

# Towards Autonomous Robotic Coral Reef Health Assessment

Travis Manderson, Jimmy Li, David Cortés Poza, Natasha Dudek, David Meger, and Gregory Dudek

**Abstract** This paper addresses the automated analysis of coral in shallow reef environments up to 90 ft deep. During a series of robotic ocean deployments, we have collected a data set of coral and non-coral imagery from four distinct reef locations. The data has been annotated by an experienced biologist and presented as a representative challenge for visual understanding techniques. We describe baseline techniques using texture and color features combined with classifiers for two vision sub-tasks: live coral image classification and live coral semantic segmentation. The results of these methods demonstrate both the feasibility of the task as well as the remaining challenges that must be addressed through the development of more sophisticated techniques in the future.

## 1 Introduction

In this paper we describe a system for the automated detection and video identification of coral growths using a marine robot. Our objective is to develop a fully autonomous system that can swim over coral reefs in open water, collect video data of live coral formations, and make an estimate of coral abundance. The video is intended for examination by human specialists, but the system needs to be able to both remain resident on the reef surface and recognize coral as it is encountered to perform its mission.

Coral reefs are delicate marine environments of immense importance both ecologically and socio-economically, and yet they are under substantial threat almost everywhere they occur. One preliminary step to retaining these environments is to be able to objectively record their presence, their change over time, and their health. Such records are critical not only to any remediation effort, but also in order to

---

Travis Manderson, Jimmy Li, David Cortés Poza, David Meger and Gregory Dudek at McGill University, School of Computer Science · Natasha Dudek at University of California at Santa Cruz, Dept. of Ecology and Evolutionary Biology

present a compelling case to law makers and law enforcement officials regarding the preservation of these ecosystems. While human divers are commonly deployed to observe reefs and measure their health, the requisite measurements need to be performed using scuba gear under conditions that present a risk to the divers involved.

In the work reported here, we use a small, portable, and high mobility underwater vehicle which is able to swim over the surface of a coral reef, hover in place, navigate in confined spaces, and collect video data from multiple cameras operating simultaneously. In our current experimental configuration the vehicle is accompanied by a human supervisor, but our approach and target scenario does not require a human operator to be present while data is being collected. This vehicle is ideally suited for reef surveillance since it can be deployed manually by a single user either from shore or in the water, does not require an associated tender (ship), can maneuver even in very shallow water, and can even land on a set of legs on sand or a reef surface with limited physical contact. Our approach to covering coral reefs requires the vehicle to be initialized over or near a reef. It can subsequently circumnavigate the reef and cover its interior using inertial navigation. In prior work we have also employed GPS data, acquired by allowing the vehicle to surface, to assist in the navigation task, but in this work navigation is accomplished while remaining underwater at the expense of global localization. This paper does not focus on coverage and navigation, but rather on the system architecture, the nature of the data we collect, and our ability to detect and recognize living coral using this vehicle.

In this paper we propose and evaluate two critical components of the visual processing pipeline used both the guidance and data collection for our vehicle. These operations are the classification of images that are observed as either containing live coral or not, and the subsequent segmentation of the live coral within the image. Several structured data sets used in our evaluation are described below and is available to the community<sup>1</sup>.

## 2 Background

As coral health is an issue of worldwide importance, its monitoring has been studied by many authors previously, both in the field of biology and intelligent systems. This section describes several of the most relevant contributions.

### 2.1 Coral Reef Biology and Reef Health

Coral reefs are majestic structures crucial to ecosystem functioning. They are home to roughly 25% of the oceans' inhabitants, and act as a nursery, feeding ground, and

---

<sup>1</sup> Dataset hosted at: <http://www.cim.mcgill.ca/mrl/data.html>

shelter for thousands of marine organisms [1]. To humans, they represent approximately US\$30 billion annually in goods and services, and are the focus of many studies searching for novel biochemically active drug compounds [2]. Optimistic reports estimate that at the current rate, by 2050 some 75% of the world's remaining reefs will be critically threatened [3]; more pessimistic estimates predict that all of Earth's coral reefs will be dead by the end of the century [4].

