

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

A deep learning model for measuring coral reef halos globally from multispectral satellite imagery

Simone Franceschini^{a,*}, Amelia C. Meier^a, Aviv Suan^a, Kaci Stokes^a, Samapriya Roy^{a,b}, Elizabeth M.P. Madin^a

^a *Hawai'i Institute of Marine Biology, University of Hawai'i at Mānoa, Kāne'ohe, HI 96744, USA*

^b *Planet Labs, San Francisco, CA 94107, USA*

ARTICLE INFO

Edited by Menghua Wang

Keywords:

Remote sensing
Machine learning
Object detection
Instance segmentation
Coral reef monitoring

ABSTRACT

Reef halos are rings of bare sand that surround coral reef patches. Halo formation is likely to be the indirectly result of interactions between relatively healthy predator and herbivore populations. To reduce the risk of predation, herbivores preferentially graze close to the safety of the reef, potentially affecting the presence and size of the halo. Reef halos are readily visible in remotely sensed imagery, and monitoring their presence and changes in size may therefore offer clues as to how predator and herbivore populations are faring. However, manually identifying and measuring halos is slow and limits the spatial and temporal scope of studies. There are currently no existing tools to automatically identify single reef halos and measure their size to speed up their identification and improve our ability to quantify their variability over space and time. Here we present a set of convolutional neural networks aimed at identifying and measuring reef halos from very high-resolution satellite imagery (i.e., ~0.6 m spatial resolution). We show that deep learning algorithms can successfully detect and measure reef halos with a high degree of accuracy ($F1 = 0.824$), thereby enabling faster, more accurate spatio-temporal monitoring of halo size. This tool will aid in the global study of reef halos, and potentially coral reef ecosystem monitoring, by facilitating our discovery of the ecological dynamics underlying reef halo presence and variability.

1. Introduction

Ecological systems are often characterized by conspicuous patterns resulting from complex interactions between biotic and abiotic variables (Tarnita et al., 2017). In recent decades, the development of remote sensing technologies and aerial surveys has allowed scientists to highlight and describe an increasing number of these patterns in several different ecosystem types, such as species vegetation patterns in semi-arid ecosystems (Barbier et al., 2006) and desert fairy circles (Jurgens, 2013). One striking example of such regular spatial patterns are reef halos (also known as grazing halos). Reef halos consist of ring-like patterns of bare sand that occur around coral patch reefs (i.e., reefs spatially separated from other reefs by expanses of sand or other substrates) in a variety of primary producer habitats (e.g., benthic macroalgae; seagrass, etc.) and their presence is readily visible from satellite images (Madin et al., 2011).

Reef halos (hereafter halos) were first reported in the 1960s in the

Caribbean Sea (Ogden et al., 1973; Randall, 1965), and they were later observed in several geographic areas, such as Australia's Great Barrier Reef, the Atlantic Ocean, and the Red Sea (Madin et al., 2011). Many hypotheses for halo patterns have been proposed, such as nutrient availability (Alevizon, 2002), bioturbation (Bilodeau et al., 2021), and foraging behavior through indirect fear effects (Madin et al., 2019b). However, predation risk is a widely attributed explanation for the creation and maintenance of halos. It is believed that predators control not only the abundance of herbivores, but also their foraging behavior, which is driven in part by a perception of risk. In the case of reef halos, predation risk constrains the foraging activity of grazing species near a central shelter, so that overall herbivory declines with increasing distance from patch reefs, in turn creating a sand ring pattern that is a spatially-explicit indicator of herbivory intensity (DiFiore et al., 2019; Madin et al., 2011). However, this pattern is expected to change when reef communities are overfished. For example, overfishing of top predators is thought to reduce the overall predation risk imposed on

* Corresponding author.

E-mail address: simonefr@hawaii.edu (S. Franceschini).

<https://doi.org/10.1016/j.rse.2023.113584>

Received 10 November 2022; Received in revised form 6 April 2023; Accepted 12 April 2023

Available online 18 April 2023

0034-4257/© 2023 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

herbivores and may, in turn, lead to a lack of halo formation (Madin et al., 2019a) or may expand the width of halos (but see Madin et al., 2019a). Alternatively, if fishing is concentrated on herbivores, halos may not be expected to form at all. If one or more of these expectations are true, halos may present a unique opportunity to quickly and inexpensively monitor fish community structure for a given geographical area.

Halos can therefore be considered highly visible seascape-scale footprints of herbivory, ecological interactions, and ecosystem processes which, because of their large spatial scale, can be observed and measured from satellite imagery. Madin et al. (2019a) measured halos using freely available satellite images and compared fished versus unfished (fully protected) coral reefs over thousands of km of coastline in Australia. Meanwhile, DiFiore et al. (2019) and Madin et al. (2019b) linked herbivore behavior with halo patterns at local scales by merging satellite imagery observations with in-situ and remote camera-based fish surveys, respectively. Combining field studies with remote sensing observations has proven highly useful in understanding halo occurrence and formation. Moreover, developing freely available remote sensing tools to monitor ecological processes over large scales is an important step in improving the management and conservation of coral reef ecosystems (Hedley et al., 2016), and reef halos may represent an emerging opportunity to monitor reef ecosystems' function at large scales, including in otherwise inaccessible areas. There are currently no tools to automatically detect and measure halos from satellite images. The process of manually annotating halos from satellite imagery is also slow and limits the scale over which such monitoring could be applied (Madin et al., 2019a, 2019b, 2011).

In recent years, deep learning and computer vision techniques have

been increasingly used in ecological studies (Christin et al., 2019). In particular, due to the increasing availability of high-resolution remotely sensed imagery, convolutional neural networks (CNNs; Albawi et al., 2017) have become a powerful tool for remote sensing studies (González-Rivero et al., 2020; Lei et al., 2020) and large-scale ecosystem analysis (Brodrick et al., 2019). For example, in the context of coral reef ecosystem monitoring and conservation, CNNs outperformed previous texture-based models (e.g., support vector machines, random forests) by providing more accurate results in remote sensing classification (Wan and Ma, 2020), object detection (Song et al., 2019), and semantic segmentation (King et al., 2018).

Here we present a deep learning method consisting of a mask region-based convolutional neural network (Mask R-CNN; He et al., 2018) and a U-Net model (Isensee et al., 2018) that are able to identify and measure reef halos from high-resolution (sub-meter pixel size) satellite imagery. We selected different geographical areas where halo occurrence has been reported, and, through the use of Planet SkySat images, we trained and tested a deep learning framework that can detect halo presence globally and measure their size. Our goal was to develop a tool to enable faster, automated, halo detection and measurement from satellite imagery to pave the way for developing a low-cost, globally-applicable coral reef conservation tool.

2. Materials and methods

2.1. Study areas and dataset

Our study area included 20 areas of interest (AOI) from 6 countries (Fig. 1). Geographical areas were selected based on the occurrence of

