JOURNAL OF
EMERGING INVESTIGATORS

Article

Machine learning on crowd-sourced data to highlight coral disease

Rithika Narayan[1] and Anthony J. Pellicano[2]
[1]Elwood John Glenn High School, Elwood, NY
[2]15 Bayberry Ave, Garden City, NY

SUMMARY
Triggered largely by the warming and pollution of oceans, corals are experiencing bleaching and a variety of diseases caused by the spread of bacteria, fungi, and viruses. Identification of bleached/diseased corals enables implementation of measures to halt or retard disease. Benthic cover analysis, a standard metric used in large databases to assess live coral cover, as a standalone measure of reef health is insufficient for identification of coral bleaching/disease. Proposed herein is a solution that couples machine learning with crowd-sourced data – images from government archives, citizen science projects, and personal images collected by tourists – to build a model capable of identifying healthy, bleached, and/or diseased coral. We collected hundreds of images of corals from open source archives, including the National Oceanic and Atmospheric Administration's records and the XL Catlin Seaview Survey. The image annotation platform Labelbox was used to highlight regions of interest in each of these images and label them as "healthy", "bleached", "black band disease", "dark spot disease", "white syndrome", or "yellow band disease". These annotations were then used to build, train, and validate a Python-based image classification model, adapted from an open-source Mask R-CNN (region-based convolutional neural network) algorithm. Use of the model on a test set of coral images yields over 85% accuracy in distinguishing healthy versus unhealthy coral. This machine learning-based model has the potential to rapidly analyze a large and growing database of images to identify coral bleaching/disease around the world, thereby enabling effective allocation of resources for preservation of our marine ecosystems.

INTRODUCTION

Coral reefs are among the most biodiverse marine ecosystems, housing millions of organisms across thousands of species within just 1% of the ocean floor (1). In fact, the largest living structure is the Great Barrier Reef, a coral reef off the coast of Australia (2). In addition to providing living habitats, coral reefs protect coastal areas from tidal waves and erosion. Healthy reefs absorb 97% of a wave's energy, buffering shorelines from currents, waves, and storms, thereby mitigating loss of life and property damage (1). In addition to providing tremendous ecological benefits, coral reefs support millions of people that rely on them for food and draw in tourists with over 70 million trips made annually, making these fragile and beautiful organisms a powerful engine of coastal and marine tourism (3). In one estimate, the net annual benefit of the world's coral reefs is approximately $30 billion in the forms of tourism, recreation, coastal protection, fisheries, and biodiversity (4).

Despite their incredible ecological and economic benefits, corals and coral reefs are in crisis. Environmental changes including warming oceans could result in changes to corals such as bleaching. This phenomenon is associated with the expulsion of dwelling algae by coral polyps causing them to lose more than just their bright coloration (5). Coral polyps typically live in an endosymbiotic relationship with these algae, with the latter fulfilling the majority of the coral's energy requirements, crucial to both coral and accompanying reef health. While bleached corals might survive in the short-term, starvation eventually sets in (5). The leading cause of coral bleaching is increasing water temperatures, with mass bleaching events occurring across hundreds of miles or more (6–9). In 2016, heat stress encompassed more than half the coral reefs globally and other mass bleaching events have eliminated swathes of healthy coral in the Great Barrier Reef, the western Indian Ocean, and the Seychelles (6-8).

In addition to warming waters, the impact from land-based sources of pollution – including coastal development, deforestation, agricultural runoff, oil spills, and chemical spills – can impede coral growth and reproduction, disrupt overall ecological function, and cause disease and mortality in sensitive species (6-9). It is now well accepted that many serious coral reef ecosystem stressors originate from land-based sources, most notably toxicants, sediments, and nutrients. Under the Endangered Species Act, 22 coral species are currently listed as threatened and 3 are listed as endangered (10). With more than 40% of the world's live coral cover lost in the last three decades, approximately 500 million people have been adversely impacted (11).

In addition to bleaching, corals face a plethora of other issues including disease and predation. The most common coral diseases globally include black-, yellow-, and white-band disease, dark spot disease, white plague, and white pox (1). Many of these diseases are direct results of human actions: white pox is linked to the pollution of the oceans with human fecal matter, and white plague, which thrives on bleaching vulnerable coral, also has viral and bacterial origins that can be traced to humans (12, 13). Other diseases are primarily attributed to fungi and bacteria, including dark spot disease, which is characterized by patches of brown or purple tissue on the surface of coral and is caused by endolithic fungi that calcify the coral (14). Ocean warming will greatly exacerbate the spread of these fungi, placing further strain on coral reefs. Banded diseases, which are caused by bacteria that flourish in warmer waters and algal overgrowths, will also occur more frequently due to pollution and ocean warming (15).