Some of the major driving forces behind coral decline worldwide include increasing water temperatures, ocean acidification, increase in frequency and intensity of coral diseases, and damage due to natural disasters such as hurricanes. Many anthropogenic activities are also causing direct harm to reefs, including the overfishing of essential herbivorous species of fish, increasing amounts of water pollution from terrestrial runoff, and increasing sedimentation from coastal construction [3]. Arrival of invasive species can further exacerbate the situation and lead to a dramatic decrease in reef diversity and health, such as the invasion of Indo-Pacific lionfish in the Caribbean Sea and of the crown-of-thorns seastar in Australia [5].

While little can be done on a regional scale about issues such as global warming and increasing ocean temperature, there is an increasing focus on local management and conservation of coral reefs [6]. One critical component of any successful conservation effort is being able to assess whether a particular conservation strategy results in beneficial outcomes on the system in question. In order to protect what remains of the world's coral reefs, it is essential that we design accurate and precise methods to assess the health of coral reefs without undue risk to human participants. This will not only allow us to see when conservation efforts work, but will also help determine which reefs should be conservation priorities and provide evidence to policy makers and the general public that conservation efforts are necessary to preserve the well being of coral reef ecosystems [7].

## ***2.2 Robotic Reef Surveys***

Several research groups have considered the use of autonomous underwater vehicles (AUVs) for data collection in marine environments, and even in coral reefs. Reefs are challenging environments since they are both valuable and physically delicate, and they have complex morphologies. A few vehicles have been developed that can make close approaches to the ocean floor, corals, or aquatic structures [8, 9]. This can be challenging due to several factors: (a) the propulsion systems may be unsafe to operate close to sensitive underwater environments; (b) otherwise "gentle" devices such as gliders have limited maneuverability; (c) it is difficult for humans to produce pre-planned trajectories since sensor feedback underwater is often poor, communications are difficult and terrain models are rarely complete; (d) many propulsion systems are prone to disturbing bottom sediments which reduces visibility.

The problem of designing and controlling stable AUVs has been studied by several authors [10, 11] on a variety of platforms. In prior work with the Aqua class of

vehicles developed in our lab, we have demonstrated a combination of small size, low weight, and high maneuverability with diverse gaits [12, 13].

Several authors have also considered using towed or autonomous surface vehicles to perform visual data collection over marine environments [14], although in the context of coral reefs such an approach is feasible only for the shallowest reef structures and depends critically on very good visibility. Deep water AUVs have been used to map the ocean floor, inspect underwater structures, and measure species diversity [15].

Australia's Integrated Marine Observation System (IMOS) is carrying out a project to deliver precisely navigated time series of seabed imagery and other variables at selected stations on Australia's continental shelf [16]. They are using UAVs to make this endeavor scalable and cost efficient.

In [17], the authors present a structure from motion framework aided by the navigation sensors for building 3D reconstructions of the ocean floor and demonstrate it on an AUV surveying over a coral reef. Their approach assumes the use of a calibrated camera and some drifting pose information (compass, depth sensor, DVL). They use the SeaBED AUV, an imaging platform designed for high resolution optical and acoustic sensing. [18]

In previous work [19] we have developed a controller to allow our vehicle to autonomously move about over coral reef structures using visual feedback. In this paper we restrict our attention to the analysis of the data collected by such a system, and consider the sensing issues that arise.

### 2.3 Visual Coral Categorization

Our methodology has been inspired by recent successes of previous biologically relevant visual data sets. For example, the Fish Task of the recent LifeCLEF contest [20] supported progress on detecting moving fish in video and fish species identification through the release of nearly 20,000 carefully annotated images. The identification of coral using visually equipped AUVs has been studied previously [21]. While we share similar motivations to this work, we differ in deployment and algorithmic objectives. Nonetheless, the relationship is a motivation for the public release of our training and test images which could facilitate comparisons. Additionally, Girdhar *et al.* [22] has demonstrated a system which modifies swimming behavior on-line to follow novel visual content.

## 3 The MRL Coral Identification Challenge

The first contribution of this paper is a robot-collected data set of visual images from environments proximal to a number of coral reefs. This data was collected by the Aqua swimming robot during a series of field deployments in the Caribbean,

where the robot’s existing navigation technologies were exercised to cover each reef and its surroundings. Although our robot did not use vision to inform its navigation strategies during these trials, the images that it collected are representative of the challenge that faces a coral-seeking robot. Therefore, we have organized and annotated them to form two visual challenge tasks: live coral image classification and live coral segmentation. The remainder of this section describes the components of this effort.