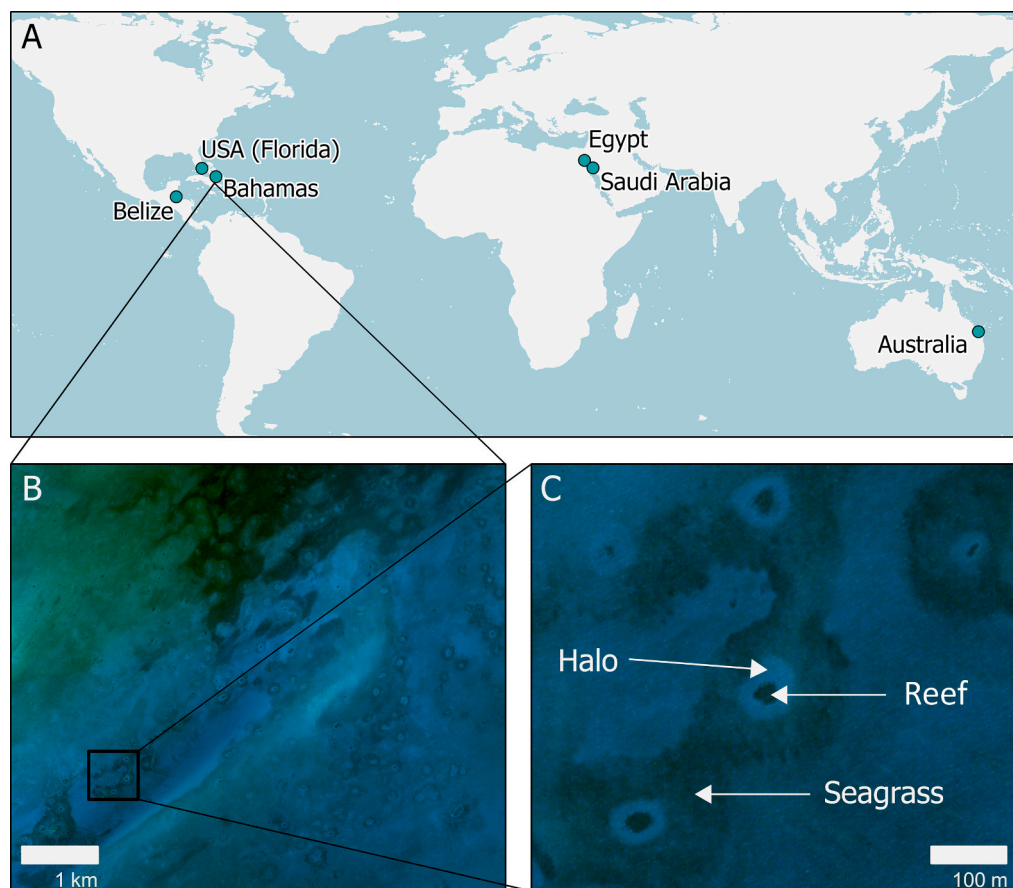


Fig. 1. A) Overview of the study regions. Number of unique AOIs per country: Australia = 1, Bahamas = 6, Belize = 7, Egypt = 3, Saudi Arabia = 2, USA (Florida) = 1. Detailed information for each AOI is reported in Table 1, section 2.3. B) Example of SkySat satellite image containing halos from the Bahamas. C) Zoomed example (a subset of B)) of halos.

halos known from previous studies (Downie et al., 2013; Madin et al., 2019a; Ogden et al., 1973; Randall, 1965) and by exploring areas with a potential occurrence of halos using Google Earth Pro. Four-band (Blue: 450–515 nm, Green: 515–595 nm, Red: 605–695 nm, NIR: 740–900 nm) SkySat satellite images were acquired through Planet Inc. (Planet Team, 2017) Planet Explorer Catalogue. Obtained as a SkySat Collect product, each image was roughly 20 km × 5.9 km, with a spatial resolution ranging between 0.5 m and 0.8 m. All surface reflectance products are orthorectified using fine digital elevation maps (30–90 m posting) and ground control points. Planet conducts atmospheric correction with the 6SV2.1 radiative transfer model which accounts for atmospheric absorption and scattering (Vermote et al., 2006), with aerosol optical depth, water, vapor and ozone inputs from MODIS near-real-time data (MOD09CMA and MOD09CMG). All calculations were done in R (R Core Team, 2022).

2.2. Deep learning model workflow overview

Our ultimate goal was to develop a tool that could identify reef halos from multispectral satellite images and automate their measurement. To achieve this goal, we developed and merged two different types of CNNs, a Mask R-CNN and U-Net model (described in detail below).

First, we labeled all non-overlapping halos present in the AOIs (Fig. 2, step 1), resulting in 4127 manually annotated halos. Merged objects were excluded from the training process since, where two reefs share some portion of a halo, it is extremely difficult to discriminate which sand area belongs to one or the other patch reef, thereby leading to an inaccurate estimate of its halo's size. Comparisons of halo sizes among locations was one of the goals of this study, thus excluding merged halos was necessary to generate accurate conclusions. Halos were labeled using ArcGIS software (version 2.9.1), allowing the geo-referenced information for all objects to be retained. To avoid biases

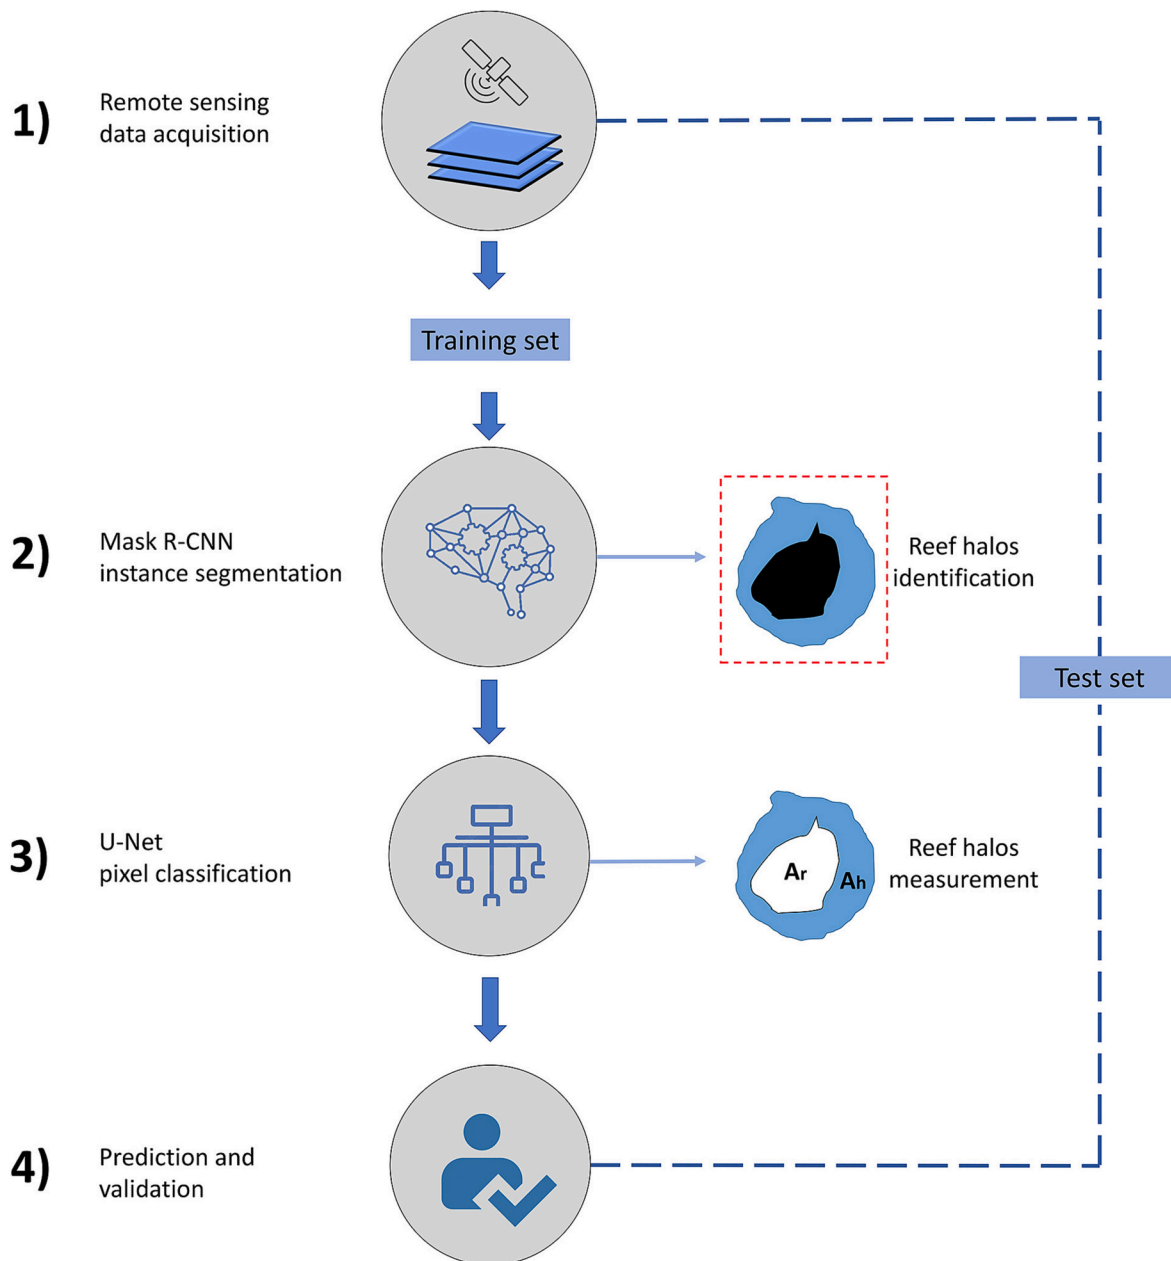


Fig. 2. Graphical summary of the deep learning model development. 1) Remote sensing data acquisition: halos have been manually labeled in ArcGIS Pro and then divided into training and test sets. 2) Mask R-CNN model development for reef halos identification. 3) U-Net model development for computing “reef” and “halo” areas (A_r and A_h respectively). 4) Prediction on test set samples and computation of performance metrics.

due to a single user's perception of halos size, five users were trained in the same labeling procedure – zooming into each halo at a 1:600 scale and tracing light contours with a mouse – and labeled the same areas of interest (AOIs). Results were compared, demonstrating an 84% area overlap between users' halo area annotations, and a Wilcoxon test showed non-significant differences in halo size among the five users (p -value = 0.73). We also assessed the depth of halo areas; Table S1 reports the mean values of halo depths across different geographic regions.