While there is tremendous cause for concern, there is much more need for action. Historical evidence indicates that reefs are able to recover, even from mass bleaching or predation events (16). However, their ability to recover is dependent on factors such as herbivory, algal cover, and coral recruitment (16). There are already initiatives in place, such as the of Mars Symbioscience, which places complex structures in open spaces in depleted reefs to serve as recruitment locations for new corals, that have successfully begun rebuilding our marine ecosystems (17). These initiatives rely on the support of data collection measures in order to direct resources to where they are most needed.

In order to mitigate and track coral disease, global efforts including the XL Caitlin Seaview Survey, the Allen Coral Atlas, and CORALNET are using underwater and satellite-based images of the coral benthic cover (18-20). For instance, the top view of the coral canopy can be used to create a coral map and assess coral health in terms of the percentage of live coral. These analyses are useful to track the health of a reef over time; for example, seeing a decrease in live coral cover while noticing an increase in algal cover could indicate that corals are losing the battle for space, sunlight, and oxygen to algae (21). Traditionally, these analyses have been done by human annotators, but this is both an expensive and laborious process. As described in Automated annotation of coral reef survey images, of the millions of images of reefs taken each year, only 1% is annotated by humans, which leaves a wealth of information untapped (22).

To bridge this gap, machine learning (ML) approaches have recently been unveiled that analyze benthic cover – including types of corals in a reef, other reef invertebrates, and sand – and achieve the same levels of accuracy as traditional methods but more efficiently (18-20). Provided empirical data, ML makes new predictions or recognizes new patterns. The field is very broad and includes many types of algorithms with various advantages and disadvantages for a given task, such as the support vector machine (SVM) used as the architecture for CORALNET or convolutional neural networks (CNNs). CNNs are a type of deep learning architecture most often used in image classification which apply convolutions to small groups of pixels to extract features and generate functions that map the input features to a classification output (22, 23). ML has been making headway in medicine, where it has already been successfully utilized in order to diagnose skin cancers, diabetic retinopathy, lymph node metastasis, pneumonia, emphysema, and many more pathologies (24). Indeed, most diagnoses can trigger intervention and this is the goal for use of ML in environmental, and more specifically

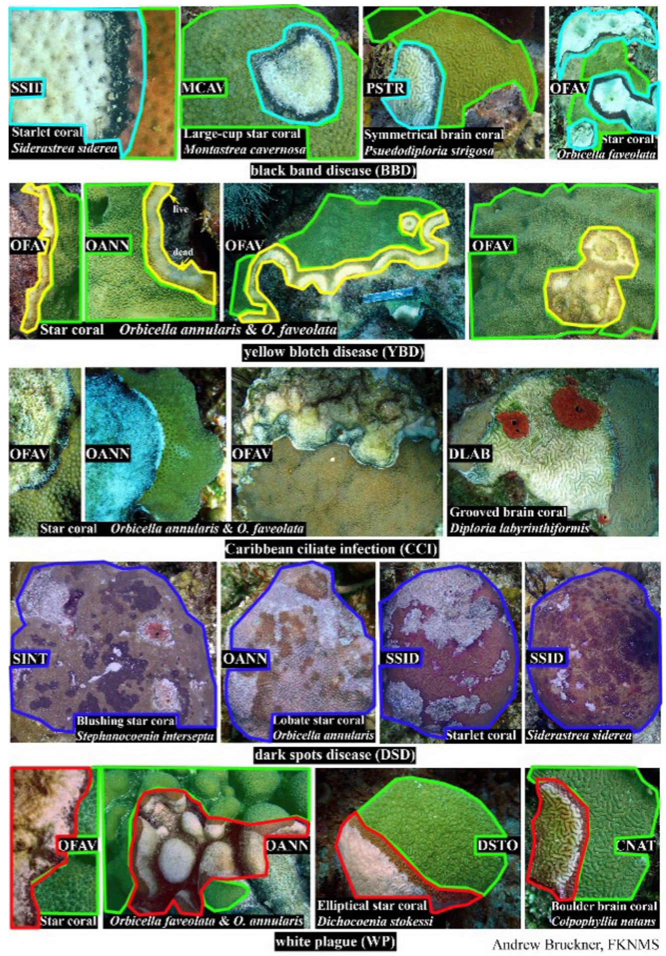

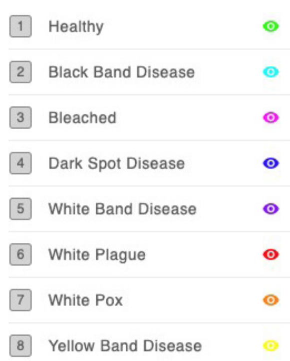

Figure 1. Annotation within Training Set. Example of image data annotated in Labelbox and used for model-training. Healthy and diseased corals were annotated with different colors using the polygonal selection tool. The image shows annotations for black band disease, yellow band disease, dark spot disease, and white plague as well as the surrounding healthy tissue. The annotations are stored in the JSON file for the model to learn from.

coral reef, monitoring. However, assessment of the benthic cover alone or coral variety within a reef alone, is insufficient to fully diagnose the extent of reef disease. Satellite imagery might show healthy coral tentacle tops whereas bleaching might be from bottom-up. Furthermore, satellite imagery may lack depth penetration or resolution to detect smaller lesions characteristic of yellow band disease or white pox.