### 3.1 Robotic Data Collection

As mentioned previously, operating robots near coral formations requires specialized robot hardware and capabilities. We utilized the Aqua robot [23], an amphibious hexapod that swims using the oscillations of its flippers. Aqua has been designed for use as a visual inspection device and is equipped with four cameras with a variety of properties. A forward-facing stereo pair with narrow field-of-view allows recovery of depth, a front fish-eye camera captures the wider scene, and finally, a 45 degree mirror allows the fourth camera to capture the ocean floor directly below the robot.

In order to achieve broad coverage of the underwater environment, our robot executed a coverage pattern repeatedly over the reef. We set the parameters of this motion by hand so that the robot would pass completely over the reef as well as an equal portion of the sandy surroundings. This gives our data set a roughly equal split between the coral images we target and less desirable content, which poses an interesting classification problem for the visual processing component.

Two attitude strategies were employed, each targeted to induce ideal viewpoints for a different sub-set of Aqua’s cameras. First, a *flat-swimming* maneuver controlled the robot to be aligned with gravity in both the roll and pitch rotational axes. With this attitude, the downward looking camera views the ocean bottom with an orthogonal viewpoint and the front fish-eye camera views the horizon at roughly half the image height. Second, we considered swimming with a downwards pitch of thirty degrees. This strategy allowed the narrow-view stereo pair to view the ocean bottom slightly in front of the robot. The depths observed at this angle would allow fixed-altitude operations, which are desirable in order to prevent accidental collisions with the coral.

The robot executed five data collection runs at four distinct reef locations (one reef was visited twice). We selected reefs within the Folkstone Marine Preserve and in Heron Bay, both of which are located on the western coast of St. James, Barbados. During each run, the robot covered an area of approximately 100 square meters. Each reef location was an instance of the spur-and-groove coral formations that tend to present the widest range of diversity of coral species, and are thus ideal regions for collection of biologically relevant data.

## Data Statistics

All of the videos are taken at 15 frames per second, with VGA resolution. The total size of visual data collected over the five collection runs is 104 gigabytes consisting of 164 minutes of video. Depth and IMU data are also recorded throughout.

### 3.2 Data Annotation

A marine biologist manually annotated the coral within a subset of the images we collected. The results of this annotation have been made available in a standardized format, and the data is being released publicly for the purposes of comparison of results and classifier training. As a variety of tasks can be considered, depending on the goals of the robot platform, we define two coral-related visual tasks and accompanying evaluation criteria. We continue by describing our annotation procedure.

#### Annotation for Image Classification

The first sub-task that we define is coral image classification. Given an image, the system outputs whether there is live coral in the image. To create training and testing data for this task, we extracted images at 5 second intervals from all of the videos taken by the downward-looking mirrored camera while the robot was swimming flat. Each image was then subdivided into four  $320 \times 256$  quadrants to limit the diversity and facilitate ease of labeling. The biologist labeled 3704 images into one of three categories:

- **Yes:** There is live coral in the image
- **No:** There is no live coral in the image
- **Reject:** The image should be discarded because it is too difficult to tell whether there is live coral or not. This could be because the image is too blurry or the coral is too small to see clearly.

This provided us with 1087 **Yes** images, 2336 **No**, and 281 **Reject** images. Figure 1 shows some examples of **Yes** and **No** images.

#### Annotation for Segmentation

Secondly, we define the coral segmentation task, where the coral regions within an image must be identified, through creation of a coral mask. While some existing segmentation data sets contain pixel-wise ground truth, we lacked the resources to produce this detailed data. Instead, we have manually annotated rectangular coral regions for each of the 1087 **Yes** images from our classification data set. Examples of the selected image regions are shown in Figure 2. Rectangular regions cause a

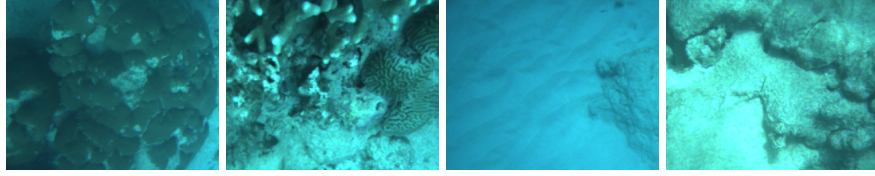


Fig. 1: Annotated images used for training a detector for the image classification task. The left two images are labeled as having coral and the right two images are labeled as not containing coral.

small approximation error at region boundaries, but this task is still a reliable proxy for coral segmentation, as will be demonstrated in our results section.