The dataset generated from the imagery annotation was divided into training and test sets (~70% and ~30%, respectively). The training set was used for model implementation and optimization, while the test set was used for comparing model-predicted vs. manually-annotated halos. In addition, we selected independent areas (i.e., AOIs where no halos were used for the training process) for the test set (Table 1; section 2.3) to better estimate the model generalization properties (Franceschini et al., 2018).

The second step of this work was to train the Mask R-CNN model (Fig. 2, step 2) using the training data set, thereby automatizing the identification of reef halos from the test set of satellite images and extracting the shape of the reef halos from the imagery background (i.e., extracting both the patch reef and its surrounding halo). The Mask R-CNN was trained with 3322 reef halos from 13 AOIs, while the remaining 805 halos from seven AOIs were kept to evaluate the model's performance (Table 1). After object extraction, we automated halo measurement using a U-Net pixel classification model (Fig. 2, step 3) which discriminated the halo (sand ring) from the interior reef patch. A total of 6428 annotations from the training set were used for model training. Annotations consisted of pixel areas manually classified by users, divided into "patch reef" and "halo" classes, respectively. Lastly, we compared the CNN's predictions and measurements to the observed values from the manually-generated test set (Fig. 2, step 4).

Satellite images were processed in one batch rather than individually by location. This approach can be more efficient in terms of processing time and computational resources, as the model can perform the necessary computations for multiple images simultaneously. Additionally, the model can learn from a larger and more diverse dataset, potentially improving its accuracy and ability to generalize to new data. Several metrics were taken into account in order to validate model accuracy on independent observations (described below, section 2.5).

Both the Mask R-CNN and U-Net models were developed using Tensorflow (ver. 2.5.1) (Abadi et al., 2016), Keras (ver. 2.4.3), and PyTorch (ver. 1.8.2) libraries in Python 3.9 and ArcGIS 2.9.1. The

machine running the deep learning model had an Intel® Core™ i9-9900K CPU @3.60 GHz processor with 128 GB of installed RAM and a PNY NVIDIA Quadro A6000 GPU.

2.3. Mask R-CNN model for halo segmentation

To identify the edges of the reef halos, we applied the computer vision technique of instance segmentation (Watanabe and Wolf, 2019). This method consists of a set of supervised learning algorithms that assign a label for each pixel by considering other nearby objects and the background, thereby identifying the object's boundaries. The Mask R-CNN segmentation model (He et al., 2018) is one of the most advanced instance segmentation techniques available to date and predicts a segmented region in parallel with bounding box recognition (Ren et al., 2015). The Mask R-CNN architecture consists of two different networks: the backbone and the region proposal network. These networks run once per image to give a set of region proposals containing the object. Specifically, our Mask R-CNN computes three outputs for each predicted object: (i) a class label, (ii) a bounding box that delimits the object, and (iii) a mask that distinguishes pixels belonging to the object from those to the background. For a more detailed description of the Mask R-CNN model structure see He et al. (2018). The rationale behind the use of an instance segmentation model like the Mask R-CNN is that the deep learning algorithm can not only recognize the color pattern of halos from multispectral information, but it can also determine the shape of the object based on the surrounding pattern and the background. In this case, the model only recognized a halo when it had the characteristic ring shape around a reef which was then surrounded by algae/seagrass beyond the outer edge of the halo.

In this study, we used the Residual Network 34 (ResNet34) backbone network for our Mask R-CNN (He et al., 2016; Wu et al., 2019). ResNet34 consists of a 34-layered network that has been pre-trained with millions of labeled, high-resolution images belonging to the ImageNet global dataset (Deng et al., 2009). In recent years, ResNet34 has outperformed previous models in a variety of tasks, providing accurate predictions for both object detection (Dai et al., 2016) and semantic image segmentation (Chen et al., 2018). In order to improve the ability of the model to learn and recognize reef halo objects, the weights of the backbone network were unfrozen during the training process. This step allowed the network to modify the weights during the backward pass of training, thereby improving the model's ability to predict a new class of objects (i.e., coral reef halos) (LeCun et al., 2015). The model's

Table 1

List of selected AOIs and number of observed halos for each area. Country codes: AUS = Australia, BHS = Bahamas, BLZ = Belize, EGY = Egypt, SAU = Saudi Arabia, USA = United States of America (Florida). Mean latitude and longitude are expressed in decimal degrees. Each AOI was assigned to the training or the test set, respectively.

AOI	Image ID	Mean LAT	Mean LON	Date (yyyy-mm-dd)	Area (km ²)	No. halos obs.	Subset
AUS01	20201107_025354_ssc9_u0001	23.46287	151.91099	2020-11-07	43.1	10	Test
BHS01	20190510_152255_ssc4_u0003	22.86968	-75.96340	2019-05-10	101.2	373	Training
BHS02	20190625_182659_ssc7_u0001	23.71767	-77.48061	2019-06-25	91.4	136	Training
BHS03	20190309_183114_ssc11_u0001	23.91426	-77.51960	2019-03-09	94.4	275	Training
BHS04	20190531_182931_ssc6_u0001	25.03904	-76.83888	2019-05-31	83.9	355	Training
BHS05	20200919_182915_ssc7_u0002	24.78926	-76.84203	2020-09-19	19.1	84	Test
BHS06	20190626_182402_ssc8_u0003	24.97612	-77.48089	2020-06-26	15.2	52	Test
BLZ01	20201129_191646_ssc8_u0001	16.84286	-87.82219	2020-11-29	63.8	211	Training
BLZ02	20201022_191415_ssc6_u0001	17.01935	-88.03578	2020-10-22	96.3	392	Training
BLZ03	20201022_191415_ssc6_u0002	16.77202	-87.84856	2020-10-22	97.9	371	Test
BLZ04	20200626_192104_ssc6_u0001	17.20029	-87.54292	2020-06-26	43.2	274	Training
BLZ05	20200626_192104_ssc6_u0002	17.44689	-87.52561	2020-06-26	36.4	164	Training
BLZ06	20200627_190037_ssc7_u0001	17.34428	-87.47563	2020-06-27	53.7	291	Training
BLZ07	20200627_191435_ssc8_u0001	17.50868	-87.67782	2020-06-27	93.8	801	Training
EGY01	20190413_081243_ssc13d3_0004	27.74041	34.20744	2019-04-13	8.2	29	Training
EGY02	20190413_081243_ssc13d3_0005	27.74803	34.19379	2019-04-13	10.2	22	Test
EGY03	20190413_081243_ssc13d3_0006	27.73911	34.19365	2019-04-13	7.4	4	Training
SAU01	20200202_110140_ssc7_u0001	28.03414	35.03130	2020-02-02	7.8	17	Training
SAU02	20210627_034355_ssc19_u0001	25.49153	36.75708	2021-06-27	25.1	13	Test
USA01	20210127_154738_ssc12_u0001	17.24515	-87.54992	2020-10-27	62.5	253	Test

learning rate was optimized using PyTorch function *torch.optim* and set to 0.0013, with batch size set to 64, the maximum number of iterations set to 10,000, and with early stopping when validation set error stopped decreasing to avoid overfitting (Franceschini et al., 2019).

A SoftMax function was used in the final layer of the network (Gao and Pavel, 2018). Thus, the final output of the neural network was converted into a normalized probability distribution (0–1), representing the probability that an object belonged to the class reef halo. In some cases, optimization of the threshold value is needed in order to discriminate if an object belongs to a class or not (Lei et al., 2020). In this study, the threshold value for reef halos detection was optimized according to the performance metrics computed on the test set (section 2.5).

