SPECTRAL UNMIXING AND MAPPING OF CORAL REEF BENTHIC COVER

Rohan Zeng[1], *Eric J. Hochberg*[2], *Alberto Candela*[3], *and David S. Wettergreen*[1]

[1]The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, United States
[2]Bermuda Institute of Ocean Sciences (BIOS), St.George's, Bermuda
[3]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, United States

ABSTRACT

Coral reefs are an important ecosystem to the local communities and indigenous wildlife that rely on them. However, reefs have greatly degraded in recent decades with the remaining at increasing risk of loss. Quantitatively mapping these reefs would provide a resource for us to monitor changes and understand their health. We explore methods leveraging limited spectral data and resources for efficient global scale modeling of coral reefs. We then evaluate performance on a Deep Neural Network and our previously developed Deep Conditional Dirichlet Model. Regions of high uncertainty based on the model output prediction are used to determine informative *in situ* sampling. An ergodic planner is implemented to generate a path through these regions to acquire samples that best improve the coral map. The result is a resource efficient learning based pipeline that augments existing spectral data and maps coral reefs globally to improve our understanding of their condition.

Index Terms— coral reef, unmixing, remote sensing, limited data, ergodic planning

1. INTRODUCTION

Coral reefs play a vital role in supporting both local communities and indigenous wildlife [2, 3]. Reefs worldwide are in critical and rapid decline, with 33-55% of reefs degrading over recent decades and about 35% of the remaining reefs at risk of loss over the next few decades [4, 5]. Despite these rising issues, only 0.1% of the world's reefs have been studied quantitatively [6].

To address this issue, we aim to create a valuable resource for monitoring changes and assessing the health of coral reefs by quantitatively mapping them. Data is often expensive to process, so we explore how previously developed machine learning methods can make efficient use of limited spectral data and resources to model coral reefs on a global scale. In order to improve upon model predictions, *in situ*, or on site, expeditions can validate and verify the output.

We develop upon our prior work to demonstrate how clustering-based approaches allow our applied classification and unmixing models to perform comparably to prior results even on limited and biased datasets [7, 8, 9]. After classification, we apply informative path planning to determine efficient sampling locations for *in situ* expeditions to validate and improve coral models. Thus, we present a process to classify and unmix spectral data to identify coral reefs (Figure 1).

2. APPROACH

Mapping of coral reefs has typically consisted of *in situ* missions. However, these are costly expenditures and are difficult to reproduce over time. The introduction of remote sensing for coral mapping made maps of greater scale and consistency possible. We seek the ability to produce benthic classifications or bathymetric maps from remotely sensed spectral data. This is significant because the benthic cover of reef ecosystems has proven to be an effective metric for monitoring coral reef health. By characterizing the spectral reflectance of reef benthic communities using remote sensing, different benthic components can be classified and mapped over large areas [10, 11].

Our prior work focused on mapping the individual distributions of the benthic cover classes of coral, algae, and sand. We utilized labeled Portable Remote Imaging SpectroMeter (PRISM) data of Heron Island, aligned the labels with input high resolution Worldview-3 and low resolution Landsat-8 data, then trained Support Vector Machines (SVM), Deep Neural Networks (DNN), and Deep Conditional Dirichlet Models (DCDM) through supervised learning [12, 13, 9]. The DCDM has the same architecture as the DNN for classification while also incorporating a Dirichlet distribution based loss function to support unmixing. In our current study, we train the same input data on updated PRISM labels that exhibit class imbalances and limited data representation (Figure 1b). We approach this problem by applying different clustering and balancing methods on the dataset, such as k-means++ clustering, to achieve fair representation for all classes [14]. The ability to train models on limited amounts of data and successfully classify coral regions would reduce costs while retaining global coverage and efficiency for mapping.