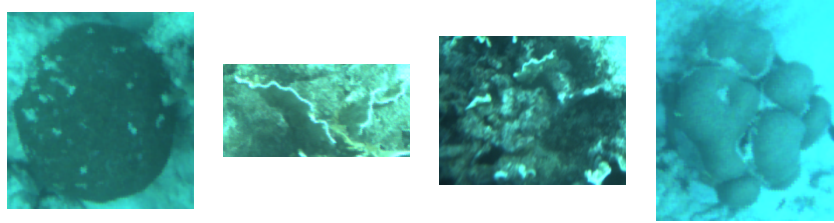


Fig. 2: Positive training images cropped to contain only coral, which is useful for training a detector for the coral segmentation task.

### Annotation Statistics

The final annotated data set produced by our labeler was reduced in size from the raw robot footage due to the rejection of poor quality and ambiguous images. We separated the annotated data into a training set (416 positive examples and 701 negative examples) and a test set (492 positive examples and 1544 negative examples). The training set contains images from three data collection runs at three unique reefs, and the test set contains images taken from two data collection runs at the fourth reef location. Thus, there is no overlap between the training and test sets.

We have additionally defined evaluation protocols for the use of this data, following best-practices from existing challenges such as the ILSVRC [24]. Broadly, we measure performance on each binary categorization task as prediction accuracy, normalized by the data set size. For the categorization task this represents the number of images, and for segmentation this is measured in image area. Methods cannot be optimized directly on the test data set. Rather, parameters should be refined by splitting the training set into folds and then the performance reported after only a single run on the test set. This data is being released to the public alongside with

this paper and we will maintain a record of the best performing techniques over time as other authors attempt the task. We now continue by describing several baseline techniques that we have developed.

## 4 Method

Coral identification in the ocean shares many of the typical challenges that face terrestrial vision systems, as well as several unique to this task. The lighting conditions in the shallow ocean include caustics caused by the water’s surface, inter-reflections and the absorption of low-frequency colors. This makes brightness invariance essential. The robot changes its orientation during the survey, which implies the need for orientation invariance. Small floating particles are ubiquitous in the underwater domain, causing an optical snow effect. Additionally, the appearance of the coral itself has a wide diversity and there are local variations between reef locations, so generalization must be the focus of learned methods.

In the face of these challenges, our approach to identification of coral is to encode the visual data in robust feature representations that capture canonical appearance properties of coral, such as its color and texture and to learn coral classifiers from training data on top of these features. We develop two processing streams – one for each of the visual tasks described above. Our classification process employs Gabor functions and global processing to compute aggregate statistics. Segmentation is achieved through local computations on sub-regions of the image. Each approach will be described in detail in the remainder of this section.

### 4.1 *Global Image Statistics for Coral Classification*

The classification pipeline uses both global color and aggregate texture features in a classifier subsystem to learn from labeled example images and subsequently predict whether an image contains live coral. This subsystem computes 2 types of attributes over the entire (global) images to produce a characteristic feature vector. These vectors are then classified using a support vector machine (SVM) trained with our manually classified data. Figure 3 (top) illustrates the classification pipeline.

Our method represents texture through the use of the well-known Gabor transform. The Gabor function [25] is a sinusoid occurring within a Gaussian envelope and has inspired a class of image filters particularly suited to describing texture [26]. Our method automatically selects a sub-set of Gabor wavelets from a large family by selecting those with frequency and spatial support parameters that optimize task performance, using cross-validation on the training set.

Applying filters results in a stack of transformed images and we extract robust energy statistics from these in order to produce vector suitable for classification. The amplitude histogram of each Gabor filter provides a characterization of the image