2.4. U-net model for halo measurement

A U-Net convolutional network was performed on the Mask R-CNN extracted objects to automatically measure the identified halos. U-Net is a neural network designed primarily for image segmentation and pixel classification (Ronneberger et al., 2015). U-Net architecture consists of two paths. The first one, the contracting path, acts as a regular CNN (Albawi et al., 2017) with repeated application of convolutional layers and max pooling operations for downsampling (Kattenborn et al., 2021). The second path is an expansion path, which consists of up-convolutions and concatenations with features from the contracting path. This expansion allows the network to learn localized classification information and contextualize each pixel’s classification according to its surrounding area (Siddique et al., 2021).

We applied the U-Net model to classify and predict the area of two classes: patch reef and halos. The model’s learning rate was set to 0.0015, and batch size was set to 64. It was important to keep both output classes since, as both DiFiore et al. (2019) and Madin et al. (2019a) showed, patch reef geometry is a key determinant of halo size and thus should be taken into account when quantifying variation in halo size. The final output of the model was pixel assignments to either the patch reef or halo class.

It is important to point out that for the purpose of this final step of the model development, we considered “patch reef” as any object within the halo sand ring. In fact, the central core of the halo can sometimes be the result of a mixture of substances, like reef associated with seagrass or macroalgae (Madin et al., 2011; Ogden et al., 1973). This did not affect the definition of halo in sensu *latu* and allowed an optimization of the final prediction of the model, which was to differentiate the central patch from the halo ring pattern and to measure their areas. For convenience, in the following sections, we will refer to the core of the halo as “reef” and we acknowledge that it could be the result of different substances, such as the presence of primary producers, dead coral/rubble, or rocks.

2.5. Performance metrics

Mask R-CNN model performance was evaluated by computing precision, recall, and an F1-metric on the test data set. These metrics are commonly used for evaluating object detection and instance segmentation model performance (Goodwin et al., 2022; Guirado et al., 2021). The threshold value for predicting halo occurrence from test set images set was optimized according to these parameters:

$$Precision = \frac{True\ Positives}{True\ Positives + False\ Positive} \tag{1}$$

$$Recall = \frac{True\ Positives}{True\ Positives + False\ Negative} \tag{2}$$

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{3}$$

An inherent trade-off exists between the number of false positives and false negatives, and the F1 metric represents the most effective method for selecting the optimal threshold (Zhao et al., 2018). For both Mask R-CNN and U-Net model, we selected the dice coefficient (DICE; Dice, 1945) as our criteria to evaluate the pixel overlap between the prediction and the manually labeled region:

$$DICE = \frac{2 |A \cap B|}{|A| + |B|} \tag{4}$$

Where *A* is the predicted region and *B* is the ground truth. DICE is analogous to the F1-measure for estimating pixel overlap. The F1-metric and DICE values, therefore, represent the best compromise between the precision and recall measures of a classifier. Their scores can range from 0 to 1, with 1 representing a model that perfectly classifies each observation into the correct class and 0 representing a model that is unable to classify any observation into the correct class.

3. Results

3.1. Mask R-CNN model

Table 2 shows the precision, recall and F1-measure values of the Mask R-CNN model for each AOI and each threshold value.

As expected, model precision was inversely related to the threshold (Fig. 3). When the model is more conservative in assigning the “halo”

Table 2

Performance metrics of Mask R-CNN model according to each AOI and threshold value.

AOI	Precision	Recall	F1	Threshold	Mean
AUS01	1.000	0.400	0.571	0.8	Precision = 0.979 Recall = 0.550 F1 = 0.661
BHS05	0.868	0.392	0.540	0.8	
BHS06	1.000	0.076	0.142	0.8	
BLZ03	1.000	0.956	0.978	0.8	
EGY02	1.000	0.590	0.742	0.8	
SAU02	1.000	0.615	0.761	0.8	Precision = 0.957 Recall = 0.711 F1 = 0.794
USA01	0.985	0.818	0.894	0.8	
AUS01	0.857	0.600	0.705	0.7	
BHS05	0.920	0.690	0.789	0.7	
BHS06	1.000	0.269	0.424	0.7	
BLZ03	1.000	0.970	0.984	0.7	Precision = 0.871 Recall = 0.860 F1 = 0.857
EGY02	0.944	0.772	0.850	0.7	
SAU02	1.000	0.769	0.869	0.7	
USA01	0.978	0.905	0.940	0.7	
AUS01	0.642	0.900	0.750	0.6	
BHS05	0.919	0.797	0.853	0.6	Precision = 0.831 Recall = 0.931 F1 = 0.870
BHS06	0.944	0.653	0.772	0.6	
BLZ03	0.976	0.986	0.981	0.6	
EGY02	0.857	0.818	0.837	0.6	
SAU02	0.800	0.923	0.857	0.6	
USA01	0.959	0.940	0.950	0.6	Precision = 0.736 Recall = 0.964 F1 = 0.824
AUS01	0.555	1.000	0.714	0.5	
BHS05	0.917	0.952	0.935	0.5	
BHS06	0.937	0.865	0.900	0.5	
BLZ03	0.958	1.000	0.978	0.5	
EGY02	0.750	0.818	0.782	0.5	Precision = 0.650 Recall = 0.966 F1 = 0.761
SAU02	0.750	0.923	0.827	0.5	
USA01	0.945	0.964	0.955	0.5	
AUS01	0.454	1.000	0.625	0.4	
BHS05	0.833	0.952	0.888	0.4	
BHS06	0.776	1.000	0.873	0.4	Precision = 0.650 Recall = 0.966 F1 = 0.761
BLZ03	0.946	1.000	0.972	0.4	
EGY02	0.600	0.818	0.692	0.4	
SAU02	0.619	1.000	0.764	0.4	
USA01	0.925	0.980	0.952	0.4	
AUS01	0.333	1.000	0.500	0.3	Precision = 0.650 Recall = 0.966 F1 = 0.761
BHS05	0.757	0.964	0.848	0.3	
BHS06	0.634	1.000	0.776	0.3	
BLZ03	0.896	1.000	0.945	0.3	
EGY02	0.500	0.818	0.620	0.3	
SAU02	0.541	1.000	0.702	0.3	Precision = 0.650 Recall = 0.966 F1 = 0.761
USA01	0.889	0.984	0.934	0.3	




Fig. 3. Mean values of model precision (green), recall (pink) and F1-metric (blue) computed on test set AOIs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

label, the probability of mis-labeling an object as a halo (i.e., false positive) decreases. The opposite phenomenon occurs with recall, conversely determined by the proportion of false negatives. The optimal threshold was then selected according to the mean F1 between all AOIs (Fig. 3), and model performance showed a maximum F1 value with a threshold of 0.5. While locally-optimal results could be achieved for each geographical area by considering each site's metrics, but such models would lack global generalizability, hence all locations were pooled. By selecting a 0.5 threshold, our Mask R-CNN model correctly identified all the halos for most of the AOI (see recall, Table 2), with a range from 80% to 100% of detected objects.

The median pixel accuracy for the Mask R-CNN model was 79%, with a minimum value of 65% for our Bahamas sites and a maximum of 83% for our Florida sites. The median DICE value for the Mask R-CNN model was 0.731, with a minimum value of 0.698 for our Bahamas sites and a maximum of 0.781 for our Belize sites. Table 2 shows the effect of confidence thresholding on model performance. Recall values decreased for all models with increasing thresholds. In particular, for areas hosting a low apparent density of seabed vegetation, the model's ability to correctly identify the shape of the halo resulted in the lowest recall values for the higher threshold values (0.7–0.8). In fact, Fig. 4 shows how the absence of dense macrophytes in some geographic areas can alter the perception of halo edges, making their identification difficult even for a human observer. Furthermore, the low number of observed halos in AUS01 (Table 1) significantly affected the variance of the performance metrics. The model showed high accuracy (mean F1 = 0.8) in Red Sea areas (i.e., EGY02 and SAU02, Table 2), for which the number of observed halos in the training set was lower relative to the other AOIs (Table 1).

Most false positive predictions consisted of the identification of merged halos (see Fig. 6, section 3.2). Merged halos occur when two or more patch reefs share the same sand ring pattern because of their proximity (Fig. 4D). From a conceptual point of view, halos were correctly identified; however, we considered merged halo identification an error because we did not include merged halos in our training dataset and because the overlap between halos prevents explicit measurement of individual halo size, rendering these estimates inaccurate. No significant relationship was observed between the mean depth of the halos (S1) and the performance of the Mask R-CNN model across the different

regions.

