# Rugosity, Slope and Aspect from Bathymetric Stereo Image Reconstructions

Ariell Friedman, Oscar Pizarro & Stefan B. Williams
Australian Centre for Field Robotics (ACFR)
University of Sydney,
NSW 2006, Australia
Email: a.friedman@acfr.usyd.edu.au

***Abstract*—This paper demonstrates how multi-scale measures of rugosity, slope and aspect can be derived from fine-scale bathymetric reconstructions created using geo-referenced stereo imagery collected by an Autonomous Underwater Vehicle (AUV). We briefly describe the 3D triangular meshes generated from the stereo images and then present a detailed overview of how rugosity can be derived by considering the area of triangles within a window and their projection onto the plane of best fit. By obtaining the plane of best fit, slope and aspect can be calculated with very little extra effort. The results are validated on a simulated surface and the effects of mesh resolution and window size are explored. The technique is demonstrated on real data gathered by an AUV on surveys that cover several linear kilometres and consist of thousands of images. The ability to distinguish habitat types based on rugosity and slope are demonstrated through K-means cluster analysis. A human labelled data set is then used to train a SVM classifier that exhibits promising habitat classification potential based on rugosity and slope.**

## I. Introduction

Terrain complexity is strongly correlated to biodiversity in marine environments [1], [2]. Even when represented as bathymetry, it is necessary to abstract this structural information into simpler representations in order to perform analytical work. Ecologists typically use indices, such as rugosity, slope and aspect to describe habitat complexity.

Rugosity is traditionally measured *in situ* by divers along a single, linear profile using chain-tape methods [1], [3], [4] or profile gauges [1]. In these methods, rugosity is typically calculated to be the ratio between the length of the contoured surface profile and the linear distance between the end points. These traditional methods are labour intensive, depth limited and put humans at risk. As a result, surveys tend to be spatially and temporally sparse and not easily repeatable. These measurements are performed using SCUBA in less than 30m of water, which means that the majority of marine habitats cannot be described by this measure. Furthermore, the outputs of transects using the traditional approach are calculated at a single, predefined resolution and scale imposed by the link-size (or gauge spacing) and the transect length. This is an important limitation since some spatial patterns and processes operate at scales not well resolved by the particular choice of chain or gauge [4].

Alternative work has investigated the possibility of deriving these measures from bathymetric maps collected during ship borne surveys [5], [6]. However, these methods cannot resolve fine scale structure due to the resolution of the survey data. Other studies have used airborne LIDAR to measure topography [7], but unfortunately these measurements are depth limited due to the poor penetration of laser in water.

Autonomous Underwater Vehicles (AUVs) capable of high precision navigation and equipped with stereo cameras can recover bathymetry at fine resolutions over relatively large, contiguous extents of seafloor compared to diver based assessments. Measures derived from an AUV make it possible to obtain dense coverage over larger extents beyond diver depths [8]. Given that the survey and calculations are performed without humans, a potential source of bias is eliminated. Furthermore, AUVs with proper navigation systems provide the ability for easily repeatable transects, making it is possible to revisit an area of interest for monitoring purposes.

Rugosity for a 3D surface is typically defined as the ratio between the area of the contoured or draped surface and the area of its orthogonal projection onto a plane. A method for calculating rugosity on raster-formatted digital elevation grids has been proposed by Jenness in [9], however, forcing an irregular mesh into a raster grid causes reconstructions to be less accurate. Furthermore, Jenness's proposed rugosity calculation is subject to edge-effect problems and by using the horizontal planimetric area, slope is encoded into the result.

The method proposed in this paper uses the geo-referenced stereo imagery obtained using an AUV to generate fine-scale bathymetric reconstructions with centimetre resolution in the form of irregular 3D triangular meshes [8]. Rugosity can be derived by calculating the sum of the area of the triangles that make up the surface and dividing that by the sum of their projections onto the plane of best fit. Fitting a plane to the data ensures that rugosity and slope are decoupled. As a consequence of fitting a plane, obtaining slope and aspect is trivial.

The remainder of this paper is organized as follows. Section II briefly describes the AUV *Sirius* and how 3D triangular meshes are generated from stereo data. Section III presents a detailed derivation of how rugosity, slope and aspect can be calculated from the triangular meshes. The results are validated in Section IV, which includes running the algorithm on a simulated surface and exploring the effects of mesh resolution and window size. Section V demonstrates the method on real data and shows its potential for discriminating habitats.

Section VI presents conclusions and outlines potential future work.

## II. Data Collection

The Australian Centre for Field Robotics (ACFR) at the University of Sydney operates an AUV called *Sirius*. It is a modified version of the SeaBED AUV [10] built at the Woods Hole Oceanographic Institution and was designed for high-resolution, georeferenced survey work. For navigation and localisation the vehicle uses DVL, USBL, GPS, a pressure sensor, a compass and inclinometers. It is passively stable in pitch and roll and controls yaw with a pair of horizontal aft-facing thrusters that are also used to control forward and backward motion. Depth is controlled by a single vertical thruster. For survey work, the AUV is equipped with multibeam sonar, a suite of oceanographic sensors and a stereo camera pair. The downward looking camera pair has a baseline of approximately 7 cm and pixel resolutions $1360 \times 1024$ with a field of view of $42 \times 34$ degrees. The AUV is typically programmed to maintain an altitude of 2 m and travel at a velocity of 0.5 m/s capturing stereo image pairs at a frequency of 1-2 Hz.

Using the automated system developed by Johnson-Roberson et. al. [8], the stereo imagery is combined with vehicle poses to deliver fine-scale 3D, texture mapped terrain reconstructions. The processing pipeline for generating the stereo meshes is broken down into the following steps:
1) Data Acquisition and Preprocessing. The stereo imagery is acquired by the AUV and preprocessed to partially compensate for lighting and wavelength-dependent colour absorption.
2) Stereo Depth Estimation. 2D feature points are extracted from each image pair, correspondences are proposed, outliers are rejected, and 3D position is determined by triangulation. The 3D point clouds are converted into Delaunay triangulated meshes.
3) Mesh Aggregation. The individual stereo meshes are put into a common reference frame using SLAM-based poses and fused into a single mesh using volumetric range image processing (VRIP) [11]. Discontinuities between integrated meshes are minimised and simplified versions of the mesh are produced to allow for fast visualization at broad scales.
4) Texturing. The polygons of the complete mesh are assigned textures based on the projection of overlapping imagery.