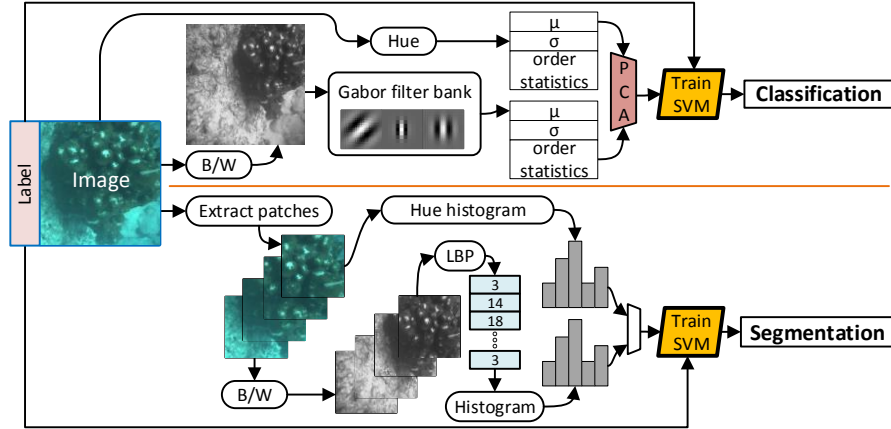


Fig. 3: Image processing pipeline. (top) Gabor-based classification. (bottom) Local binary pattern (LBP)-based segmentation.

content including the presence of outlier objects. Order statistics can effectively characterize such a signal [27] and are robust to much of the noise present in our task. For this reason we characterized the energy distribution with several statistics of each Gabor filter: the mean energy, the variance of the energy distribution, and the energy at a specific set of percentiles of the cumulative distribution (5th, 20th, 80th and 95th percentiles). In order to capture color information, we additionally extract the same robust statistics for the distribution of hue values observed in the image.

The result of both the Gabor texture filters and the color summary are concatenated into a fixed-length vector. Depending on the number of active Gabor components, this representation can have between 24 and several hundred dimensions. In order to reduce computation and simplify the learning, we perform principal components analysis on these vectors to find the subspace that captures 99.99 *per cent* of the variance.

The final step in this pipeline is to predict the label of an image (live coral or not). We learn a support vector machine (SVM) from the training images described previously, and apply the resulting learned model to make coral predictions on new images.

## 4.2 Local Binary Pattern based Coral Segmentation

Our coral segmentation pipeline uses LBPs [28] and color information as image descriptors, and a support vector machine to detect whether small patches of the images correspond to live coral or not. Unlike the Gabor filters, which are applied

globally, our features and classifier are applied on small image patches, which allows fine-grained segmentation of coral regions. Figure 3 (bottom) illustrates the segmentation pipeline.

For a given pixel in the image, its LBP is computed by comparing its gray level  $g_c$  with that of a set of  $P$  samples in its neighborhood,  $g_p$  ( $p = 1, 2, \dots, P$ ). These samples are evenly spaced along a circle with radius  $R$  pixels, centered at  $g_c$  (see Figure 4). For any sample that doesn't fall exactly in the center of a pixel, its gray value is estimated by interpolation. The LBP is computed according to

$$LBP_{P,R} = \sum_{p=0}^{P-1} \mathbb{1}_{\{g_c - g_p \geq 0\}} 2^p, \quad (1)$$

where  $\mathbb{1}_{\{\cdot\}}$  is the indicator function.

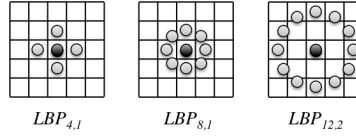


Fig. 4: Local binary pattern neighbor sets for  $(P = 4, R = 1)$ ,  $(P = 8, R = 1)$  and  $(P = 12, R = 2)$ .

To achieve rotational invariance, Ojala et al. [28] proposed to label the LBPs according to their number of 0/1 transitions. LBPs with up to 2 transitions are called *uniform* and they are assigned a label corresponding to the number of 1's in the pattern. LBPs with more than 2 transitions are called *nonuniform* and they are all assigned the label  $P + 1$ . Finally, the rotation invariant LBP image descriptor is a  $P + 2$  bin histogram of these labels computed across all pixels in the image. Uniform patterns are assigned to unique bins, while nonuniform patterns are all assigned to a single bin. As color is also an important feature for coral segmentation, we appended the LBP histogram with an 8 bin histogram of the hue values of the pixels in the image patch.

During robot operation time, our learned model is used to segment an image by splitting it into patches with the same size as those used during training. Features are extracted from each patch and scored with the SVM, producing a coral segmentation mask that can be used to guide the robot during its mission.

