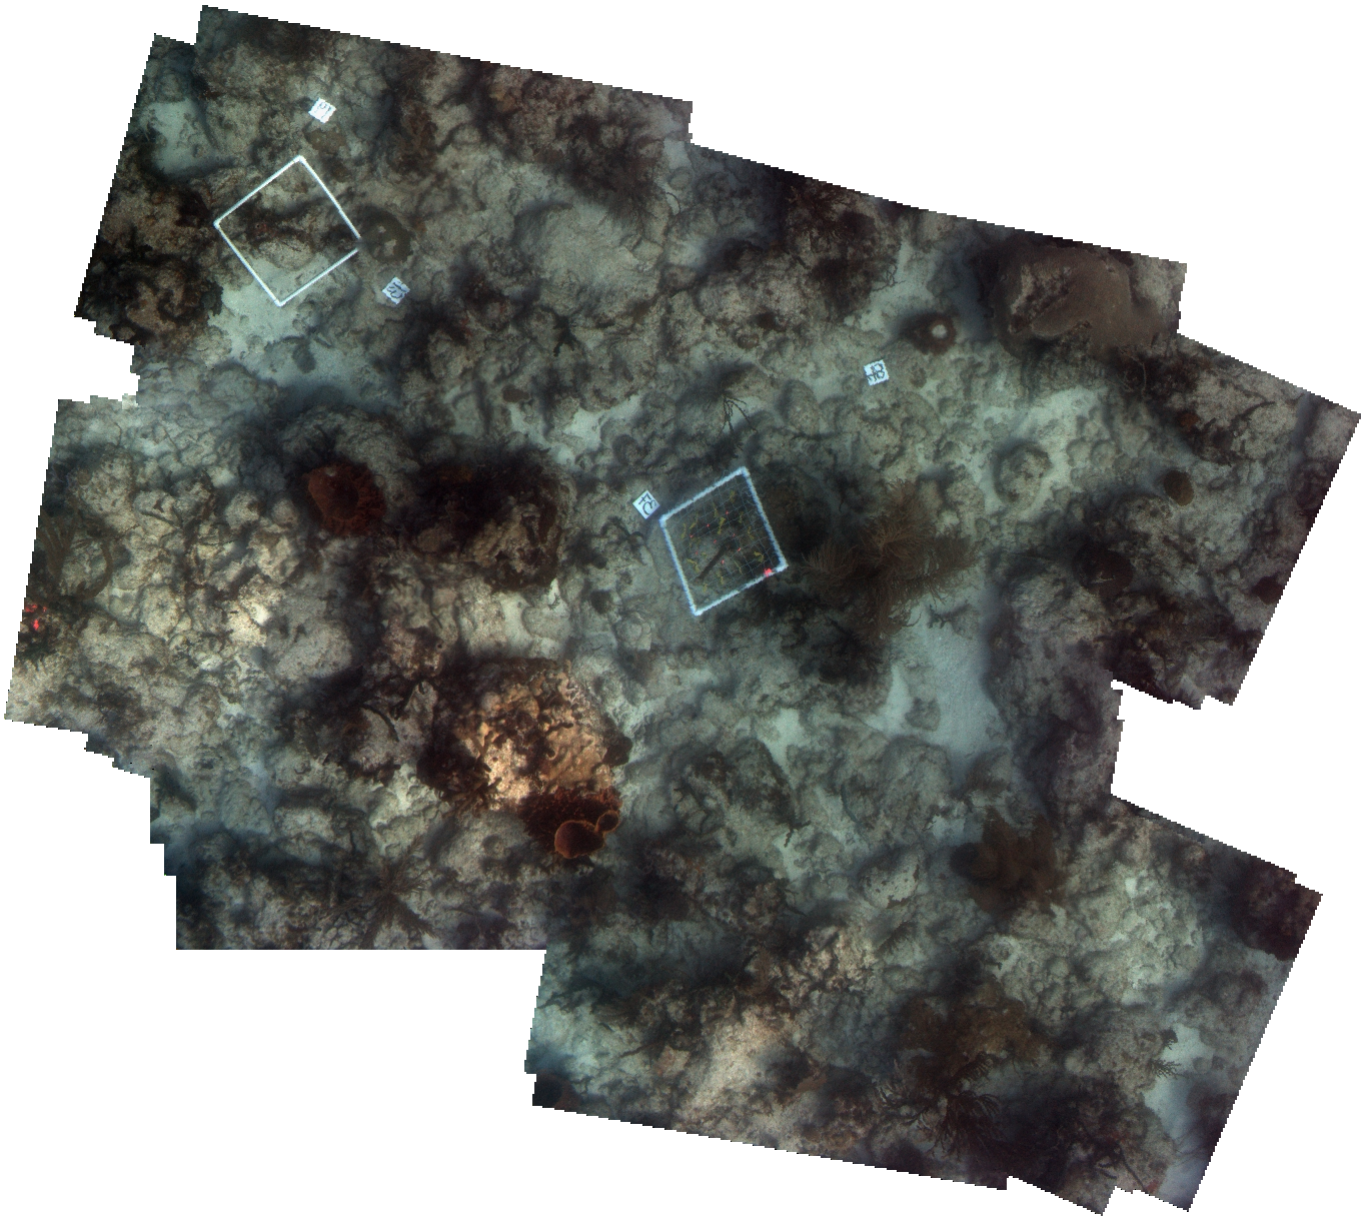

Figure 4. Resulting final mosaic image. After global alignment, the final mosaic was blended using graph cut algorithms [24]

difficult to measure in situ by divers and too small or costly to be quantified effectively using aerial and satellite remote sensing tools. Documenting the extent of damage caused by physical disturbance is one of the main challenges of post-damage surveys in coral reef habitats. In the case of vessel groundings, the effective and accurate assessment of the extent of the damage caused is a crucial first step in the Habitat Equivalency Analysis (HEA) required to determine the amount of compensatory restoration required [29], [30]. Grounding scars are commonly measured in situ by divers using flexible tapes following the "fishbone" method described by [31]. In addition, the boundaries of the damaged areas or the positions of objects of interest (e.g., injured corals) are surveyed using surface-deployed GPS units positioned over specific locations, and the extent of the damage is later calculated from the polygon delineated by the GPS locations.

Landscape mosaics are advantageous for assessment of damage and recovery operations because they permit simultaneous mapping of both the scale of the entire injury as well as the fine scale appropriate to assess individual colony damage. Lirman et al. [10] showed that an estimate of the damaged area derived from a landscape mosaic agreed within 2% with an estimate produced by a diver using differential GPS. Gleason et al. [32] mapped a large scar ($>3,000m^2$) in Puerto Rico by assembling multiple individual landscape mosaics acquired by divers. Despite the huge area, the Puerto Rico mosaic was rendered at 1 cm spatial resolution, allowing the assessment of individual coral colonies (Fig. 8).