After the spectral data has been processed, we seek to improve the classification and unmixing confidences. We choose

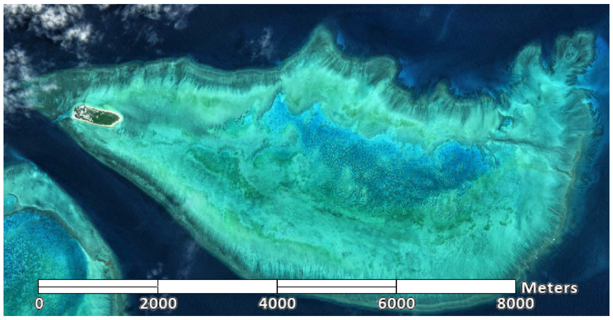

(a) Heron Island, small in upper left, and surrounding coral reef [1].

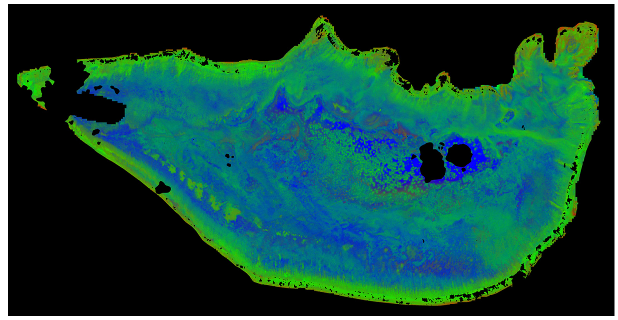

(b) Labeled PRISM data, with Coral (Red), Algae (Green), and Sand (Blue). These labels are biased and imbalanced, with 56.93% algae, 42.55% sand, and just 0.5296% coral.

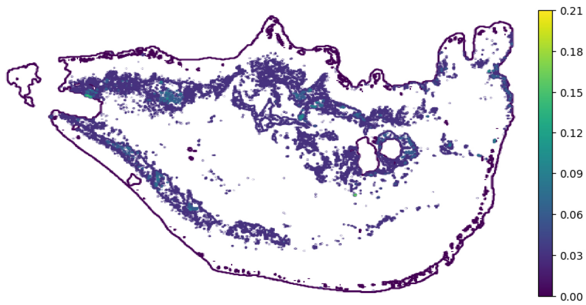

(c) Variance (uncertainty) information prior derived from the trained DCDM model outputs.

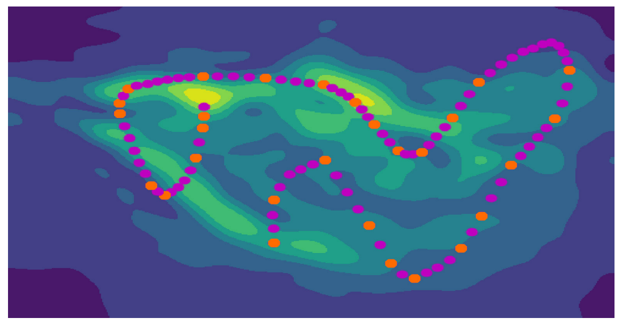

(d) Planner optimized trajectory (magenta) with sparse sampling locations (orange) on DCDM information prior.

Fig. 1: Coral detection and sampling plan on Heron Island. Models are trained using spectral data of the scene alongside labeled data. The variance in the model outputs is used as an information prior for the ergodic planner. The planner yields an information maximizing trajectory for *in situ* expeditions to efficiently validate and refine the model.

to supplement the model predictions with *in situ* expeditions. However, these expeditions are still costly expenditures. By leveraging an approximate knowledge of the map from the model outputs and quantifying the classification variance, we can determine a more efficient trajectory for the *in situ* expedition that makes the best use of provided resources.

We first quantify areas of high uncertainty, or variance, in the model predictions. The weighted sum of these variances per class constitutes our information prior. This determines regions with the highest potential information gain from the predictions (Figure 1c). We then apply an ergodic planner which takes its information prior from the model to optimize the information gain to expedition resource and cost trade off [15, 16].