3.2. U-net model and halo measurement

Reef halos extracted by the Mask-RCNN were used as input for prediction by the pixel classification model. Our U-Net model showed a median pixel accuracy values of 82% and 77% for reef and halo classes, respectively. The median DICE value was equal to 0.823, with a minimum of 0.789 for the AUS01 AOI and a maximum of 0.851 for the USA01 AOI. Fig. 5 and Fig. 6 show the final output of the deep learning model in different sample areas of the test set. Mean reflectance spectra of the central reef patch and the halos are provided in Fig. S2.

In particular, Fig. 6 provides further detail on the model's performance in identifying halos when they are isolated versus when they are merged. As described above, merged halos were not taken into account in this study given the difficulty of estimating halo sizes around individual reefs.


The integration of our method in ArcGIS Pro allowed us to retain the geographic information and the size of each object, thereby automating the area computation (Fig. 7). "Halo" and "reef" areas were computed only for the Mask R-CNN true positive halos.

The Australia AOI showed the smallest mean areas for both patch reef and halos, respectively (Fig. 7A) (patch reef: mean = 101.2 m², sd = 55.1; halos: mean = 356.5 m², sd = 188.5). The Bahamas and Florida AOIs tended to have larger average patch reefs and halos (patch reef: mean = 625.5 m², sd = 610.9 for Bahamas AOI; halos: mean = 1411.7 m², sd = 888.3 for Florida AOI). For all AOIs, the mean individual halo area was greater than the mean patch reef area, which matched what we observed during the labeling process. In particular, the Red Sea regions are characterized by large halos surrounding relatively small patch reefs.

4. Discussion


This paper presents a deep learning method to automatically identify and measure coral reef halos from multispectral satellite images. Given the potential that grazing halos have shown in monitoring the health status of reef ecosystems via large-scale 'footprints' of ecological interactions (Madin et al., 2019a), it is necessary to automate and optimize

Upper Left: 151°54'29"E 23°27'43"S




0 0.02 0.04 0.07 0.11 0.14 Kilometers

Upper Left: 77°27'49"W 24°58'15"N




0 0.05 0.1 0.2 0.3 0.4 Kilometers

Upper Left: 87°51'55"W 16°43'36"N



0 0.04 0.07 0.14 0.21 0.28 Kilometers

Upper Left: 80°11'14"W 25°22'11"N



0 0.05 0.1 0.2 0.3 0.4 Kilometers

Fig. 4. Examples of reef halos in different geographic areas. A) AUS01, B) BHS06, C) BLZ03 and D) USA01. The density of macrophytes on the seabed clearly affects the perception of halo edges, resulting different levels of accuracy by the Mask R-CNN among geographic areas (Table 2).

their identification in order to reveal spatial and temporal changes over reef systems globally. Manually annotating halos is extremely time- and labor-intensive. Our deep learning approach can identify and measure roughly 300 halos over an area of 100km² in ~2 min with very high accuracy (F1 = 0.978), whereas the same task requires approximately

10 h for a human annotator to accomplish.

The best model was selected by manually setting different threshold values and comparing F1 scores. In a global model perspective, the use of an ROC curve (Gonçalves et al., 2014) would have led to the similar result, but, in this instance, our goal was to compute the performance of

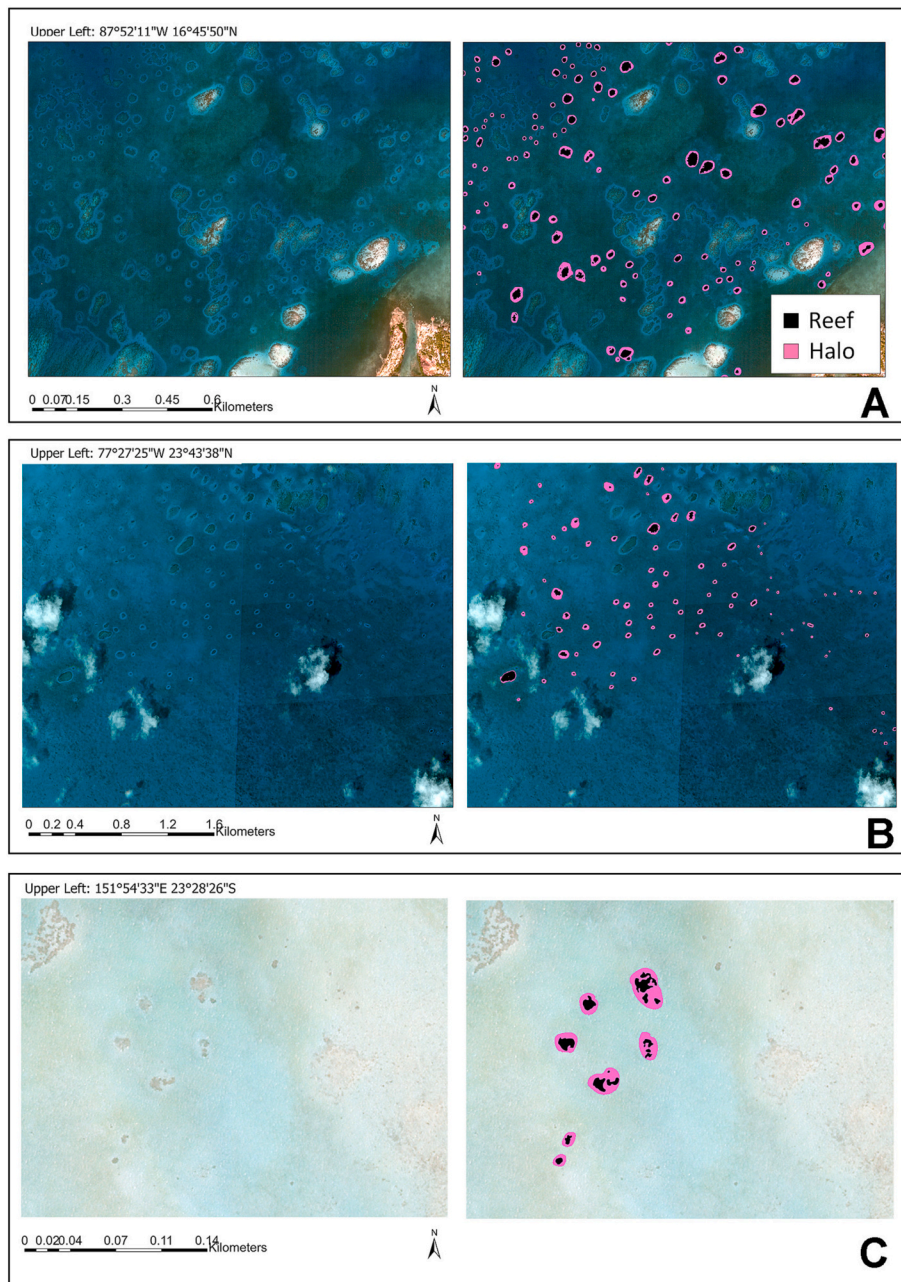


Fig. 5. An illustration of the model output in different AOIs. As a satellite image is processed (left side), the Mask R-CNN model identifies the halos and U-Net model classifies pixel according to the “reef” (black) and “halo” (pink) classes (right side). Each panel shows a zoomed-in (subset) area of different test set AOIs: A) BLZ02, B) BHS03 and C) AUS01. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the model in each geographic region based upon the value of the global threshold in order to visually demonstrate how model accuracy varies across regions. Accuracy values for halo prediction were higher in the Atlantic Ocean (e.g., Belize and Florida). These locations had higher numbers of halos than other areas. This result may have been due in part to the relatively high densities of seagrass/algae on the seabed in these areas, which increases the contrast of halo boundaries (i.e., the sand/macrophyte boundary; Fig. 4C,D) and therefore the ability of the model to correctly identify the halos. Nevertheless, even with lower values (of what?), our model demonstrated high levels of accuracy in detecting halos in geographic areas where seagrass/algae cover was more sparse (e.g., Australia; Figs. 4A, 5C). The seagrass/algae density is obviously a key factor in determining the continuity and contrast of halo contours, altering the perception of the halo boundary even for a human

observer. Despite this limiting factor, our model correctly predicted the presence of halos with a precision of 55% for Australia and 70% for the Red Sea. In the future, we plan to increase the training dataset for these and other areas to reach the same sample size as we had for the Atlantic Ocean. As expected, macrophytes density also affected pixel classification, with the lowest DICE value of 0.789 for Australia (Fig. 5C). Nevertheless, the model performance in estimating the area of patch reef and halo classes was very close to that of a human observer, as the mean overlap between observed and predicted area was 0.823 for the deep learning model, and it was 0.840 between the halos labeled by different users in the quality control phase.