One alternative approach to benthic cover analysis is using the hundreds of thousands of photographs routinely obtained by professional and amateur scuba divers, tourists, and other deep-sea forays. These images represent a warehouse of untapped information that can speak to coral

	Healthy	Bleached	Black Band Disease	Yellow Band Disease	Dark Spot Disease	White Syndromes
Training	456	184	54	142	62	250
Validation	114	14	14	14	16	36

Table 1. Distribution of individual instances of coral health conditions across training and validation data sets.

health across the globe. This approach, which encourages citizen science and empowers stakeholders, couples ML with crowd-sourced images to complement existing efforts to identify the extent of worldwide coral bleaching and coral diseases. Analyzing images uploaded by users across the globe for evidence of bleaching or disease by traditional methods would be a Herculean task. By contrast, ML could be deployed to detect and identify disease in corals by first training the machines on an annotated set of images.

Thus, our overarching hypothesis is that coupling ML with crowd-sourced images can be used to analyze existing and new databases to identify coral disease with greater accuracy and efficiency than human analyses. As described below, a CNN has now been trained on images of corals with varying degrees of health in order to create a model that can identify disease. This model, as an ML-based platform, can now be used to analyze incoming crowd-sourced images of corals along with the location of these corals. With a high accuracy rate of 85% in detecting unhealthy corals, this model can demonstrate the utility of ML in performing in-depth analysis of underwater images, beyond the scope of benthic cover analysis. In order to further develop this model, we need heightened public awareness regarding the health of the coral reefs, because this can translate to increased data and intensified conservation efforts. The data collected could be integrated into the National Oceanographic and Atmospheric Administration's (NOAA's) Deep Sea Coral Data Portal (DSCDP) or into citizen science projects such as the Great Reef Census, which asks volunteers to upload 10 photos of their diving site (25, 26). This coral health database can be rapidly used to identify trends in the health of different reefs and anticipate outbreaks. Based on these trends, resources for treatment of the diseases can be allocated, protection zones can be set up, and local awareness can be raised to lower risk factors.

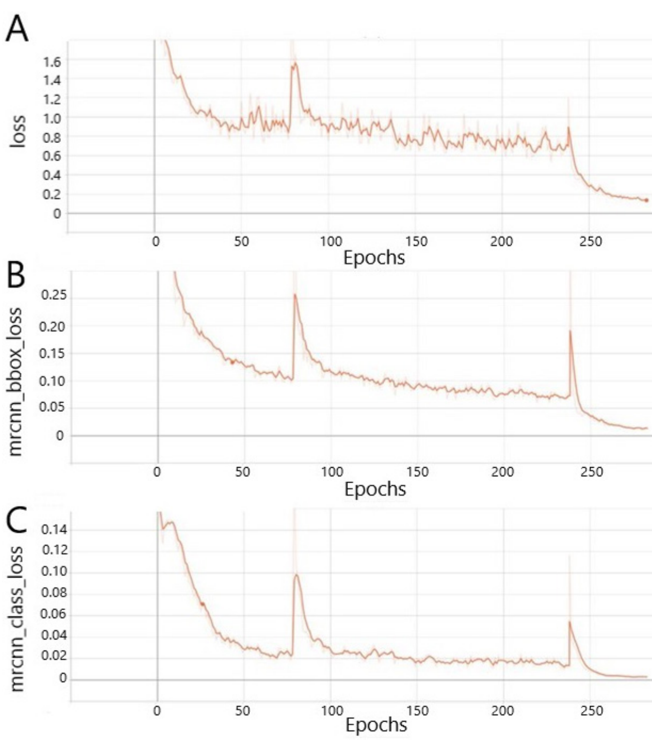

Figure 2. Tensorboard Loss Statistics following Model Training. Loss is defined as the difference between the results of the ground-truth "user" annotation training data and the model's prediction attempt. Pictured here: segmentation loss (A), binding-box loss (B), and class loss (C). Segmentation and binding-box identifies coral features while class loss predicts which class those features belong to healthy or diseased coral.