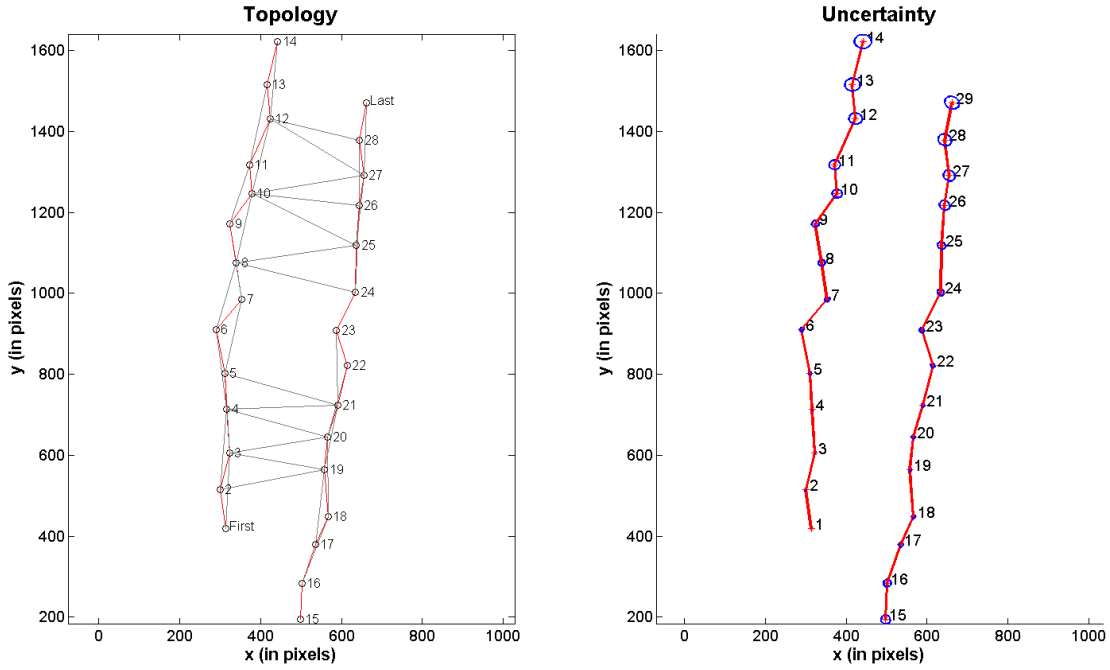

Figure 6. Axes are in pixels and approximately 200 pixels per meter. (a) Final trajectory obtained by the proposed method. Red lines are links between time consecutive overlapping images while the black ones are between the non-time consecutive. Blue lines show the non-overlapping time consecutive images. (b) Uncertainty on the final trajectory. Uncertainty ellipses are drawn with 95% confidence level.

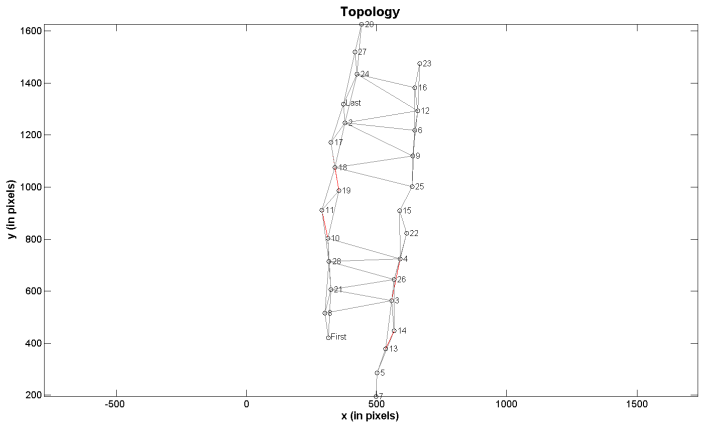

Figure 7. Axes are in pixels and approximately 200 pixels per meter. Final trajectory obtained by the proposed scheme using randomly ordered images in the dataset.

Monitoring individual coral colonies requires establishing a permanent site and periodically measuring the size and condition of each colony within that site. Currently, state-of-the-art assessment techniques rely on divers to measure colony sizes using tape measures or meter sticks, and to estimate colony condition visually. The drawbacks of this technique are, first, that the divers must tag each colony so the specific coral can be identified in the future, and, second, that the divers must have the relevant biological/ecological training to identify corals and assess their condition in the field. This diver-based tagging method is the state-of-the-art method used to establish new permanent plots today.

Landscape mosaics have two key advantages relative to the diver-based method that improve colony-based monitoring. First, tags are not necessary because repeat mosaics taken over the same area can be registered to one another. Removing the reliance on tags eliminates the need for physical contact with corals, thereby greatly reducing the potential for inadvertent damage and the amount of gear that must be permanently attached to the seafloor (e.g., nails, tags, markers). Tagging is