Fig. 5 shows the final output of the deep learning model in three sample locations (i.e., Belize, Bahamas, and Australia). Although the benthic characteristics of the seabed appear different as the density and

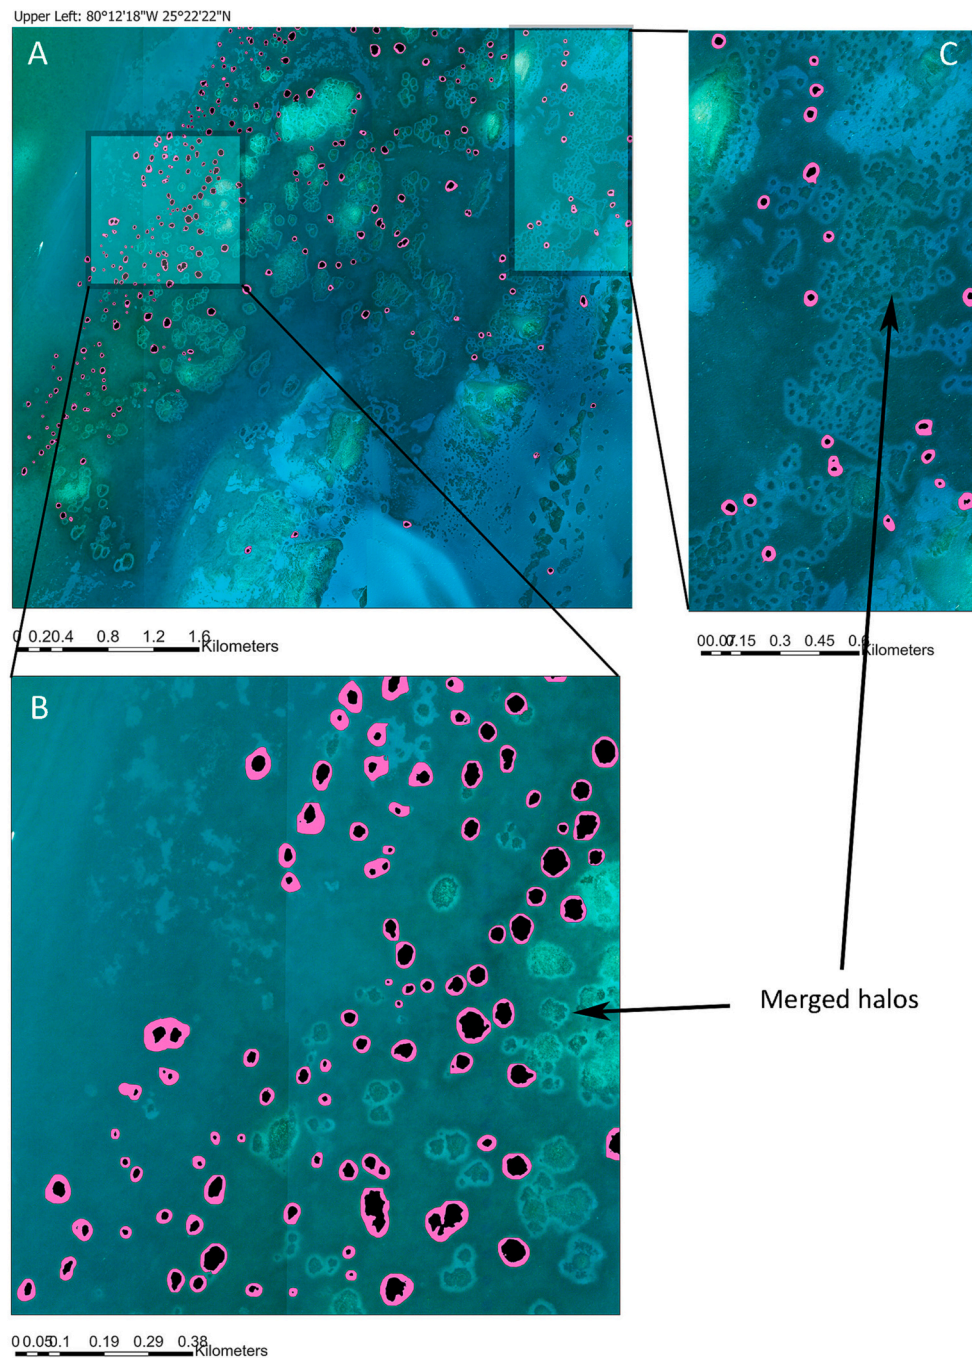


Fig. 6. A. Model prediction in the USA01 area. The model predicts single halos with high accuracy (B) ($F1 = 0.955$), while it tends to exclude most merged halos (C).

composition of the primary producers varies across the different areas, the model distinguishes the unique ring shape of the halo (“pink” color) from the background. In particular, Fig. 5 shows that the instance segmentation model can identify the characteristic ring shape around the patch reef regardless of the nature of the primary producer, which is known to vary across geographical areas. The performance of the model is most heavily affected by the density of the primary producers on the seabed, which limits the ability of the algorithm to identify halo boundaries. One of the qualities that renders the Mask R-CNN an efficient object detector of halos is that it bases its prediction not only on the color bands of the object with respect to the background, but also on its size and its shape. This enables the model to function in a variety of geographical areas where halos may occur in (even rather sparse) macroalgae, seagrass or other macrophytes. The primary objective in

developing this tool was to automate the measurement of halos and determine their area relative to their central patch reefs, regardless of the primary producer surrounding the shape of the halo, and the fact that our model can overcome this limitation renders it broadly useful for locations around the world with widely different dominant macrophytes.

Our model also tends to exclude very large patch reefs (Fig. 5A, Belize), as these tend to approach the size of smaller islands (radius > 0.5 km). This is not necessarily a shortcoming of the model since, although halos can assume a range of sizes (e.g., up to ~90 m width) (Downie et al., 2013), the literature suggests that most are in the range of tens of meters in width (Madin et al., 2022).

As mentioned above, we intentionally excluded merged (i.e., overlapping) halos from our model development, given the uncertainty when estimating the reef:halo area ratio for such halos. As a result, the