RESULTS

In order to develop a model capable of differentiating between healthy and unhealthy corals we selected an algorithm, the Mask R-CNN architecture, and modified it to be compatible with the training and validation datasets. These datasets were comprised of 335 images of healthy and unhealthy corals and were annotated to feature 1356 total individual instances of healthy, bleached, black-, yellow-, and white band disease, dark spot disease, white plague, and white pox and corals (Table 1, Figure 1). Descriptions from the image sources were used to judge the condition of each coral. The algorithm was then trained on these images and annotations and its performance was evaluated over the course of the training period until it achieved over 85% accuracy in distinguishing between healthy and bleached corals.

The model was trained on this dataset for 283 epochs (or passes) with heavy image augmentation to prevent overfitting. Loss values eventually approached zero, as shown by the Tensorboard plots (Figure 2). All training losses were minimized – segmentation loss approached 0.136, binding-box loss approached 0.013, and class loss approached .003. Peaks in the loss indicate a change in learning rate or the addition of image augmentation by the researcher. Both image augmentation and learning rate are configurations of the model that can be manually changed by the researcher (turned on or off in the case of augmentation and increased or decreased in the case of learning rate) over the course of training in order to drive the progression of the model's training. These two configurations are part of a broader set of

Backbone	Resnet101
Number of Classes	7 (includes background)
Batch Size	2
Steps per Epoch	1000
Validation Steps	50
Train ROIs per Image	200
Image Resize Mode	Square
Maximum Image Dimension	1024
Minimum Image Dimension	800
Mask Pool Size	14
Pool Size	7
Learning Rate	.001
Weight Decay	.0001
Learning Momentum	.9

Table 2. Selected hyperparameters from training. Hyper-parameters are values used by a researcher to control the training process and can be tuned during training to achieve the best results.

hyperparameters that were subject to tuning over the course of training in order to improve the model's performance (Table 2).

As expected, model accuracy was initially poor (Figure 3). For example, confidence levels started low (0.373, for example) and the confidence threshold was increased to 0.8 in order to display only predictions that the model was confident of. The ML model improved as training progressed, until it produced more accurate predictions of coral detection and class separation (Figures 4 and 5). In order to test the model's performance, the ML-based algorithm was run on a test set comprised of approximately 30 images of healthy and bleached corals, which were compiled by an independent source and not included in the training or validation sets. The model performed with >85% diagnostic accuracy in differentiating the two classes overall (Figure 6).

DISCUSSION

The research presented herein couples crowd-sourced data with ML to provide a deliverable that complements existing strategies to diagnose the extent of coral bleaching and identify diseased corals. It has the potential to spur citizen science and empower stakeholders as active participants and trigger corrective measures to halt or slow coral bleaching and disease.

Almost on a seasonal basis, we are experiencing irreversible ecological tipping points. The cascade of changes sparked by global warming could threaten the future of humanity, say scientists, who warned that more than half of the climate tipping points identified a decade ago are now "active" (16). Human footprints, including pollution and global warming, are starting to leave permanent marks on our ecosystem, including the calcium carbonate skeletons of bleached and starved corals. Eventually, the coral skeletons will erode, causing the reef structure to collapse. The mass bleaching events in recent years have been triggered by global warming and, if current trends continue, corals are expected to become increasingly rare on reef systems.

Fortunately, bleaching is reversible if the stressors are removed quickly (27). There is a global ongoing effort to harness satellite imagery to map benthic cover (18-20). However, additional resources must be taken advantage of. Fortunately, there is a large crowd-sourced databank in the form of videos and photographs of corals and coral reefs across the globe. These close-range photographs hold valuable information on the health of the corals. The research described herein utilizes a ML-based approach to diagnose coral bleaching based on these photographs. ML has precedent in diagnosis: the technology has revolutionized medicine, now routinely outperforming and outpacing human experts in making diagnoses at a fraction of the cost and saving lives across the world (24).

Our ML algorithm is currently able to identify healthy and bleached corals at > 85% accuracy. More data of the various disease classes (banded disease, dark spot disease, and the white syndromes) must be accumulated and used