## 5 Experiments and results

### 5.1 Global Coral Classification

Our global classifier was tested on the data sets above using distinct testing and training sets collected over different reefs. We were able to achieve a net classifier accuracy of 89.9% on balanced sets of images containing coral and not containing coral. This accuracy generally increased with the number of Gabor basis functions, however since these are the primary source of computational cost, we are interested in a compromise between performance and the number of filters user. The trade-off between accuracy and the size of the filter bank is illustrated in Fig. 5. While using a bank of 24 or more filters provides maximal performance, the 80.6% rate achieved with just 20 filters appears quite acceptable for our applications.

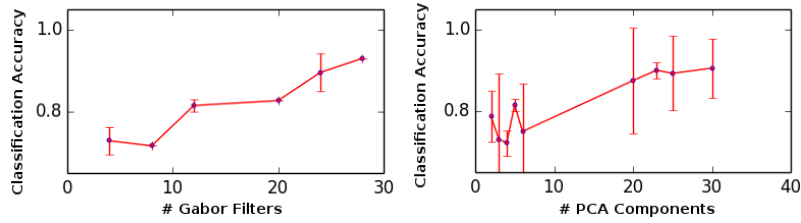


Fig. 5: Classification accuracy increases with both: (left) number of Gabor filters; and (right) number of PCA components. This reflects the trade-off of computational effort and performance.

### 5.2 LBP-based Coral Segmentation

To study the effect of varying the number of points and radius ( $P, R$ ) of the LBPs and the size of the patches on the segmentation, we performed a grid search on these parameters. Also, to optimize the performance of the support vector machine (SVM), we ran a grid search on the gamma, tolerance and regularization constant ( $C$ ) parameters of the radial basis function (RBF) kernel.

The LBP parameters had very small impact on the accuracy of the classifier. We tested over the values  $(P, R) = (8, 1), (16, 2), (16, 3), (24, 3), (32, 5)$  and found that the difference in accuracy between them was less than 2.1%, regardless of the patch size. Given such a small impact, we decided to use  $(P = 8, R = 1)$  for the remaining experiments.

The patch size, on the other hand, had a much larger impact on the classification accuracy, which is illustrated in Figure 6. The maximum classification accuracy

achieved was 81.16% with the RBF kernel parameters set to  $\gamma = 0.0001$ ,  $tol = 2$  and  $C = 10000000$ . The optimal patch size was found to be 30 pixels.

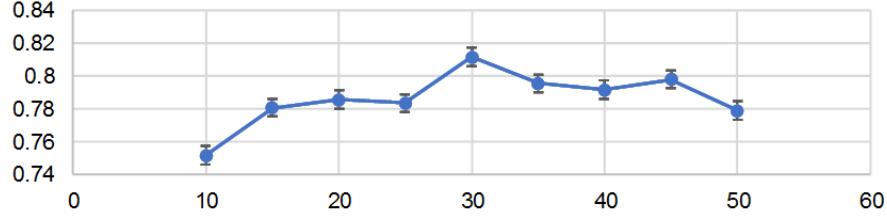
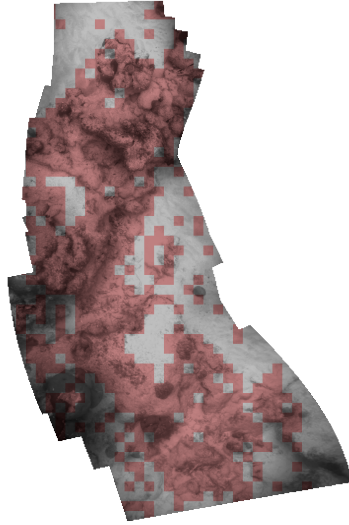
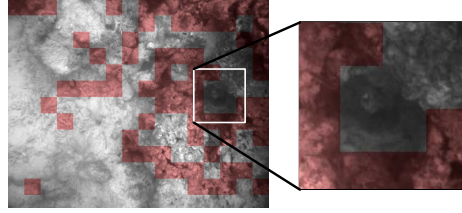


Fig. 6: Classification accuracy vs. patch size (pixels)

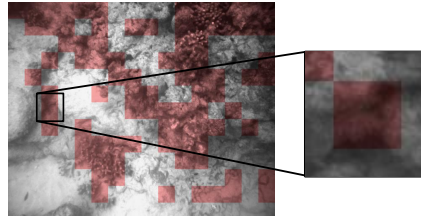
In Figure 7, we present some examples of images from the test set with an overlay (in red) showing the segmented live coral. Figure 7a is a stitched image created from several consecutive frames from the original video. We observe that the segmentation pipeline correctly finds areas of the image with live coral. We also observe areas where the classifier has problems detecting coral such as when the texture is uniform (with an example of a false negative shown in Figure 7b). Likewise, live coral can be incorrectly detected when variations in texture (or shadows) match that of live coral (with an example of a false positive shown in Figure 7c).