## III. Calculating Rugosity, Slope & Aspect

In order to calculate rugosity, slope and aspect, it is necessary to first define the surface over which the calculations will be performed. The surface is defined by a Delaunay triangulated mesh which is made up by a set of triangular faces that connect vertices to make a 3D surface. The vertices of the surface are contained in the set $\boldsymbol{V} = \{\boldsymbol{v}_m\}$, such that $\boldsymbol{v}_m \in \mathbb{R}^3$ and $m = 1, ..., M$. $\boldsymbol{v}_m = (x_m, y_m, z_m)$ represents the vertex $m$ described by its $x, y, z$ coordinates. The triangles of the surface are contained in the set $\boldsymbol{T} = \{\boldsymbol{t}_n\}$, where $n = 1, ..., N$, such that $\boldsymbol{t}_n \subset \boldsymbol{V}$. $\boldsymbol{t}_n = (\boldsymbol{v}_{1_n}, \boldsymbol{v}_{2_n}, \boldsymbol{v}_{3_n})$ represents a triangle defined by three vertices in $\boldsymbol{V}$.

The rugosity index $r$ for a particular location in the surface mesh is calculated by centring a window, of specified size over the location and dividing the area of the contoured surface bounded by the window, by the area of the orthogonal projection of the surface onto a plane

$$r = \frac{A}{A'} \tag{1}$$

where $A$ is the area of the contoured surface within the window, and $A'$ is the area of the orthogonal projection of that surface onto a plane.

The window can be described by the subset of triangles and vertices that it encloses. The subset of vertices are contained in $\boldsymbol{X} = \{\boldsymbol{x}_k\}$, such that $k = 1, ..., K$ and $\boldsymbol{X} \subseteq \boldsymbol{V}$. A vertex is only included in $\boldsymbol{X}$ if it forms part of a triangle that falls entirely within the window. The subset of triangles within the window are contained in $\boldsymbol{W} = \{\boldsymbol{w}_j\}$, where $j = 1, ..., J$. $\boldsymbol{w}_j = (\boldsymbol{x}_{1_j}, \boldsymbol{x}_{2_j}, \boldsymbol{x}_{3_j})$ represents a triangle comprised of three vertices in $\boldsymbol{X}$, such that $\boldsymbol{w}_j \subset \boldsymbol{X}$.

The area of the contoured surface bounded by the window $A$, is equal to the summation of the areas of all the individual triangles that are contained within the window

$$A = \sum_{j=1}^{J} a_j. \tag{2}$$

The area of an individual triangle $a_j$ in the contoured surface can be calculated to be half the magnitude of the cross product of the vectors representing two adjacent sides of the triangle. The intuition for this calculation is as follows: let a triangle in the surface be defined by the vertices $\boldsymbol{v_1} = (x_1, y_1, z_1)$, $\boldsymbol{v_2} = (x_2, y_2, z_2)$, $\boldsymbol{v_3} = (x_3, y_3, z_3)$, and the adjacent vectors $\vec{\boldsymbol{v_2 v_1}}$ and $\vec{\boldsymbol{v_2 v_3}}$ to be:

$$\vec{\boldsymbol{v_2 v_1}} = [x_1 - x_2]\mathbf{i} + [y_1 - y_2]\mathbf{j} + [z_1 - z_2]\mathbf{k}$$
$$\vec{\boldsymbol{v_2 v_3}} = [x_3 - x_2]\mathbf{i} + [y_3 - y_2]\mathbf{j} + [z_3 - z_2]\mathbf{k}$$

The area of a parallelogram with sides $\vec{\boldsymbol{v_2 v_1}}$ and $\vec{\boldsymbol{v_2 v_3}}$ is equal to the magnitude of the cross product of vectors representing two adjacent sides. The area of an individual triangle $a_j$ is then half of this, and can be expressed as

$$a_j = \frac{1}{2} || \vec{\boldsymbol{v_2 v_1}} \times \vec{\boldsymbol{v_2 v_3}} ||. \tag{3}$$

Next we need to consider the projected area $A'$, which is the area of the orthogonal projection of the surface contained within the window, onto a plane. The correct choice of plane is an important consideration. Simply projecting the points onto the $x, y$ plane by setting the $z$ components to zero, for example, tethers rugosity to slope. This would mean that flat, steep terrain would exhibit an overstated rugosity index. Ideally, we would like to have rugosity decoupled from slope. Therefore, we require the area of the orthogonal projection of the surface onto the plane that best fits its vertices (contained in $\boldsymbol{X}$). The plane that best represents the data can be obtained using Principle Component Analysis (PCA).

PCA is used to determine the orthogonal projection of the data onto the *principle subspace* (a lower dimensional linear space) such that the variance of the projected data is maximized [12]. It involves evaluating the mean and the covariance matrix of the data $\boldsymbol{X}$ and then finding the eigenvectors and corresponding eigenvalues of the covariance matrix. By ordering the eigenvectors in the order of descending eigenvalues, an ordered orthogonal basis $\boldsymbol{u}$ is created containing the eigenvectors

$$\boldsymbol{u} = (\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{c}}) \tag{4}$$

where $\hat{\boldsymbol{a}}$ is the principle component and has the direction of largest variance of the data, $\hat{\boldsymbol{b}}$ is the secondary component, and $\hat{\boldsymbol{c}}$ is the third component and has the direction of the least variance of the data, and is orthogonal to the principle and secondary components. Consequently, $\hat{\boldsymbol{c}}$ is a direction vector normal to the principle plane of the data, however, it is ambiguous as to whether it is inward or outward facing. Given that the data is obtained from overhead imagery, it is assumed that the outward facing normal will always have an upward facing component. This is enforced by checking the sign of the dot product between $\hat{\boldsymbol{c}}$ and the upward facing unit vector, $\hat{\boldsymbol{k}}$

**if** $(\hat{\boldsymbol{c}} \cdot \hat{\boldsymbol{k}} >= 0)$ **then**
  $\hat{\boldsymbol{p}} = \hat{\boldsymbol{c}}$
**else**
  $\hat{\boldsymbol{p}} = -\hat{\boldsymbol{c}}$
**end if**

where, $\hat{\boldsymbol{p}}$ is the outward-facing normal to the principle plane of the data.