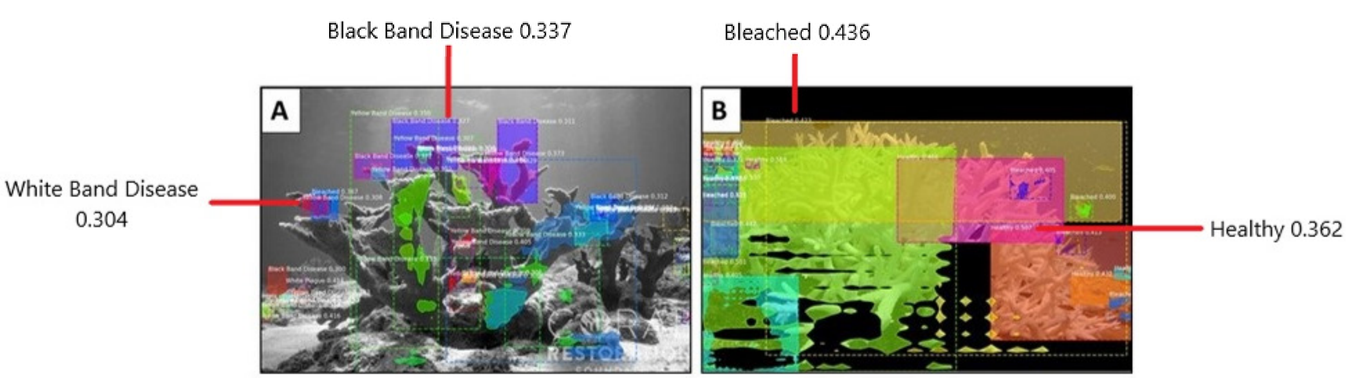

Figure 3. Early Model Predictions. Some coral features are picked out (A), but overall features (B) are poorly segmented and poorly identified). Each box represents a class prediction by the model with the predicted class and confidence level noted in the upper left-hand corner. Selected predictions are highlighted.

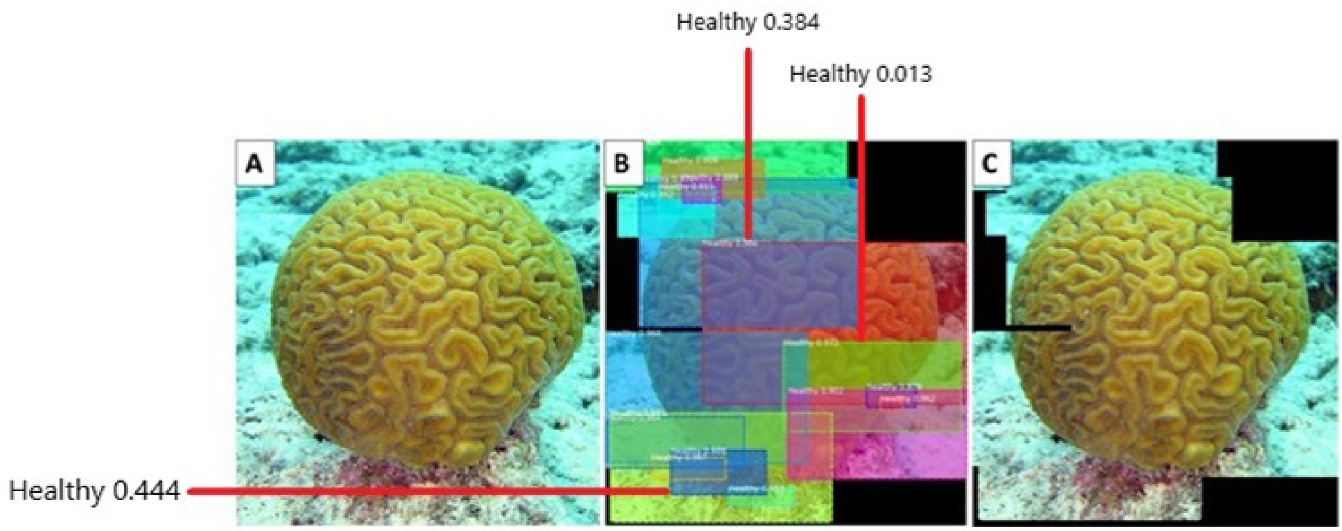

Figure 4. Healthy Coral Detection. Model detection of healthy coral. Images of healthy coral (A) are analyzed with the model which makes predictions (B) breaking the image down into segmented regions that the computer recognizes as healthy coral (visualized in C). Selected predictions are highlighted.

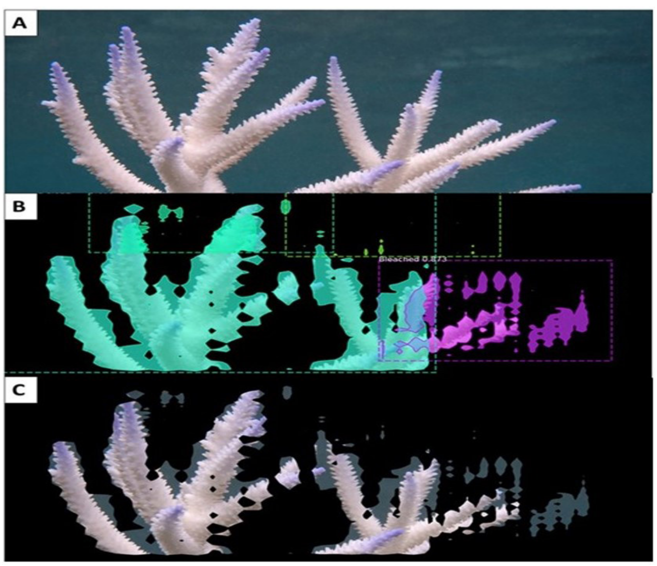

Figure 5. Bleached Coral Detection. Model detection of a bleached coral showcasing the developing segmentation capabilities of the model. Images of unhealthy coral (A) are analyzed with the model that makes predictions (B) breaking the image into segmented regions that the computer recognizes as bleached coral (visualized in C). Black areas are those that the model recognizes as background