(a) Reef 3 segmentation



(b) Live coral segmentation. (right) false negative



(c) Live coral segmentation. (left) false positive

Fig. 7: Samples of live coral segmentation

## 6 Discussion

We have described a robot-vision system for performing automated coral surveys of the sea floor. We learn coral predictors that are able to robustly detect live coral patches and segment them from the background, agreeing with the assessments of an experienced coral biologist with an accuracy of 80 to 90 *per cent*. These results are based on a data set of thousands of labeled images of only moderate quality, confounded by the typical phenomena that confront any diver or AUV. Our data set is being made available in conjunction with this submission.

In future, we plan to study the disambiguation of coral from other zooxanthellae-containing organisms and on the automated labeling of different coral subspecies. This will require suitably labeled training data, as well as more diverse raw data sets, potentially including active illumination. Additionally, we hope to integrate coral mapping into the navigation stack of our vehicle, as we have successfully done in the past with other vision-guided navigation methods [22]. The resulting system has the potential to perform autonomous longitudinal surveys, providing biologists with an easy, quick, and accurate way of monitoring reef health. Such methods are critical for understanding how these ecosystems respond to environmental disturbances, documenting the efficacy of novel coral reef conservation and restoration efforts, and convincing policy makers to enact stringent protection measures for coral reef ecosystems.

## References

1. M. D. Spalding, C. Ravilious, and E. P. Green, *United Nations Environment Programme, World Conservation Monitoring Centre. World Atlas of Coral Reefs*. University of California Press: Berkeley, 2001.
2. "Millennium ecosystem assessment: Ecosystems and human well-being - Synthesis report," World Resources Institute, Washington, DC, 2005.
3. L. M. Burke, K. Reyter, M. Spalding, and A. Perry, *Reefs at risk revisited*. World Resources Institute Washington, DC, 2011.
4. O. Hoegh-Guldberg, P. J. Mumby, A. J. Hooten, R. S. Steneck, P. Greenfield, E. Gomez, C. D. Harvell, P. F. Sale, A. J. Edwards, K. Caldeira, N. Knowlton, C. M. Eakin, R. Iglesias-Prieto, N. Muthiga, R. H. Bradbury, A. Dubi, and M. E. Hatziolos, "Coral reefs under rapid climate change and ocean acidification," *Science*, vol. 318, no. 5857, pp. 1737–1742, 2007.
5. M. A. Albins, M. A. Hixon, *et al.*, "Invasive indo-pacific lionfish *pteris volitans* reduce recruitment of atlantic coral-reef fishes," *Marine Ecology Progress Series*, vol. 367, pp. 233–238, 2008.
6. D. R. Bellwood, T. P. Hughes, C. Folke, and M. Nystrom, "Confronting the coral reef crisis," *Nature*, vol. 429, no. 6994, pp. 827–833, 06 2004.
7. C. R. Margules and M. B. Usher, "Criteria used in assessing wildlife conservation potential: A review," *Biological Conservation*, vol. 21, no. 2, pp. 79 – 109, 1981.
8. J. Eskesen, D. Owens, M. Soroka, J. Morash, F. Hover, C. Chrysostomidis, J. Morash, and F. Hover, "Design and performance of odyssey iv: A deep ocean hover-capable auv," MIT, Tech. Rep. MITSG 09-08, 2009.
9. M. Woolsey, V. Asper, A. Diercks, and K. McLetchie, "Enhancing niust's seabed class auv, mola mola," in *Proceedings of Autonomous Underwater Vehicles*, 2010, pp. 1–5.