The projected area $A'$ can now be expressed as a summation of the areas of the individually projected triangles bound by the window

$$A' = \sum_{j=1}^{J} a_j(|\hat{\boldsymbol{p}} \cdot \hat{\boldsymbol{n}}_j|) \tag{5}$$

where

$$\hat{\boldsymbol{n}}_j = \frac{\vec{\boldsymbol{v_2 v_1}} \times \vec{\boldsymbol{v_2 v_3}}}{||\vec{\boldsymbol{v_2 v_1}} \times \vec{\boldsymbol{v_2 v_3}}||}$$

is the unit vector normal to the face of triangle $j$ and $|\hat{\boldsymbol{p}} \cdot \hat{\boldsymbol{n}}_j|$ gives a ratio for the projected area of the triangle on the plane to its actual contoured area in 3D space. From this, it is possible to compute the rugosity index shown in Equation 1.

Given that we now have the vector, $\hat{\boldsymbol{p}}$, normal to the plane of best fit, it is relatively straight forward to obtain slope and aspect.

Slope, denoted by $\theta$, refers to the angle between the plane of best fit and the horizontal plane. This angle is equivalent to the angle between the normal vectors of the two planes and can be obtained from their dot product, which is $\hat{\boldsymbol{p}} \cdot \hat{\boldsymbol{k}} = \cos\theta$ (noting that $\hat{\boldsymbol{p}}$ and $\hat{\boldsymbol{k}}$ are both unit vectors). Thus, slope can be calculated as

$$\theta = \cos^{-1}(\hat{\boldsymbol{p}} \cdot \hat{\boldsymbol{k}}). \tag{6}$$

The slope is an angle in the range $(0, \frac{\pi}{2})$.

Aspect, denoted by $\psi$, refers to the direction that the surface slope faces. It is defined as the angle between the positive $x$ axis and the projection of the normal onto the $x, y$ plane. It can be calculated as

$$\psi = \tan^{-1}\left(\frac{p_x}{p_y}\right) \tag{7}$$

where $p_x$ and $p_y$ are the components of $\hat{\boldsymbol{p}}$ in the $x$ and $y$ directions, respectively, and $\tan^{-1}$ is the 4-quadrant inverse tangent that outputs an angle in the range $(-\pi, \pi)$.

## IV. Validation Results

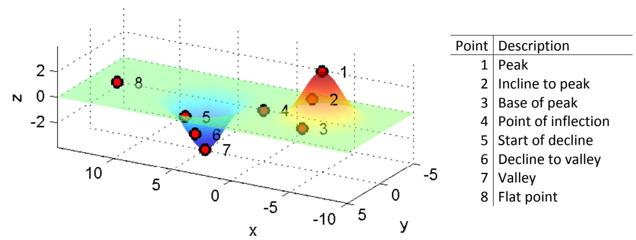


Fig. 1. Simulated Gaussian surface mesh showing some selected points of interest. The surface is made up of two Gaussian distributions where one has been shifted in x and subtracted from the other.

The simulated Gaussian surface shown in Figure 1 was used to test the method for validation purposes and demonstrate that it gives reasonable results before applying the method to real data. A number of points have been selected at different locations on the surface to provide an idea of how the rugosity, aspect and slope behave at different areas on the curve.

Figure 2 shows results for rugosity aspect and slope calculated at every vertex in the surface mesh. The triangle areas $a_j$ and normals $\hat{\boldsymbol{n}}_j$ are calculated over the entire surface mesh once at the beginning of the algorithm and then the window is moved across the surface, centring it over every vertex performing PCA and the necessary calculations for rugosity, aspect and slope. It can be seen that the rugosity is largest at the peak and valley in the surface and the slope is largest along the incline/decline to the peak/valley in the surface. The aspect needs to be considered with reference to the the slope, i.e. at regions where the slope is close to zero, the aspect is relatively erratic since the normal vector points almost directly up and the direction of the component of the normal projected onto the $x - y$ plane, changes dramatically with a small change in any of the variables in the calculation. It should also be noticed that aspect is subject to angular wraparound issues. It is a discontinuous angle measurement that varies between $\{-180°, +180°\}$, where a value of $-180°$ should be interpreted to be the same as a value of $+180°$. This needs to be taken into consideration when considering the results. Direction vectors have been shown on the aspect plot to help explain the result.

### A. Effects of Window Size

Figure 3(a) shows plots for the different points of interest in the simulated surface at selected window sizes to show how

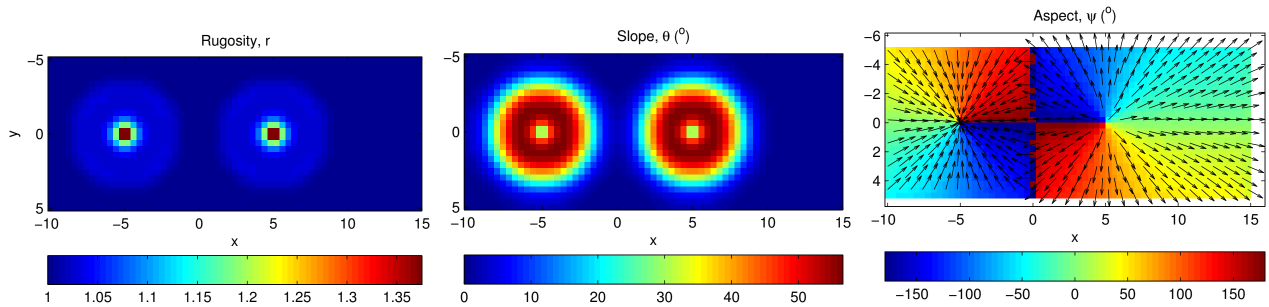


Fig. 2. Rugosity, slope and aspect calculated at every vertex for the simulated Gaussian surface mesh in Figure 1 with a mesh resolution of 50mm and a window size of $1m \times 1m$.

window size affects the results. The window size needs to be chosen with reference to the spatial scales of the environmental features to be considered. It can be likened to the chain/transect length in the conventional chain-tape method, of which the importance of scale has been outlined in [13], [2], [14]. It can be seen that, at the feature scales considered, a window size of $4m \times 4m$ provides the highest level of discrimination for the rugosity calculation amongst the selected features. Smaller window sizes do not capture as much variation in the ruggedness of the surface and the larger $8m \times 8m$ window size essentially "smooths out" the features (as seen by the comparatively low slope values). At the scales considered, the slope increases with decreasing window size and the aspect is unaffected (although, it is noisy at regions with zero slope).

### *B. Effects of Mesh Resolution*

Figure 3(b) shows plots for the different points of interest in the simulated surface at selected window sizes to illustrate the effects of mesh resolution (or level of detail). The mesh resolution is analogous to the link-size for the chain-tape method. The importance of link size is explored in [4]. Over the scales considered, it can be seen that slope and aspect are insensitive to the suface mesh resolution, however, rugosity increases with decreasing resolution.