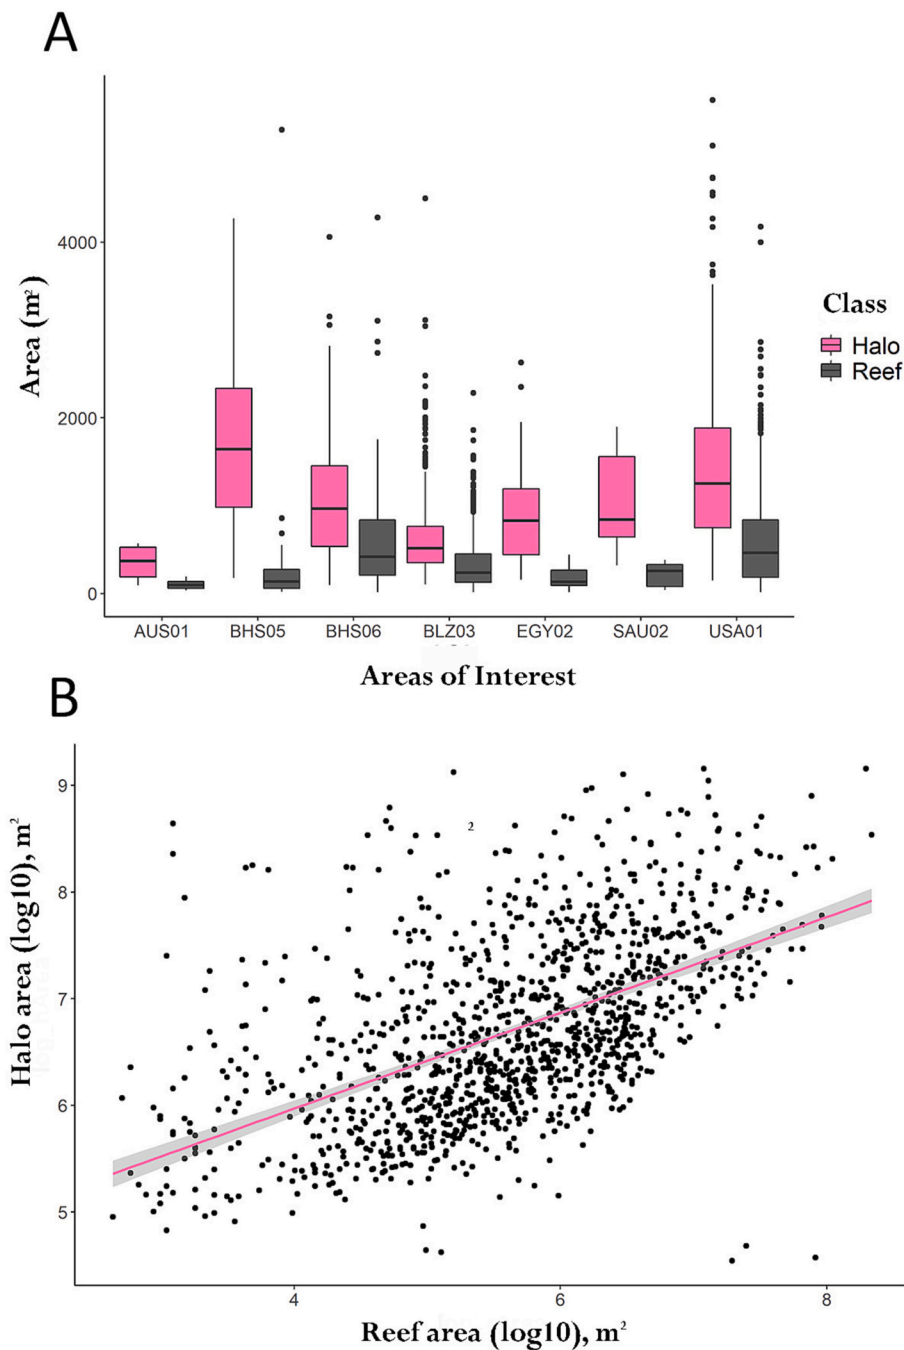


Fig. 7. A) Box plots of model-predicted “reef” and “halo” areas for each AOI. B) Scatter plot showing the (log-transformed) relationship between reef and halo areas over all AOIs.

Mask R-CNN model tends to exclude merged halos (Fig. 6B,C). Although this represents a limitation in terms of prediction of the number of halos in a given location, it allows the model to obtain more accurate estimates of the average size of halos across time and space. Inclusion of merged halos would have biased the model’s estimates of the ratio of reef to halo area (Fig. 7), as it was not possible to discern with precision what constitutes an individual reef’s halo. This would have undermined one of our main objectives of the work, namely to provide the basis for a tool to monitor the presence and average size of halos over space and time.

As patch reefs increase in size, larger halos are more likely to occur (Fig. 7B). This result is in agreement with Madin et al.’s (2019a) findings for the Great Barrier Reef and suggests that the same pattern tends to hold across geographic regions globally. As the patch reef area:

perimeter ratio increases, more herbivores are expected to forage along a functionally smaller perimeter, resulting in wider halos due to the increased impact of higher herbivore foraging density (DiFiore et al., 2019; Madin et al., 2019b). Fig. 2A shows that patch reefs and halo sizes vary greatly by geographical area. Although investigating the relationships between the size of halos and geographic variables was outside the scope of this work, we did explore the relationship between halo size and latitude, but we found no clear relationship. Given the wide geographical range and large climatic differences among our AOIs, it is plausible to assume that many factors may affect the size of the halos globally, such as the growth rates of the dominant species of seagrass/algae, the size structure of the fish and invertebrate community, and the physical and chemical characteristics of each region (e.g., hydrodynamics and

nutrient regimes). Each of these factors remain speculative, but they emphasize how our deep learning model can be used in the future to untangle the relationships between halo occurrence and size and ecological, environmental and/or anthropogenic variables.

Our model framework was trained and optimized using high-resolution images (~0.6 m), thus using satellite images with a lower resolution as input (i.e. > 2.0 m) could lead to an underestimate of objects because halo boundaries necessarily become less pronounced at lower pixel sizes. Future iterations of the model will include other spatial resolution imagery as inputs to extend and test the model's applicability to coarser (and generally less expensive) satellite imagery. Mask R-CNN models generally require a large amount of training data to learn pattern characteristics and perform segmentation, yet our deep learning model started to show a good predictive performance with a minimum number of ~1000 halos – a possible result of the unique nature of halos relative to their background. Our model is freely available, and its integration into ArcGIS as a user-friendly platform provides easy access to other researchers and organizations working in coral reef conservation. This tool serves as a point of departure to quickly identify and measure changes in halos across space and time at a global scale, supporting future studies to further understand the underlying mechanisms behind the creation of halos and to potentially increase their use them as indicators of key ecological processes, such as herbivory and predation, that are crucial factors for the health of coral reef ecosystems.

Author contributions

SF was responsible for Conceptualization, Data curation, Formal analysis, Methodology, Validation, Visualization, and Writing - the original draft. ACM was responsible for Data curation, Formal analysis, Validation, Visualization, and Writing – review & editing. AS was responsible for Data curation, Validation, and Writing – review & editing. KS was responsible for Validation. SR was responsible for Data curation. EMPM was responsible for Conceptualization, Funding acquisition, Data curation, Investigation, Project administration, Resources, Supervision, and Writing – review & editing.

Declaration of Competing Interest

Corresponding authors declare on behalf of all authors that there is no conflict of interest. We declare that we do not have any commercial or associative interest that represents a conflict of interest in connection with the work submitted.

Data availability

Data will be made available on request.