The uncertainty based information prior defines the likelihood of generating informative measurements at any point within the map. Within ergodic search, the planner minimizes the difference between the Fourier spectral decomposition of this expected information distribution and the agent's, or the planned trajectory's, time-averaged statistics in order to optimize its trajectory. In other words, the objective of the planner is to reduce the uncertainty in the agent's map by optimizing the dynamic behavior of the agent [15]. The resultant trajectory spends time in regions proportional to the expected information gain, which makes more efficient use of resources than an uninformed trajectory such as a random walk or grid-based coverage [16].

In turn, *in situ* measurements can be taken at optimized locations to verify and update regions where the model is most uncertain, thus increasing the efficiency of the expedition. In addition to the agent dynamics, a sampling variable is incorporated into the optimization problem to determine where sensing measurements should be taken along the trajectory. An L_1 metric is used to regularize the sampling variable in order to promote sparsity during optimization [17] (Figure 1d).

The evaluation of the process is quantified by three primary metrics. Accuracy helps us understand how well the models are at determining the hard classifications of the dataset. The Kullback-Leibler Divergence (KLD) and the Mean Absolute Error (MAE) serve as an evaluation of how well the model is able to match the ground truth compositions, specifically how well the model is able to unmix the dataset.

Data	Model	Accuracy	KLD	MAE
Landsat	SVM	72.68%	—	—
Landsat	DNN	72.07%	0.0393	0.1867
Landsat	DCDM	72.22%	0.0419	0.1878
WorldView	SVM	69.79%	—	—
WorldView	DNN	70.38%	0.0487	0.2054
WorldView	DCDM	69.91%	0.0512	0.2056

(a) Model Performance on All Classes

Data	Model	Accuracy	KLD	MAE
Landsat	SVM	0%	—	—
Landsat	DNN	25.60%	0.0937	0.3078
Landsat	DCDM	22.34%	0.1015	0.3163
WorldView	SVM	0%	—	—
WorldView	DNN	6.170%	0.1127	0.3535
WorldView	DCDM	22.14%	0.1025	0.3215

(b) Model Performance on Coral Class

Table 1: DNN and DCDM performance on data clustered through k-means++ clustering. Performance is given on classifying all classes in the benthic cover as well as on just the coral class as the primary class of interest.

3. MACHINE LEARNING RESULTS

Despite the limited and biased dataset with minimal coral representation, we found that the SVM as a standard machine learning approach does give good accuracy overall however, this is due to the SVM classifying all pixels as algae or sand without any coral classifications. The high abundance of the former two classes in the dataset even after clustering allows the SVM to attain a high accuracy with 0% coral classifications, as seen in Table 1a and Table 1b.

Between a DNN model and DCDM trained on the same dataset, the performance across all metrics appear to be roughly similar with regards to all of the classes (Table 1a). The DNN and DCDM accuracies are quantitatively similar, though in all other metrics the DCDM appears to perform slightly worse in comparison to the DNN model. Landsat achieves an accuracy of approximately 72.15% while WorldView performs slightly worse with an accuracy of about 70.15%. These results all seem to contradict the improvement in performance with high spatial-resolution WorldView data alongside the unmixing capabilities of the DCDM observed in our previous research [9]. However, this is considering evaluation on all of the classes together, where both algae and sand dominate the test set.

Differences in performance can be observed with the DNN versus DCDM and Landsat versus WorldView when examining only the coral test set (Table 1b). The DNN has an average accuracy of 15.89% while the DCDM has an average accuracy of 22.24%, thus the DCDM offers more consistent performance in classifying and unmixing coral. The Landsat models still tend to outperform the WorldView models. However, between the DDN and DCDM, WorldView shows increased accuracy by almost 3 times while also decreasing its KLD and MAE which is indicative of successful unmixing. Though the evaluation on all of the classes together do not appear to meet our expectations, there is supportive evidence for improved performance as a result of the higher spatial-resolution data and the DCDM together. This is demonstrated by the successful isolation and improvement of classifications on the rare coral class.

While distinctions are observed between models and datasets, performance across the board is quantitatively quite close. More data is needed to establish statistical significance for these conclusions, however notably we have verified that comparable results to previous contributions, with accuracies of around 70-75%, are observed despite dataset limitations. Additionally, though this accuracy may appear low, this is expected due to the difficulty of unmixing spectral data compared to simple classification tasks.