## V. Results on Real Datasets

The AUV *Sirius* has been used to collect data at various sites around Australian waters. The results presented in this paper are from Scott Reef off Western Australia and the Tasman Peninsula off Tasmania.

Figure 4 shows the results for an AUV survey performed at Scott Reef that densely covered an area of $50m \times 75m$ with 9,831 stereo image pairs. This survey featured a partially populated substrate boundary between dense coral and barren sand, as illustrated by Figure 4(b).

Figure 5 shows the effect of different window sizes on the calculation of rugosity, slope and aspect. A larger window provides more spatial smoothing, however too much smoothing causes information loss, so it is important to select a window size that is appropriate for the scale of the processes to be observed.

It can be seen from Figure 5 that rugosity appears to be a good indicator for the different substrate types and it outlines the boundary between the different substrates shown in Figure 4(a) quite closely. Figure 6 shows unsupervised classification (clustering) results based normalised rugosity and slope. Using K-means to generate 3 clusters, these results show the ability to discriminate sand from reef from the boundary between the two using rugosity and slope. It should be pointed out that the aspect measurement is subject to angular wraparound issues, which is problematic for the purposes of machine learning. Furthermore, measurements of aspect are only likely to be useful for classification purposes when framed in context with water currents and environmental conditions. As such, aspect will not be used a descriptor in this paper.

Figure 7 shows clustering results for a survey performed on the O'Hara Reef in the Tasman Peninsula, Tasmania. This dataset covers several linear kilometres and consists of 11,243 images. Given that the survey coverage is sparse, a dense reconstruction is not possible with this data. Therefore, it does not make sense to attempt to use large window sizes. The choice of window size should be comparable to the swath width. Consequently, rugosity and slope are calculated using a window defined by the size of an individual stereo pair reconstruction (approximately $1.5m \times 1.2m$). Using K-means on normalised rugosity and slope, it was possible to extract 3 distinct clusters from the data. Cluster 1 consisted mostly of barren sand, Cluster 2 featured more rugged reef-like terrain and Cluster 3 appeared to consist mostly of a mixture of coarse sand and rubble. A disadvantage of using the K-means algorithm is that the resulting clusters can depend on the initial random assignments (or seeds).

In an attempt to better validate the classification potential of rugosity and slope the 11,243 images from the O'Hara Reef dataset were individually labelled to be one of 4 habitat classes (Low relief reef, High relief reef, Sand/rubble or Ecklonia). The spatial layout of the manually labelled imagery, along with some sample images from each class are shown in Figure 8.

The labelled dataset was used to train a multi-class Support Vector Machine (SVM) that was implemented using the LIBSVM toolbox [15] in MATLAB. The SVM was trained on the first 10% of the samples from each class using a radial basis kernel function. A summary of the classification output based on normalised rugosity and slope is shown in

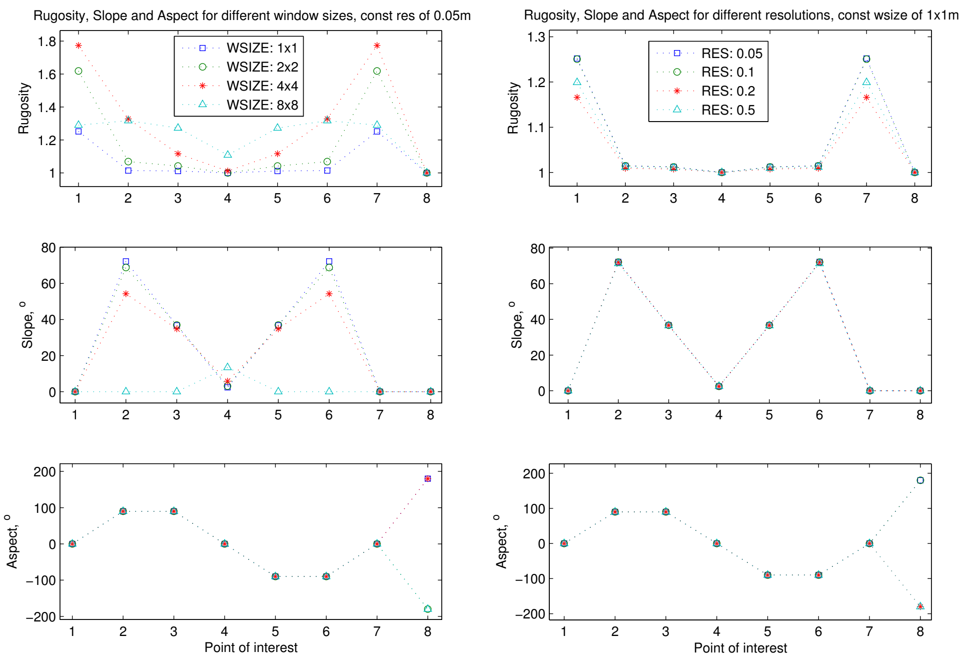


(a) (b)

Fig. 3. Rugosity, Slope and Aspect of the points outlined in Figure 1. (a) Constant surface mesh resolution of 50mm over different window sizes. (b) Constant window size of $1m \times 1m$ over different surface mesh resolutions.

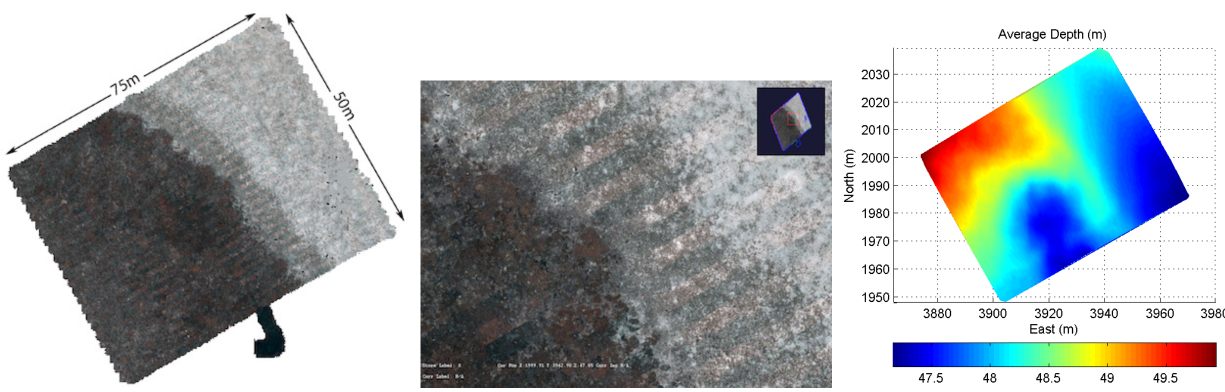


(a) (b) (c)

Fig. 4. Dense AUV grid at Scott Reef off western Australia covering $50m \times 75m$ with 9,831 stereo image pairs. (a) Textured 3D mesh overview of survey site reconstructed using the method outlined in Section II. (b) Close up of transition zone showing dense coral cover, barren sand and an intermediate, partially populated substrate class. (c) Colour map of average mesh depth.