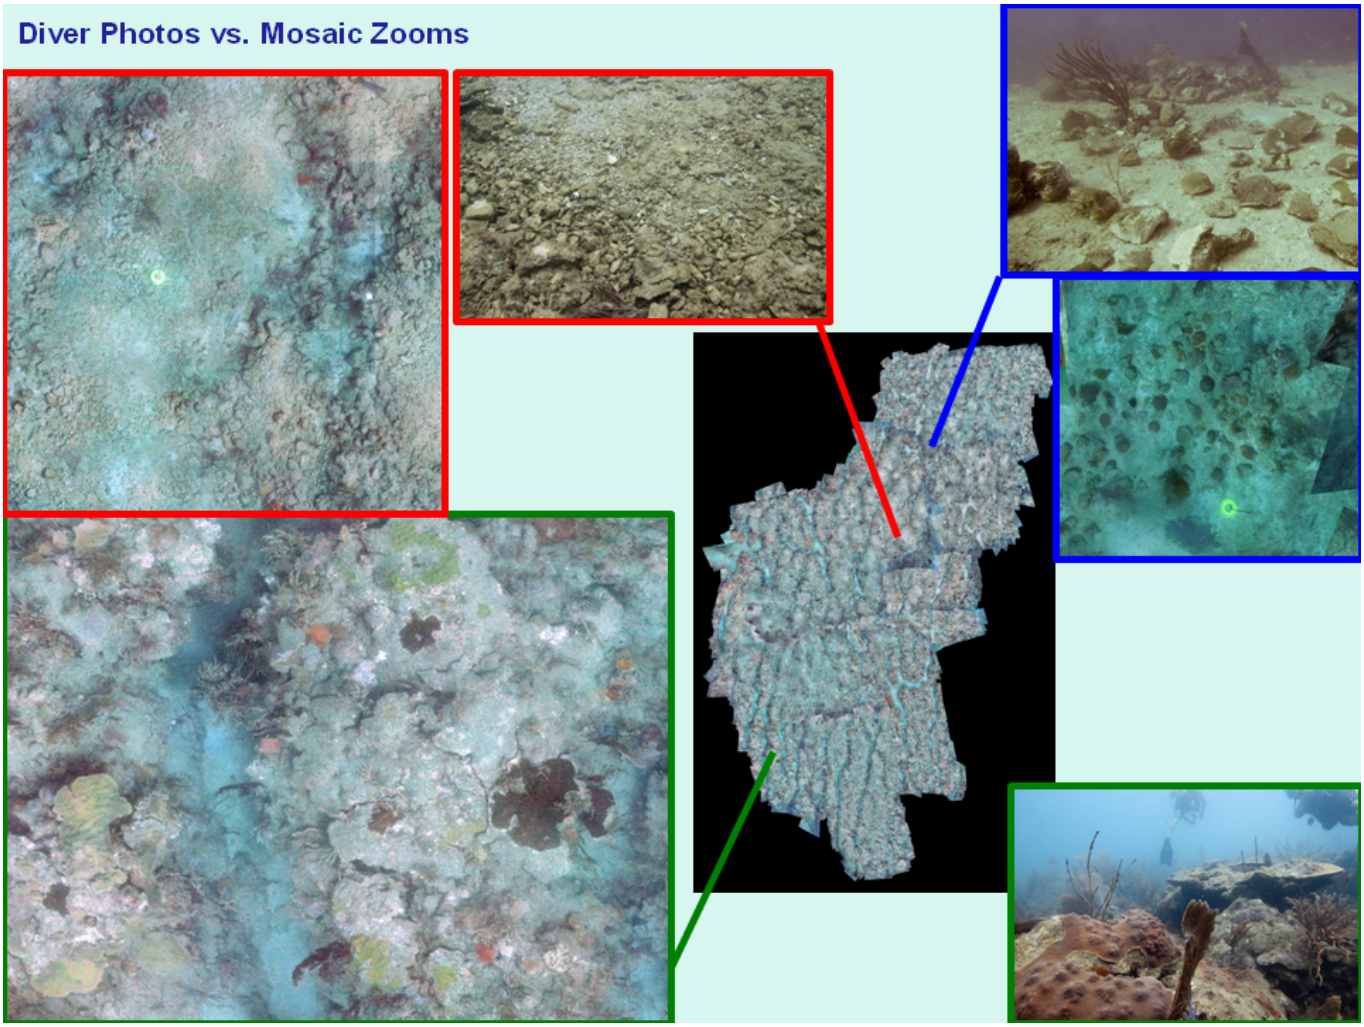

Figure 8. Landscape mosaic of a ship-grounding scar in Puerto Rico. The dimensions of the mosaiced area are 117×67 m, covering over $4{,}700$ m^2 with 1×1 cm pixels. The entire mosaic is presented here at $<$5% of its full resolution, but the insets show portions of the mosaic at full resolution to give an idea of the level of detail in comparison with oblique images acquired by divers. Note the pulverized rock within the area of maximum damage (red insets), the condition of the unaffected area surrounding the scar (green insets), and the coral fragments ready to be adhered to the substrate as part of the remediation process (blue insets).

labor intensive both during the initial establishment of the plot and during re-location of colonies through time. Furthermore, tags can get lost due to burial, failure of the attachment mechanism, biofouling, or simply diver error, and lost tags represent lost data as colonies can no-longer be identified. Second, divers who collect the data do not necessarily have to have extensive training in coral reef biology.