to train the model in order to improve its performance in detecting and segmenting such diseases. Uploading of data, including images, into cloud storage makes them available for use both remotely and continuously. ML approaches come with notable benefits including: efficiency allowing for larger sample analyses, the presence of a "trained expert" wherever there is a computer rather than relying on thinly spread marine biologists, the detection of patterns and correlations which would require a mass workforce working in unison for large amounts of time to replicate, and learning based on expert evaluations.

Coral	Bleached Probability
	85.8%
	10.8%
	7%
	92.9%

Figure 6. Heathy vs. Bleached Coral Diagnosis. In the above images selected from the test set, the model outputs the probability of the coral it has identified in the image being bleached. The probabilities for the two healthy corals (the second and third images) are 10.8% and 7%, which would be ignored by the model in implementation through the use of a higher threshold. The two images of bleached corals (the first and fourth images), however, are identified as having a much higher probability of being bleached (85.8% and 92.9% respectively) which indicates the model's success in identifying coral bleaching.

This algorithm will be made available to NOAA and private companies such as National Geographic so that it can be used freely. It can also be incorporated into an application for use by the general public in order to encourage more tourists to obtain photographs of corals, run them through the algorithm, and upload the results to coral health databases like NOAA's DSCDP (25). With the successful implementation of widespread reef monitoring and analysis will come the ability to efficiently respond to bleaching and disease events. As the Reef Resilience Network notes, once a banded disease or white syndrome is diagnosed, its spread can be controlled through aspirations and removal of infected tissue, as well as isolation of bacteria through epoxy application (28). Although there is no method of treating bleached corals, these vulnerable corals can be protected from further harm and given the opportunity to recover by halting trawling, industrial runoff, and mining in the area after bleaching events (16).

This approach does have weaknesses in that some false discovery rate is inherent with ML methods. However, this feature can be minimized through more training on additional data (29). Secondly, images taken at greater depths might not have the same resolution and quality as images obtained at shallower depths resulting in false positive findings. This issue could be tackled through the use of algorithms such as the sea-thru algorithm that reconstructs color in underwater images in the preprocessing of input images (30). Nevertheless, this approach is an important and stakeholder-empowering first step to address a global crisis because diagnosis can trigger intervention (28).

MATERIALS AND METHODS

An ML algorithm was selected after a literature review, which indicated that a CNN, which has precedent for use in object segmentation, has a high Mathew's Correlation Coefficient (MCC), a performance parameter that assesses classifiers, compared to other algorithms when applied to coral images (23, 31). The Mask Region (R) CNN, an open-source CNN, was chosen because of its successful use in the analysis of nuclei in microscope images, detection of sports fields in satellite images, and other object detection and segmentation projects (32). An R-CNN was preferable to a CNN because it not only can classify objects in an image but detect where they are located in the image and output a "mask" over the region of interest (32). The Mask R-CNN was then modified to best suit the qualities of the dataset by adding and changing the names of the classes (the health conditions of the coral) as well as making the algorithm compatible with the data annotation platform of choice, Labelbox (33).

A total of 335 images of individual corals and coral reefs with the aforementioned health conditions (healthy, bleached, black-, yellow-, and white- band disease, dark spot disease, white plague, and white pox) were obtained, which corresponded to a total of 1356 individual instances of these health conditions (Table 1). These images were then annotated in Labelbox using the segmentation polygon tool to outline masks of the shape of the corals as well as disease lesions. The accuracy of the annotation of these images was ensured by collecting them from reliable research sources such as the XL Catlin Seaview Survey and published ecological papers, among others (18, 19). These masks from the annotations are what the algorithm learned from and emulated.

The dataset was split into two groups, one for training and one for validation. The setup for both groups was identical in terms of classes and annotation type, but the data were split between the two. This split sent 80% of the images in the dataset to the training set and 20% to the validation set, which is an ideal split for midsize datasets (34). The classes that images could be classified into were healthy, bleached, black band disease, dark spot disease, white syndromes, and yellow band disease. White syndromes encompass white band disease, white plague, and white pox, which can visually take on the same appearance but affect different species of coral and are collectively linked to the Vibrio bacteria (35). The platform generated a JavaScript Object Notation (JSON) format that contains image paths and the annotations for each image. JSONs are readable by humans and computers alike, making them a viable method for providing data to a model (36). This JSON was used by the Mask R-CNN algorithm to train and validate the images.