Figure 9, which upon visual inspection appears to approximate the hand-classified output shown in Figure 8 relatively well. The classification performance was quantified using 10-fold cross validation over the entire labelled dataset. Using just rugosity the SVM achieved a 10-fold cross validation prediction accuracy of 71.10%, using just slope, it yielded a cross validation prediction accuracy of 61.73%, and combining both features, a cross validation prediction accuracy of 73.06% was achieved. Reducing the number of classes to three by combining Low relief reef and High relief reef into one class, increased the cross validation prediction accuracy to 91.85%.

From the confusion matrix shown in Table I, it is apparent that the SVM is classifying a large proportion of Low relief reef as High relief reef and a smaller proportion of High relief reef as Low relief reef. This can be attributed to the Low relief reef class including some images which have high variability in rugosity and slope, such as the interface between reef and sand.

The other notable point of confusion for the classifier is the false classification of Ecklonia as High relief reef. This is because the stereo meshes do not represent the reconstructions of the substrate, but rather a reconstruction of the substrate plus benthos. Consequently, using just rugosity and slope causes

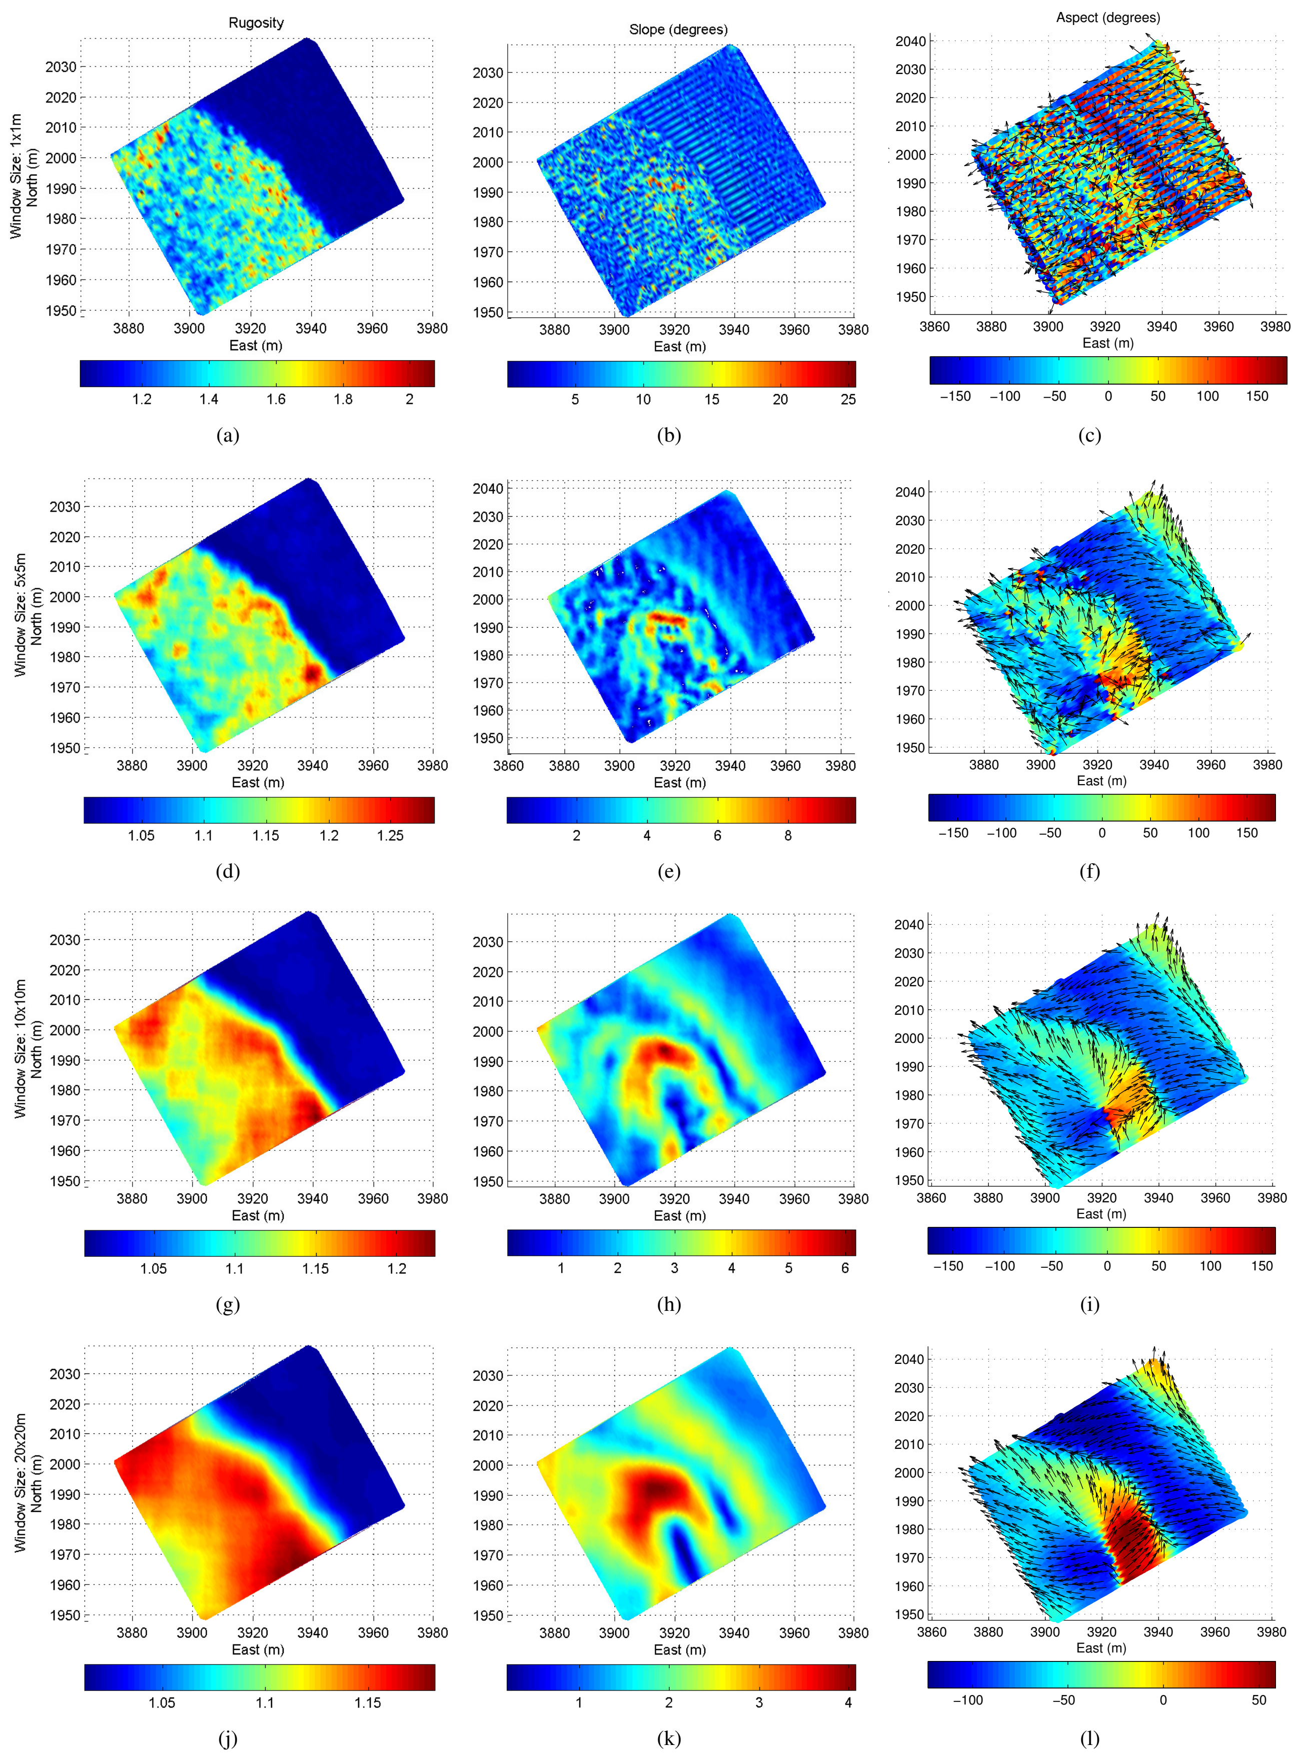


Fig. 5. Dense AUV grid completed in Scott Reef showing the effect of different window sizes on the results. (a), (b) and (c) show rugosity, slope and aspect with a window of $1m \times 1m$. (d), (e) and (f) show rugosity, slope and aspect with a window of $5m \times 5m$. (g), (h) and (i) show rugosity, slope and aspect with a window of $10m \times 10m$. (j), (k) and (l) show rugosity, slope and aspect with a window of $20m \times 20m$.

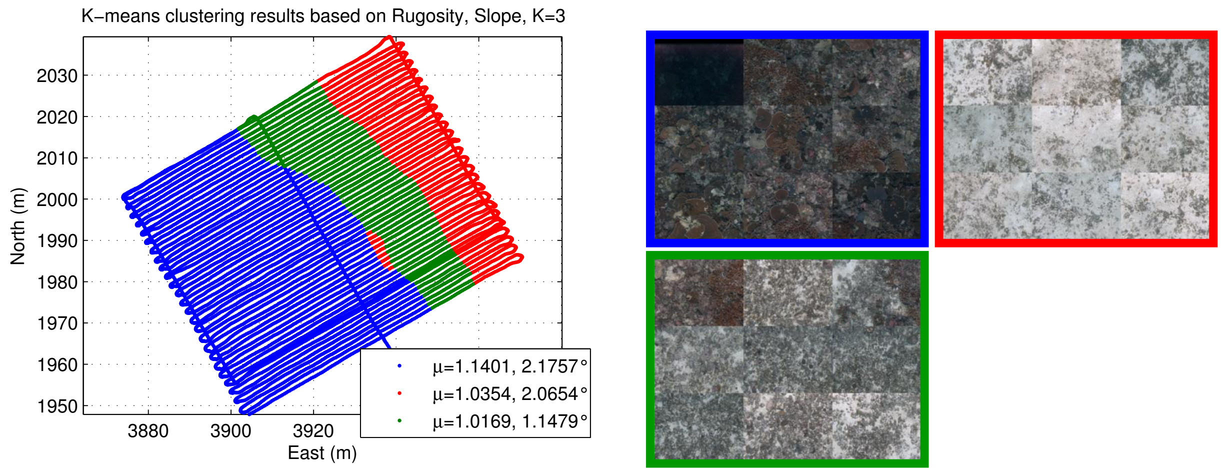


(a)

(b)

Fig. 6. K-means clustering results based on normalised rugosity and slope for dense AUV grid completed in Scott Reef, K=3, window size = $20m \times 20m$. (a) Spatial representation of habitat classes. (b) Sample images for each cluster (Cluster 1 = blue, Cluster 2 = red, Cluster 3 = green).

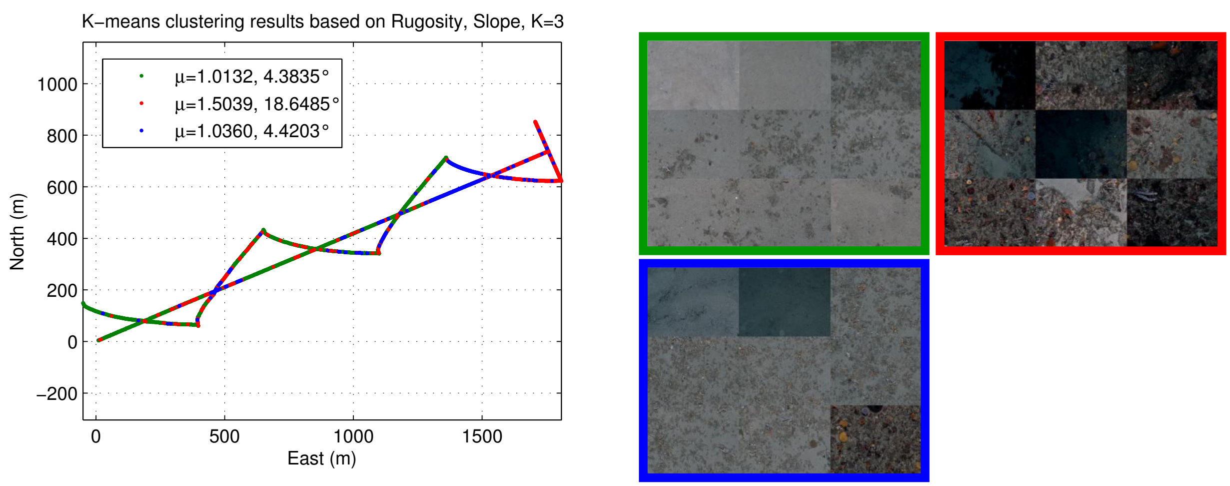


(a)