VI. Conclusions

We have presented a generic scheme for creating large-area mosaics with application to environmental mapping over areas where only image information is available. Our scheme aims at obtaining the topology with minimum number of image matching attempts as well as obtaining the best possible trajectory estimation. The proposed approach is able to deal with the cases where time consecutive images do not have overlapping areas.

Acknowledgments.

This work was partially funded through the Spanish Ministry of Education and Science (MCINN) under grant CTM2010-15216, EU Project FP7-ICT-2009-248497 and US SERDP program project CS-1333. Nuno Gracias was supported by MCINN under the Ramon y Cajal program.

References

[1] H. Sawhney, S. Hsu, and R. Kumar, "Robust video mosaicing through topology inference and local to global alignment," in *European Conference on Computer Vision*, vol. II, Freiburg, Germany, June 1998, pp. 103–119.

[2] R. Kumar, H. Sawhney, S. Samarasekera, S. Hsu, H. Tao, Y. Guo, K. Hanna, A. Pope, R. Wildes, D. Hirvonen, M. Hansen, and P. Burt, "Aerial video surveillance and exploitation," *Proceedings of the IEEE*, vol. 89, no. 10, pp. 1518–1539, Oct 2001.

[3] Z. Zhu, E. Riseman, A. Hanson, and H. Schultz, "An efficient method for geo-referenced video mosaicing for environmental monitoring," *Machine Vision and Applications*, vol. 16, no. 4, pp. 203–216, 2005.

[4] J. Escartin, R. Garcia, O. Delaunoy, J. Ferrer, N. Gracias, A. Elibol, X. Cufi, L. Neumann, D. J. Fornari, S. E. Humpris, and J. Renard, "Globally aligned photomosaic of the Lucky Strike hydrothermal vent field (Mid-Atlantic Ridge, 3718.5'N): Release of georeferenced data, mosaic construction, and viewing software," *Geochemistry Geophysics Geosystems*, vol. 9, no. 12, p. Q12009, 2008.

[5] R. M. Eustice, H. Singh, J. J. Leonard, and M. R. Walter, "Visually mapping the RMS titanic: Conservative covariance estimates for SLAM information filters," *International Journal of Robotics Research*, vol. 25, no. 12, pp. 1223–1242, 2006.

[6] K. Jerosch, A. Ldtke, M. Schlter, and G. Ioannidis, "Automatic content-based analysis of georeferenced image data: Detection of beggiatoa mats in seafloor video mosaics from the hakon mosby mud volcano," *Computers & Geosciences*, vol. 33, no. 2, pp. 202 – 218, 2007.

[7] D. Lirman, N. Gracias, B. Gintert, A. Gleason, R. P. Reid, S. Negahdaripour, and P. Kramer, "Development and application of a video–mosaic survey technology to document the status of coral reef communities," *Environmental Monitoring and Assessment*, vol. 159, pp. 59–73, 2007.

[8] O. Pizarro and H. Singh, "Toward large-area mosaicing for underwater scientific applications," *IEEE Journal of Oceanic Engineering*, vol. 28, no. 4, pp. 651–672, October 2003.

[9] A. Gleason, D. Lirman, D. Williams, N. Gracias, B. Gintert, H. Madjidi, R. Reid, G. Boynton, S. Negahdaripour, M. Miller, and P. Kramer, "Documenting hurricane impacts on coral reefs using two–dimensional video–mosaic technology," *Marine Ecology*, vol. 28, no. 2, pp. 254–258, June 2007.

[10] D. Lirman, N. Gracias, B. Gintert, A. Gleason, G. Deangelo, M. Dick, E. Martinez, and R. P. Reid, "Damage and recovery assessment of vessel grounding injuries on coral reef habitats using georeferenced landscape video mosaics," *Limnology and Oceanography: Methods*, vol. 8, pp. 88–97, 2010.

[11] O. Delaunoy, N. Gracias, and R. Garcia, "Towards detecting changes in underwater image sequences," in *OCEANS 2008-MTS/IEEE Techno-Ocean*, Kobe, Japan, 2008, pp. 1–8.