10. S. Mohan and A. Thondiyath, "A non-linear tracking control scheme for an under-actuated autonomous underwater robotic vehicle," *International Journal of Ocean System Engineering*, vol. 1, no. 3, pp. 120 – 135, 2011.
11. D. Font, M. Tresanchez, C. Siegentahler, T. Pallej, M. Teixid, C. Pradalier, and J. Palacin, "Design and implementation of a biomimetic turtle hydrofoil for an autonomous underwater vehicle," *Sensors*, vol. 11, no. 12, pp. 11 168–11 187, 2011.
12. D. Meger, F. Shkurti, D. Cortes Poza, P. Giguere, and G. Dudek, "3d trajectory synthesis and control for a legged swimming robot," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. IEEE, 2014, pp. 2257–2264.
13. D. Meger, J. C. G. Higuera, A. Xu, and G. Dudek, "Learning legged swimming gaits from experience," in *International Conference on Robotics and Autonomous Systems (ICRA)*, 2015.
14. S. Williams, O. Pizarro, M. Johnson-Roberson, I. Mahon, J. Webster, R. Beaman, and T. Bridge, "Auv-assisted surveying of relic reef sites," in *OCEANS 2008*, Sept 2008, pp. 1–7.
15. T. Maki, A. Kume, T. Ura, T. Sakamaki, and H. Suzuki, "Autonomous detection and volume determination of tubeworm colonies from underwater robotic surveys," in *OCEANS 2010 IEEE - Sydney*, May 2010, pp. 1–8.
16. S. Williams, O. Pizarro, M. V. Jakuba, C. R. Johnson, N. S. Barrett, R. C. Babcock, G. A. Kendrick, P. D. Steinberg, A. J. Heyward, P. Doherty, I. Mahon, M. Johnson-Roberson, D. Steinberg, and A. Friedman, "Monitoring of benthic reference sites using an autonomous underwater vehicle," *IEEE Robotics and Automation Magazine*, vol. 19, no. 1, pp. 73–84, 2012.
17. O. Pizarro, R. M. Eustice, and H. Singh, "Large area 3-D reconstructions from underwater optical surveys," *IEEE Journal of Oceanic Engineering*, vol. 34, no. 2, pp. 150–169, 2009.
18. H. Singh, R. Armstrong, F. Gilbes, R. Eustice, C. Roman, O. Pizarro, and J. Torres, "Imaging coral i: Imaging coral habitats with the seabed auv," *Subsurface Sensing Technologies and Applications*, vol. 5, no. 1, pp. 25–42, 2004.
19. P. Giguere, G. Dudek, C. Prahacs, N. Plamondon, and K. Turgeon, "Unsupervised learning of terrain appearance for automated coral reef exploration," in *Computer and Robot Vision, 2009. CRV '09. Canadian Conference on*, May 2009, pp. 268–275.
20. C. Spampinato, S. Palazzo, B. Boom, and R. B. Fisher, "Overview of the lifeclef 2014 fish task," in *Working Notes for CLEF 2014 Conference, Sheffield, UK, September 15-18, 2014.*, 2014, pp. 616–624. [Online]. Available: <http://ceur-ws.org/Vol-1180/CLEF2014wn-Life-SpampinatoEt2014.pdf>
21. M. Johnson-Roberson, S. Kumar, and S. Willams, "Segmentation and classification of coral for oceanographic surveys: A semi-supervised machine learning approach," in *OCEANS 2006 - Asia Pacific*, May 2006, pp. 1–6.
22. Y. Girdhar, D. Whitney, and G. Dudek, "Curiosity based exploration for learning terrain models," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 578–584.
23. G. Dudek, M. Jenkin, C. Prahacs, A. Hogue, J. Sattar, P. Giguere, A. German, H. Liu, S. Saunderson, A. Ripsman, S. Simhon, L. A. Torres-Mendez, E. Milios, P. Zhang, and I. Rekleitis, "A visually guided swimming robot," in *Proceedings of Intelligent Robots and Systems*, Edmonton, Alberta, Canada, August 2005.
24. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," 2014.
25. D. Gabor, "Theory of communication," *Journal of the IEE*, vol. 93, pp. 429 – 457, 1946.
26. I. Fogel and D. Sagi, "Gabor filters as texture discriminator," *Biological Cybernetics*, vol. 61, no. 2, pp. 103–113, 1989. [Online]. Available: <http://dx.doi.org/10.1007/BF00204594>
27. P. Meer, J. Jolion, and A. Rosenfeld, "A fast parallel algorithm for blind estimation of noise variance," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 12, no. 2, pp. 216–223, 1990.
28. T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 7, pp. 971–987, Jul 2002.