(b)

Fig. 7. K-means clustering results based on normalised rugosity and slope for the O'Hara Reef dataset, K=3, window size = $1.5m \times 1.2m$. (a) Spatial representation of habitat classes. (b) Sample images for each cluster (Cluster 1 = green, Cluster 2 = red, Cluster 3 = blue).

Ecklonia to be indistinguishable from High-relief reef in some cases, which accounts for some of the miss-classification. This is also apparent in Table II, which shows means and standard deviations of rugosity and slope for hand-labelled classes. The results are mostly intuitive. For example, High relief reef has a higher mean rugosity than Low relief reef which has a higher mean rugosity than Sand/rubble. However, Eklonia has the highest mean rugosity, even though it may occur on flat substrate. Including additional descriptors in the feature vector, such as colour and texture, would significantly improve the results, making it possible to extract more distinct classes from the data.

## VI. Conclusion and Future Work

This paper has demonstrated how multi-scale measures of rugosity, slope and aspect can be derived from fine-scale bathymetric reconstructions created using geo-referenced stereo imagery collected by an AUV. We presented a detailed overview of how rugosity can be derived by considering the area of triangles within a window and their projection onto the plane of best fit, which was found using PCA. Through obtaining the plane of best fit, slope and aspect are calculated with very little extra effort. The results were validated on a simulated surface and the effects of mesh resolution and window size were explored. The technique was demonstrated on real data gathered by the AUV *Sirius* on surveys that cover

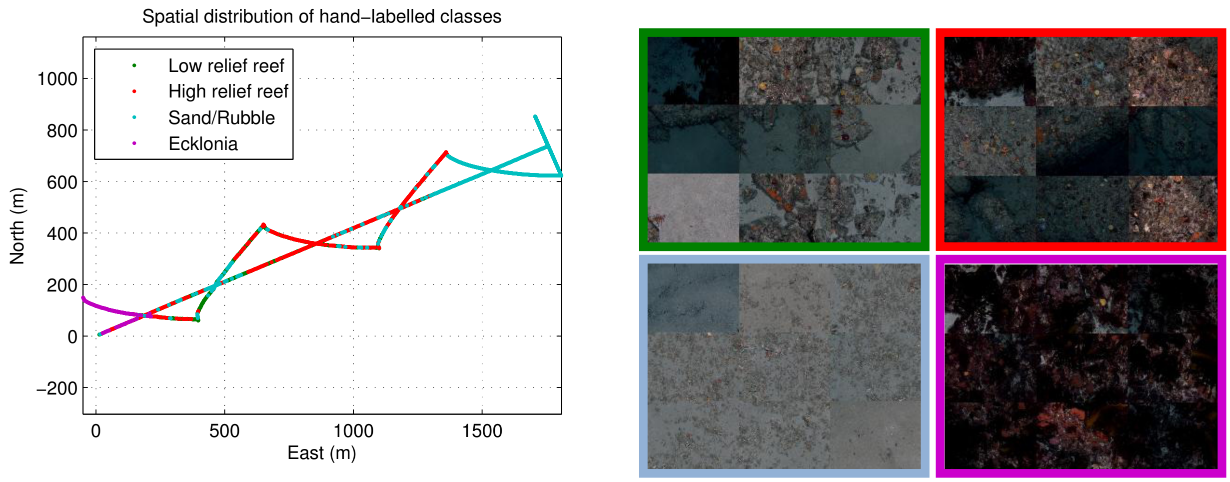


(a) (b)

Fig. 8. Hand-labelled classes for the O'Hara Reef dataset. (a) Spatial representation of habitat classes. (b) Sample images for each hand-labelled habitat class (Low relief reef = green, High relief reef = red, Sand/rubble = blue, Ecklonia = purple).

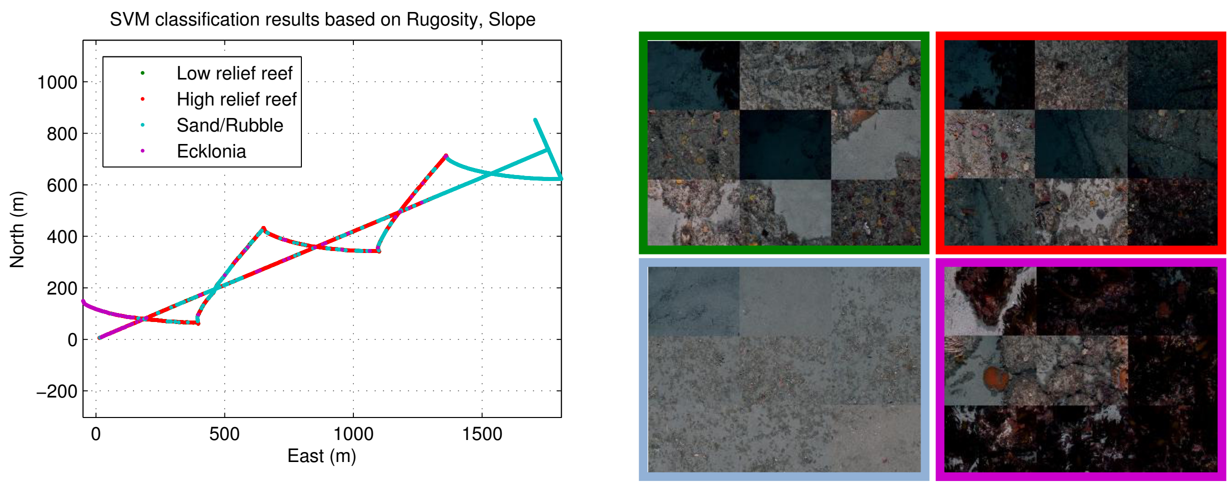


(a) (b)

Fig. 9. SVM classification results based on normalised rugosity and slope for the O'Hara Reef dataset, window size = $1.5m \times 1.2m$. (a) Spatial representation of habitat classes. (b) Sample images for each habitat class (Low relief reef = green, High relief reef = red, Sand/rubble = blue, Ecklonia = purple).

several linear kilometres and consist of thousands of images. The ability to distinguish habitat types based on rugosity and slope was demonstrated through K-means cluster analysis. A human labelled dataset was used to train a SVM classifier that exhibited promising results based on rugosity and slope.