An Amazon Web Services (AWS) EC2 instance was started in which the algorithm was developed using the images from the dataset, the JSON files containing the image annotations, and the modified algorithm code. The instance was equipped with Tensorflow graphical processing unit (GPU) capabilities, which increases computing power and decreases training time (37). Over time, the algorithm generated logs as it trained that allow it to build on what was previously learned with each training session rather than relearning each session. This ability is the basis of the neural network.

The progress of the algorithm was monitored using Tensorboard, a feature of Tensorflow. Tensorboard provides loss graphs, which ideally approach but never reach zero, and accuracy graphs, which ideally approach but never reach one, that track model performance in multiple aspects of the task, including class identification and object segmentation. The loss and accuracy graphs informed decisions on how to adjust the training in order to improve the model (Table 2). If loss or accuracy seemed to stagnate, the learning rate of the algorithm was changed to help it escape from the loss landscape. Image augmentation, or transformations of annotated images used to artificially inflate the size of the dataset, was also implemented to prevent the model from overfitting. This was done because this dataset of 335 images and 1356 instances is a smaller dataset for an ML project. Examples of image augmentation include flipping or rotating the image such that it becomes unfamiliar to the model and acts as a new image for the model to learn from.

Once the model was in use, the confidence level it displayed was used as a measure of how well the program was performing. A consistently low confidence threshold indicated that more training was needed on a larger set, perhaps for a specific health condition if that is where the model was underperforming. Once the loss on the model plateaued at a sufficiently low point, the model was tested on an independently curated set of images of healthy and unhealthy corals.