4. PATH PLANNING RESULTS

After spectral processing, we apply the ergodic planner to determine efficient sampling locations for *in situ* expeditions (Figure 1d). We evaluate the model performance pre planning as well as the updated model performance post planning. The pre-planning model is the average model predictions from spectral processing. The post-planning model is calculated by using the ground truth labels and the information maps from the ergodic planner to update the pre-planning model. The greater the information gained during planning, the closer the post-planning model matches the ground truth labels.

We observed a general increase in performance across all metrics when using information priors based on the model predictions compared to random priors to simulate updates on the model output. This was particularly improved on classifying coral as the primary class of interest. This suggests that ergodic based *in situ* expeditions would definitely improve benthic cover mapping compared to random traversal. When evaluating the ergodic planner, we considered the performance of the models pre and post planning on all classes as well as on just the coral class (Table 2).

We found that the DNN outperformed the DCDM on accuracy before and after the simulated update, though the DCDM was able to achieve lower, and thus better, KLD and MAE scores after being updated (Table 2a). This indicates the results more closely match the distribution of ground truth benthic class mixtures, which suggests that the DCDM is able to more intelligently sample locations that improve

Model	Metric	Pre-Planning	Post-Planning
DNN	Accuracy	70.36%	92.97%
DCDM	Accuracy	69.59%	92.62%
DNN	KLD	0.1331	0.0399
DCDM	KLD	0.1337	0.0385
DNN	MAE	0.4556	0.1847
DCDM	MAE	0.4557	0.1790

(a) Model Performance on All Classes

Model	Metric	Pre-Planning	Post-Planning
DNN	Accuracy	71.49%	90.50%
DCDM	Accuracy	71.70%	93.03%
DNN	KLD	0.1397	0.0180
DCDM	KLD	0.1380	0.0144
DNN	MAE	0.4840	0.0984
DCDM	MAE	0.4793	0.0876

(b) Model Performance on Coral Class

Table 2: Pre-planning and post-planning model performance for the DNN and the DCDM. The pre-planning model is the average model predictions from spectral processing. The post-planning model is the updated model using information from the ergodic planner. Performance is given on classifying all classes in the benthic cover as well as on just the coral class as the primary class of interest.

its performance than the DNN. When evaluating the limited represented coral class in Table 2b, the DCDM outperforms the DNN across all metrics before and after the update. These further support the ability of the DCDM to sample locations that best improve its performance compared to the DNN (Table 2).

Part of the improvement of the DCDM over the DNN can be observed from the initial information prior of the two models. By having regions of high information close together, optimized trajectories spend more time gathering information as opposed to traversing to new high information regions. The DCDM prior had focused high information regions compared to the DNN prior which were more spread out. This contributes to the greater increase in performance of the DCDM as well.

Additionally, we observed the impact of varying the weights of each class variance when calculating the weighted sum information prior for the planner. Let $[\omega_C, \omega_A, \omega_S]$ be the weights for the coral, algae, and sand variances respectively. We found that increasing the weight towards a class can improve the prediction of that class after sampling. For example, having coral biased weights of $[\omega_C, \omega_A, \omega_S] = [0.6, 0.2, 0.2]$ when constructing the information prior yielded higher accuracy in coral performance post planning compared to equal weights of $[\omega_C, \omega_A, \omega_S] = [0.33, 0.33, 0.33]$.

Having regions of high information sparsely distributed across the map, even if the prior is fully biased towards a single class, can cause the resulting trajectory to prioritize exploration. We observed this behavior with purely coral weights of $[\omega_C, \omega_A, \omega_S] = [1, 0, 0]$. The coral variance was highest around the perimeter of the reef region in Figure 1c. As a result, the resulting ergodic trajectory then optimized a path across the map that improved performance across all classes as opposed to the expected individual class performance.