Acknowledgements

We thank the Planet and ArcGIS teams for their support and valuable discussions. We also thank NVIDIA Academic Hardware Grant Program for providing us with a high-performance GPU running the deep learning framework. This work was generously supported by a US National Science Foundation CAREER Fellowship (award number 1941737), NOAA/Sea Grant (award #NA18OAR4170076) and a Paul M. Angell Family Foundation grant to EMPM. Maps throughout this paper were created using ArcGIS® software by Esri. ArcGIS® and ArcMap™ are the intellectual property of Esri and are used herein under license. Copyright © Esri. All rights reserved.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.rse.2023.113584>.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mane, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viegas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X., 2016. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. <https://doi.org/10.48550/arXiv.1603.04467>.
- Albawi, S., Mohammed, T.A., Al-Zawi, S., 2017. Understanding of a convolutional neural network. In: 2017 International Conference on Engineering and Technology (ICET). Presented at the 2017 International Conference on Engineering and Technology (ICET), pp. 1–6. <https://doi.org/10.1109/ICEngTechnol.2017.8308186>.
- Alevizon, 2002. Enhanced seagrass growth and fish aggregations around Bahamian patch reefs: The case for a functional connection.
- Barbier, N., Couteron, P., Lejoly, J., Deblauwe, V., Lejeune, O., 2006. Self-organized vegetation patterning as a fingerprint of climate and human impact on semi-arid ecosystems. *J. Ecol.* 94, 537–547. <https://doi.org/10.1111/j.1365-2745.2006.01126.x>.
- Bilodeau, S.M., Layman, C.A., Silman, M.R., 2021. Benthic pattern formation in shallow tropical reefscapes: does grazing explain grazing halos? *Landsc. Ecol.* 36, 1605–1620. <https://doi.org/10.1007/s10980-021-01239-1>.
- Brodrick, P.G., Davies, A.B., Asner, G.P., 2019. Uncovering ecological patterns with convolutional neural networks. *Trends Ecol. Evol.* 34, 734–745. <https://doi.org/10.1016/j.tree.2019.03.006>.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L., 2018. DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 834–848. <https://doi.org/10.1109/TPAMI.2017.2699184>.
- Christin, S., Hervet, É., Lecomte, N., 2019. Applications for deep learning in ecology. *Methods Ecol. Evol.* 10, 1632–1644. <https://doi.org/10.1111/2041-210X.13256>.
- Dai, J., He, K., Sun, J., 2016. Instance-aware semantic segmentation via multi-task network cascades. In: Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3150–3158.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. Presented at the 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>.
- Dice, L.R., 1945. Measures of the amount of ecologic association between species. *Ecology* 26, 297–302. <https://doi.org/10.2307/1932409>.
- DiFiore, B.P., Queenborough, S.A., Madin, E.M.P., Paul, V.J., Decker, M.B., Stier, A.C., 2019. Grazing halos on coral reefs: predation risk, herbivore density, and habitat size influence grazing patterns that are visible from space. *Mar. Ecol. Prog. Ser.* 627, 71–81. <https://doi.org/10.3354/meps13074>.
- Downie, R.A., Babcock, R.C., Thomson, D.P., Vanderklift, M.A., 2013. Density of herbivorous fish and intensity of herbivory are influenced by proximity to coral reefs. *Mar. Ecol. Prog. Ser.* 482, 217–225. <https://doi.org/10.3354/meps10250>.
- Franceschini, S., Gandola, E., Martinoli, M., Tancioni, L., Scardi, M., 2018. Cascaded neural networks improving fish species prediction accuracy: the role of the biotic information. *Sci. Rep.* 8, 1–12.
- Franceschini, S., Mattei, F., D'Andrea, L., Di Nardi, A., Fiorentino, F., Garofalo, G., Scardi, M., Cataudella, S., Russo, T., 2019. Rummaging through the bin: modelling marine litter distribution using artificial neural networks. *Mar. Pollut. Bull.* 149, 110580.
- Gao, B., Pavel, L., 2018. On the Properties of the Softmax Function with Application in Game Theory and Reinforcement Learning. *ArXiv170400805 Cs Math*.
- Gonçalves, L., Subtil, A., Oliveira, M.R., de Bermudez, P., Z., 2014. ROC curve estimation: an overview. *REVSTAT-Stat. J.* 12, 1–20. <https://doi.org/10.57805/revstat.v12i1.141>.
- González-Rivero, M., Beijbom, O., Rodríguez-Ramírez, A., Bryant, D.E.P., Ganase, A., Gonzalez-Marrero, Y., Herrera-Reveles, A., Kennedy, E.V., Kim, C.J.S., Lopez-Marciano, S., Markey, K., Neal, B.P., Osborne, K., Reyes-Nivia, C., Sampayo, E.M., Stolberg, K., Taylor, A., Vercelloni, J., Wyatt, M., Hoegh-Guldberg, O., 2020. Monitoring of coral reefs using artificial intelligence: a feasible and cost-effective approach. *Remote Sens.* 12, 489. <https://doi.org/10.3390/rs12030489>.
- Goodwin, M., Halvorsen, K.T., Jiao, L., Knausgård, K.M., Martin, A.H., Moyano, M., Oomen, R.A., Rasmussen, J.H., Sørtdalen, T.K., Thorbjørnsen, S.H., 2022. Unlocking the potential of deep learning for marine ecology: overview, applications, and outlook. *ICES J. Mar. Sci.* 79, 319–336. <https://doi.org/10.1093/icesjms/fsab255>.
- Guirado, E., Blanco-Sacristán, J., Rodríguez-Caballero, E., Tabik, S., Alcaraz-Segura, D., Martínez-Valderrama, J., Cabello, J., 2021. Mask R-CNN and OBIA fusion improves the segmentation of scattered vegetation in very high-resolution optical sensors. *Sensors* 21, 320. <https://doi.org/10.3390/s21010320>.
- He, K., Gkioxari, G., Dollár, P., Girshick, R., 2018. Mask R-CNN.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778.
- Hedley, J.D., Roelfsema, C.M., Chollett, I., Harborne, A.R., Heron, S.F., Weeks, S., Skirving, W.J., Strong, A.E., Eakin, C.M., Christensen, T.R.L., Ticzon, V., Bejarano, S., Mumby, P.J., 2016. Remote sensing of coral reefs for monitoring and management: a review. *Remote Sens.* 8, 118. <https://doi.org/10.3390/rs8020118>.
- Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P.F., Kohl, S., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., Maier-Hein, K.H., 2018. nnU-Net: Self-adapting Framework for U-Net-Based Medical Image Segmentation. *ArXiv180910486 Cs*.

- Juergens, N., 2013. The biological underpinnings of Namib Desert fairy circles. *Science* 339, 1618–1621. <https://doi.org/10.1126/science.1222999>.
- Kattenborn, T., Leitloff, J., Schiefer, F., Hinz, S., 2021. Review on convolutional neural networks (CNN) in vegetation remote sensing. *ISPRS J. Photogramm. Remote Sens.* 173, 24–49. <https://doi.org/10.1016/j.isprsjprs.2020.12.010>.
- King, A., Bhandarkar, S.M., Hopkinson, B.M., 2018. A comparison of deep learning methods for semantic segmentation of coral reef survey images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1394–1402.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521, 436–444. <https://doi.org/10.1038/nature14539>.
- Lei, J., Luo, X., Fang, L., Wang, M., Gu, Y., 2020. Region-enhanced convolutional neural network for object detection in remote sensing images. *IEEE Trans. Geosci. Remote Sens.* 58, 5693–5702. <https://doi.org/10.1109/TGRS.2020.2968802>.
- Madin, E.M.P., Harborne, A.R., Harmer, A.M.T., Luiz, O.J., Atwood, T.B., Sullivan, B.J., Madin, J.S., 2019a. Marine reserves shape seascapes on scales visible from space. *Proc. R. Soc. B Biol. Sci.* 286, 20190053. <https://doi.org/10.1098/rspb.2019.0053>.
- Madin, E.M.P., Madin, J.S., Booth, D.J., 2011. Landscape of fear visible from space. *Sci. Rep.* 1, 14. <https://doi.org/10.1038/srep00014>.
- Madin, E.M.P., Precoda, K., Harborne, A.R., Atwood, T.B., Roelfsema, C.M., Luiz, O.J., 2019. Multi-trophic species interactions shape seascape-scale coral reef vegetation patterns. *Front. Ecol. Evol.* 7.
- Madin, E.M.P., Precoda, K., Roelfsema, C.M., Suan, A., 2022. Global conservation potential in coral reef halos: consistency over space, time, and ecosystems worldwide. *Am. Nat.* 200, 857–871. <https://doi.org/10.1086/721436>.
- Ogden, J.C., Brown, R.A., Salesky, N., 1973. Grazing by the echinoid *Diadema antillarum* Philippi: formation of halos around west indian patch reefs. *Science* 182, 715–717. <https://doi.org/10.1126/science.182.4113.715>.
- R Core Team, 2022. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria <https://www.R-project.org/>.
- Randall, J.E., 1965. Grazing effect on sea grasses by herbivorous reef fishes in the West Indies. *Ecology* 46, 255–260. <https://doi.org/10.2307/1936328>.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN. In: *Towards Real-Time Object Detection with Region Proposal Networks*, in: *Advances in Neural Information Processing Systems*. Curran Associates Inc.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Lecture Notes in Computer Science*. Springer International Publishing, Cham, pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.
- Siddique, N., Paheding, S., Elkin, C.P., Devabhaktuni, V., 2021. U-net and its variants for medical image segmentation: a review of theory and applications. *IEEE Access* 9, 82031–82057. <https://doi.org/10.1109/ACCESS.2021.3086020>.
- Song, J., Gao, S., Zhu, Y., Ma, C., 2019. A survey of remote sensing image classification based on CNNs. *Big Earth Data* 3, 232–254. <https://doi.org/10.1080/20964471.2019.1657720>.
- Tarnita, C.E., Bonachela, J.A., Sheffer, E., Guyton, J.A., Coverdale, T.C., Long, R.A., Pringle, R.M., 2017. A theoretical foundation for multi-scale regular vegetation patterns. *Nature* 541, 398–401. <https://doi.org/10.1038/nature20801>.
- Vermote, E.F.T.D., Tanré, D., Deuzé, J.L., Herman, M., Morcrette, J.J., Kotchenova, S.Y., 2006. Second simulation of a satellite signal in the solar spectrum-vector (6SV). 6S User Guide Vers. 3, 1–55. https://ltdri.org/files/6S/6S_Manual_Part_1.pdf.
- Wan, J., Ma, Y., 2020. Multi-scale spectral-spatial remote sensing classification of coral reef habitats using CNN-SVM. *J. Coast. Res.* 102, 11–20. <https://doi.org/10.2112/SI102-002.1>.
- Watanabe, T., Wolf, D.F., 2019. Instance segmentation as image segmentation annotation. In: 2019 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2019 IEEE Intelligent Vehicles Symposium (IV), pp. 432–437. <https://doi.org/10.1109/IVS.2019.8814026>.
- Wu, Z., Shen, C., van den Hengel, A., 2019. Wider or deeper: revisiting the ResNet model for visual recognition. *Pattern Recogn.* 90, 119–133. <https://doi.org/10.1016/j.patcog.2019.01.006>.