REFERENCES

1. “Corals.” NOAA’s National Ocean Service, 1 June 2013, https://oceanservice.noaa.gov/education/tutorial_corals/welcome.html.
2. “What Is the Great Barrier Reef?” NOAA’s National Ocean Service, 1 June 2013, https://oceanservice.noaa.gov/facts/gbrlargeststructure.html.
3. Spalding, Mark, et al. “Mapping the Global Value and Distribution of Coral Reef Tourism.” Marine Policy, vol. 82, Aug. 2017, pp. 104–13, doi: 10.1016/j.marpol.2017.05.014.
4. “Value of Corals.” Coral Reef Systems, https://sioweb.ucsd.edu/projects/coralreefsystems/about-coral-reefs/value-of-corals/.
5. Pratchett, Morgan S., et al. “Changes in Bleaching Susceptibility among Corals Subject to Ocean Warming and Recurrent Bleaching in Moorea, French Polynesia.” PLOS ONE, vol. 8, no. 7, July 2013, doi:10.1371/journal.pone.0070443.
6. Hughes, Terry P., et al. “Global Warming and Recurrent Mass Bleaching of Corals.” Nature, vol. 543, no. 7645, Mar. 2017, pp. 373–77, doi:10.1038/nature21707.
7. Hughes, Terry P., et al. “Spatial and Temporal Patterns of Mass Bleaching of Corals in the Anthropocene.” Science, vol. 359, no. 6371, Jan. 2018, pp. 80–83, doi:10.1126/science.aan8048.
8. Le Nohaïc, Morane, et al. “Marine Heatwave Causes Unprecedented Regional Mass Bleaching of Thermally Resistant Corals in Northwestern Australia.” Scientific Reports, vol. 7, no. 1, Nov. 2017, doi:10.1038/s41598-017-14794-y.
9. Aeby, Greta S., et al. “Patterns of Coral Disease across the Hawaiian Archipelago: Relating Disease to Environment.” PLOS ONE, vol. 6, no. 5, May 2011, doi:10.1371/journal.pone.0020370.
10. Gregg, Rachel M. “Listing of Coral Reef Species under the U.S. Endangered Species Act.” Climate Adaptation Knowledge Exchange, 30 June 2010, https://www.cakex.org/case-studies/listing-coral-reef-species-under-us-endangered-species-act.
11. Wismer, Sharon, et al. “Young Fishes Persist despite Coral Loss on the Great Barrier Reef.” Communications Biology, vol. 2, no. 1, Dec. 2019, pp. 1–7, doi:10.1038/s42003-019-0703-0.
12. Luna, Gian Marco, et al. “Vibrio Harveyi as a Causative Agent of the White Syndrome in Tropical Stony Corals.” Environmental Microbiology Reports, vol. 2, no. 1, Feb. 2010, pp. 120–27, doi:https://doi.org/10.1111/j.1758-2229.2009.00114.x.
13. Patterson, Kathryn L., et al. “The Etiology of White Pox, a Lethal Disease of the Caribbean Elkhorn Coral, Acropora Palmata.” Proceedings of the National Academy of Sciences, vol. 99, no. 13, National Academy of Sciences, June 2002, pp. 8725–30, doi:10.1073/pnas.092260099.
14. Randall, Carly J., et al. “Does Dark-Spot Syndrome Experimentally Transmit among Caribbean Corals?” PLOS ONE, vol. 11, no. 1, Jan. 2016, doi:10.1371/journal.pone.0147493.
15. “Coral Diseases.” Florida Museum, 03 Oct. 2018, https://www.floridamuseum.ufl.edu/southflorida/habitats/corals/diseases/.
16. Hoegh-Guldberg, Ove, et al. “Coral Reef Ecosystems under Climate Change and Ocean Acidification.” Frontiers in Marine Science, vol. 4, May 2017, doi:10.3389/fmars.2017.00158.
17. “Taking Responsibility to Rebuild Coral Reefs.” Mars, https://www.mars.com/news-and-stories/articles/coral-reef-rehabilitation.
18. “XL Catlin Seaview Survey.” XL Catlin Seaview Survey, 2011, https://www.catlinseaviewsurvey.com.
19. “Allen Coral Atlas.” Allen Coral Atlas, https://allencoralatlas.org/.
20. “CoralNet BETA.” CoralNet, https://coralnet.ucsd.edu/
21. Clements, Samantha M. “The Faces and Functions of Algae on the Reef.” Scripps Institution of Oceanography, 27 Feb. 2015, https://sioweb.ucsd.edu/labs/coralreefecology/the-faces-and-functions-of-algae-on-the-reef/.
22. Beijbom, Oscar, et al. “Automated Annotation of Coral Reef Survey Images - IEEE Conference Publication.” 2012 IEEE Conference on Computer Vision and Pattern Recognition, 2012, doi:10.1109/CVPR.2012.6247798.
23. Raghav, Prabhu. “Understanding of Convolutional Neural Network (CNN) — Deep Learning.” Medium, 4 Mar. 2018, https://medium.com/@RaghavPrabhu/understanding-of-convolutional-neural-network-cnn-deep-learning-99760835f148.
24. Rajpurkar, Pranav, et al. “Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of the CheXNeXt Algorithm to Practicing Radiologists.” PLOS Medicine, vol. 15, no. 11, Nov. 2018, doi:10.1371/journal.pmed.1002686.
25. “Welcome to NOAA’s Deep-Sea Coral Data Portal.” NOAA Deep-Sea Coral Data, https://deepseacoraldata.noaa.gov/.
26. “The Great Reef Census.” The Great Reef Census, https://greatreefcensus.org.
27. Normile, Dennis. “There’s Only One Way to Save the Great Barrier Reef, Scientists Conclude.” Science, 15 Mar. 2017, https://www.sciencemag.org/news/2017/03/there-s-only-one-way-save-great-barrier-reef-scientists-conclude.

28. "Managing Coral Disease." Reef Resilience Network, https://reefresilience.org/managing-for-disturbance/managing-coral-disease/.
29. Li, Junning, and Z. Jane Wang. "Controlling the False Discovery Rate of the Association/Causality Structure Learned with the PC Algorithm." The Journal of Machine Learning Research, vol. 10, June 2009, pp. 475–514.
30. Akkaynak, Derya. "Sea-Thru." Derya Akkaynak, https://www.deryaakkaynak.com/sea-thru.
31. "Matthews Correlation Coefficient." Wikipedia, 4 July 2017, https://en.wikipedia.org/w/index.php?title=Matthews_correlation_coefficient&oldid=1005550933.
32. Abdulla, Waleed. Matterport/Mask_RCNN. Matterport, Inc, 2017. GitHub, https://github.com/matterport/Mask_RCNN.
33. "Labelbox." Labelbox, https://labelbox.com/.
34. "Pareto Principle." Wikipedia, 10 June 2005. Wikipedia, https://en.wikipedia.org/w/index.php?title=Pareto_principle&oldid=998795873.
35. Bythell, John, et al. "White Plague, White Band, and Other 'White' Diseases." Coral Health and Disease, edited by Eugene Rosenberg and Yossi Loya, Springer, 2004, pp. 351–65. Springer Link, doi:10.1007/978-3-662-06414-6_20.
36. "Introducing JSON." JSON, https://www.json.org/json-en.html.
37. "Use a GPU." TensorFlow, https://www.tensorflow.org/guide/gpu.

Article submitted: June 9, 2020
Article accepted: Septermber 3, 2020
Article published: July 26, 2021