Overall, the general increase in performance post planning suggests that *in situ* expeditions following ergodic trajectories can definitely improve benthic cover mapping with efficient sampling. The DCDM especially benefited from the ergodic planning, generally achieving better performance than the DNN model across all metrics after the update. We also found that biasing weights when constructing the information prior can yield improved individual class performance as long as the regions of high information in the prior are not too spread out.

5. CONCLUSION

We have incorporated limited data and post-planning processes to build upon our previously established machine learning method of efficient global scale benthic cover interpretation [9]. This was accomplished by first further exploring how clustering can be used to approach limited and imbalanced data. We then evaluated how our model outputs can be used to determine efficient trajectories for *in situ* expeditions and showed that ergodic planning of sampling locations improves the predictions. Thus, we have expanded upon our prior approach by completing an efficient pipeline for coral reef classification with spectral unmixing and efficient sampling.

6. ACKNOWLEDGEMENTS

This research was supported by funding from the National Aeronautics and Space Administration (NASA) through awards NNX16AB05G and 80NSSC21K1159. DigitalGlobe/Maxar data were provided by NASA's Commercial Archive Data for NASA investigators (cad4nasa.gsfc.nasa.gov) under the National Geospatial-Intelligence Agency's NextView license agreement. The research was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

7. REFERENCES

[1] Antti Lipponen, "Heronisland 2017 01 25," (Accessed November 26, 2023).

[2] Fredrik Moberg and Carl Folke, "Ecological goods and services of coral reef ecosystems," *Ecological Economics*, vol. 29, no. 2, pp. 215–233, 1999.

[3] Marjorie Reaka, "The global biodiversity of coral reefs: a comparison with rainforests," 01 1997.

[4] ISRS, "Isrs consensus statement on climate change and coral bleaching, conf. of the parties to the united nations framework convention on climate change," December 2015.

[5] Clive R Wilkinson, *Status of Coral Reefs of the World*, Global Coral Reef Monitoring Network and Reef and Rainforest Research Centre, Townsville, Australia, 2008.

[6] Eric J. Hochberg and Michelle M. Gierach, "Missing the reef for the corals: Unexpected trends between coral reef condition and the environment at the ecosystem scale," *Frontiers in Marine Science*, vol. 8, 8 2021.

[7] Alberto Candela, Kevin Edelson, Michelle M. Gierach, David R. Thompson, Gail Woodward, and David Wettergreen, "Using remote sensing and in situ measurements for efficient mapping and optimal sampling of coral reefs," *Frontiers in Marine Science*, vol. 8, 2021.

[8] Alberto Candela Garza, *Bayesian Models for Science-Driven Robotic Exploration*, Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, September 2021.

[9] Rohan Zeng, Eric J. Hochberg, Alberto Candela, and David S. Wettergreen, "Spectral unmixing and mapping of coral reef benthic cover with deep learning," in *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2022, pp. 1–5.

[10] Eric J. Hochberg, Marlin J. Atkinson, and Serge Andréfouët, "Spectral reflectance of coral reef bottom-types worldwide and implications for coral reef remote sensing," *Remote Sensing of Environment*, vol. 85, no. 2, pp. 159–173, 2003.

[11] Eric Hochberg and M. Atkinson, "Spectral discrimination of coral reef benthic communities," *Coral Reefs*, vol. 19, pp. 164–171, 07 2000.

[12] Corinna Cortes and Vladimir Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, sep 1995.

[13] Jürgen Schmidhuber, "Deep learning in neural networks: An overview," *Neural Networks*, vol. 61, pp. 85–117, Jan. 2015.

[14] David Arthur and Sergei Vassilvitskii, "K-means++: The advantages of careful seeding," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, USA, 2007, SODA '07, p. 1027–1035, Society for Industrial and Applied Mathematics.

[15] George Mathew and Igor Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4, pp. 432–442, 2011.

[16] Kevin Edelson, "Ergodic trajectory optimization for information gathering," M.S. thesis, Carnegie Mellon University, Pittsburgh, PA, October 2020.

[17] Ananya Rao, Ian Abraham, Guillaume Sartoretti, and Howie Choset, "Sparse sensing in ergodic optimization," 2022.