Future work will focus on using these measures along with others as descriptors for more advanced automated habitat classification. The use of stereo imagery to generate the meshes lends itself to combining visual features such as colour and texture which could possibly yield powerful descriptors. Combining slope and aspect with current flow fields inferred using an ADCP may provide another good indicator for benthic cover. Comparing the measures derived from stereo imagery to similar measures derived via acoustic sensors may provide useful discriminatory power in the fact that certain soft organisms such as kelp, are more transparent to acoustics. Given the method's computational tractability, it could prove useful as a virtual 'sensor' to inform adaptive surveying strategies such as delineating zones of significant change in rugosity (i.e. interface between reef and sand or healthy and damaged reef).

## ACKNOWLEDGEMENTS

This work is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government and the Integrated Marine Observing System (IMOS) through the DIISR

TABLE I
CONFUSION MATRIX FOR SVM CLASSIFIER (LRR = LOW RELIEF REEF, HRR = HIGH RELIEF REEF, SR = SAND/RUBBLE, ECK = ECKLONIA).

| | | Predicted | | | |
|---|---|---|---|---|---|
| | | LRR | HRR | SR | ECK |
| Actual | LRR | 1421 | 1255 | 327 | 2 |
| | HRR | 649 | 3040 | 39 | 33 |
| | SR | 48 | 18 | 3609 | 0 |
| | ECK | 20 | 626 | 1 | 117 |

TABLE II
MEANS ($\mu$) AND STANDARD DEVIATIONS ($\sigma$) OF RUGOSITY AND SLOPE FOR HAND-LABELLED CLASSES.

| **Class Name** | **Rugosity** $\mu_r$ | **Rugosity** $\sigma_r$ | **Slope** $\mu_s$ | **Slope** $\sigma_s$ |
|---|---|---|---|---|
| Low relief reef | 1.3994 | 0.2343 | 12.4426° | 9.5041° |
| High relief reef | 1.5711 | 0.2385 | 22.6757° | 13.3236° |
| Sand/rubble | 1.0337 | 0.0524 | 4.4666° | 2.3178° |
| Ecklonia | 1.9780 | 0.3072 | 16.9343° | 9.9628° |

National Collaborative Research Infrastructure Scheme. The authors would like to thank Jan Seiler from the University of Tasmania for providing extensively hand-labelled data for the O'Hara Reef dataset. In addition, we would also like to thank the Australian Institute for Marine Science and the Tasmanian Aquaculture and Fisheries Institute for making ship time available to support this study. The captains and crews of the R/V Solander and R/V Challenger were instrumental in facilitating successful deployment and recovery of the AUV. We also acknowledge the help of all those who have contributed to the development and operation of the AUV, including Duncan Mercer, George Powell, Matthew Johnson-Roberson, Ian Mahon, Stephen Barkby, Ritesh Lal, Paul Rigby, Jeremy Randle, Bruce Crundwell and the late Alan Trinder.

## REFERENCES

[1] M. I. McCormick, "Comparison of field methods for measuring surface topography and their associations with a tropical reef fish assemblage," *Marine Ecology Progress Series*, vol. 112, no. 1-2, pp. 87–96, 1994.

[2] J. A. Commito and B. R. Rusignuolo, "Structural complexity in mussel beds: the fractal geometry of surface topography," *Journal of Experimental Marine Biology and Ecology*, vol. 255, no. 2, pp. 133–152, 2000.

[3] N. J. Frost, M. T. Burrows, M. P. Johnson, M. E. Hanley, and S. J. Hawkins, "Measuring surface complexity in ecological studies," *Limnology and Oceanography: Methods*, vol. 3, pp. 203–210, 2005.

[4] A. Knudby and E. LeDrew, "Measuring structural complexity on coral reefs," in *American Academy of Underwater Sciences 26th Symposium*, Dauphin Island, 2007.

[5] P. Dartnell and J. V. Gardner, "Predicting seafloor facies from multibeam bathymetry and backscatter data," *Photogrammetric Engineering & Remote Sensing*, vol. 70, no. 9, pp. 1081–1091, 2004.

[6] C. D. Moustier and H. Matsumoto, "Seafloor acoustic remote sensing with multibeam echo-sounders and bathymetric sidescan sonar systems," *Marine Geophysical Researches*, vol. 15, no. 1, pp. 27–42, 1993.

[7] J. C. Brock, C. W. Wright, T. D. Clayton, and A. Nayegandhi, "Lidar optical rugosity of coral reefs in biscayne national park, florida," *Coral Reefs*, vol. 23, no. 1, pp. 48–59, 2004.

[8] M. Johnson-Roberson, O. Pizarro, S. B. Williams, and I. Mahon, "Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys," *Journal of Field Robotics*, vol. 27, no. 1, pp. 21 – 51, 2010.

[9] J. S. Jenness, "Calculating landscape surface area from digital elevation models," *Wildlife Society Bulletin*, vol. 32, no. 3, pp. 829–839, 2004.

[10] H. Singh, A. Can, R. Eustice, S. Lerner, N. McPhee, O. Pizarro, and C. Roman, "Seabed auv offers new platform for high-resolution imaging," *EOS*, vol. 85, no. 31, pp. 289, 294–295, Aug. 2004.

[11] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996.

[12] C. M. Bishop, *Pattern Recognition and Machine Learning*, M. Jordan, P. J. Kleinberg, and B. Scholkopf, Eds. Springer, 2006.

[13] T. J. Alexander, N. Barrett, M. Haddon, and G. Edgar, "Relationships between mobile macroinvertebrates and reef structure in a temperate marine reserve," *Marine Ecology Progress Series*, vol. 389, pp. 31–44, 2009.

[14] I. B. Kuffner, J. C. Brock, R. Grober-Dunsmore, V. E. Bonito, T. D. Hickey, and C. W. Wright, "Relationships between reef fish communities and remotely sensed rugosity measurements in biscayne national park, florida, usa," *Environmental biology of fishes*, vol. 78, no. 1, pp. 71–82, 2007.

[15] C.-C. Chang and C.-J. Lin, "LIBSVM: a library for support vector machines," 2001, software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.