[12] N. Gracias, S. Zwaan, A. Bernardino, and J. Santos-Victor, "Mosaic based navigation for autonomous underwater vehicles," *IEEE Journal of Oceanic Engineering*, vol. 28, no. 4, pp. 609–624, Oct. 2003.

[13] O. Pizarro, R. Eustice, and H. Singh, "Relative pose estimation for instrumented, calibrated imaging platforms," in *Digital Image Computing Techniques and Applications*, Sydney, Australia, December 2003, pp. 601–612.

[14] H. Madjidi and S. Negahdaripour, "Global alignment of sensor positions with noisy motion measurements," *IEEE Transactions on Robotics*, vol. 21, no. 6, pp. 1092–1104, 2005.

[15] Y. Rzhanov, L. Mayer, S. Beaulieu, T. Shank, S. Soule, and D. Fornari, "Deep-sea geo-referenced video mosaics," in *MTS/IEEE OCEANS Conference*, Boston, USA, sep. 2006, pp. 2319–2324.

[16] J. Ferrer, A. Elibol, O. Delaunoy, N. Gracias, and R. Garcia, "Large-area photo-mosaics using global alignment and navigation data," in *MTS/IEEE OCEANS Conference*, Vancouver, Canada, November 2007, pp. 1–9.

[17] H. Bulow and A. Birk, "Fast and robust photomapping with an unmanned aerial vehicle UAV," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, 10-15 2009, pp. 3368 –3373.

[18] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Harlow, UK: Cambridge University Press, 2004.

[19] D. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[20] S. Skiena, *Implementing Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*. Reading, MA: Addison-Wesley, 1990.

[21] R. M. Haralick, "Propagating covariance in computer vision," in *9. Theoretical Foundations of Computer Vision*, March 1998, pp. 95–114.

[22] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon, "Bundle adjustment – A modern synthesis," in *Vision Algorithms: Theory and Practice*, ser. LNCS. Springer Verlag, 2000, pp. 298–375.

[23] D. Capel, *Image Mosaicing and Super-resolution*. London: Springer-Verlag, 2004.

[24] N. Gracias, M. Mahoor, S. Negahdaripour, and A. Gleason, "Fast image blending using watersheds and graph cuts," *Image and Vision Computing*, vol. 27, no. 5, pp. 597–607, 2009.

[25] B. Gintert, N. Gracias, D. Lirman, A. Gleason, P. Kramer, and R. Reid, "Second-generation landscape mosaics of coral reefs," in *Proceedings of the 11th International Coral Reef Symposium*, 2008, pp. 7–11.

[26] R. Armstrong, "Deep zooxanthellate coral reefs of the puerto rico: Us virgin islands insular platform," *Coral Reefs*, vol. 26, no. 4, pp. 945–945, 2007.

[27] A. Gleason, N. Gracias, D. Lirman, B. Gintert, T. Smith, M. Dick, and R. Reid, "Landscape video mosaic from a mesophotic coral reef," *Coral Reefs*, vol. 29, no. 2, pp. 253–253, 2010.

[28] M. Ludvigsen, B. Sortland, G. Johnsen, and H. Singh, "Applications of geo-referenced underwater photo mosaics in marine biology and archaeology," *Oceanography*, vol. 20, no. 4, pp. 140–149, 2007.

[29] J. Milon and R. Dodge, "Applying habitat equivalency analysis for coral reef damage assessment and restoration," *Bulletin of marine science*, vol. 69, no. 2, pp. 975–988, 2001.

[30] S. Shutler, S. Gittings, T. Penn, and J. Schittone, *Compensatory restoration: how much is enough? Legal, economic, and ecological considerations*. Taylor & Francis, 2006, pp. 77–93.

[31] J. Hudson and W. Goodwin, "Assessment of vessel grounding injury to coral reef and seagrass habitats in the florida keys national marine sanctuary, florida: protocol and methods," *Bulletin of marine science*, vol. 69, no. 2, pp. 509–516, 2001.

[32] A. C. R. Gleason, D. Lirman, N. Gracias, T. Moore, S. Griffin, M. Gonzalez, B. Gintert, and R. P. Reid, "Damage assessment of vessel grounding injuries on coral reef habitats using underwater landscape mosaics," in *Proceedings of the 63rd Gulf and Caribbean Fisheries Institute*, November 2010.