REPORT

A global coral reef probability map generated using convolutional neural networks

Jiwei Li[1] · David E. Knapp[1] · Nicholas S. Fabina[1] · Emma V. Kennedy[2] ·
Kirk Larsen[3] · Mitchell B. Lyons[2] · Nicholas J. Murray[2,4] · Stuart R. Phinn[2] ·
Chris M. Roelfsema[2] · Gregory P. Asner[1]

Abstract Coral reef research and management efforts can be improved when supported by reef maps providing local-scale details across global extents. However, such maps are difficult to generate due to the broad geographic range of coral reefs, the complexities of relating satellite imagery to geomorphic or ecological realities, and other challenges. However, reef extent maps are one of the most commonly used and most valuable data products from the perspective of reef scientists and managers. Here, we used convolutional neural networks to generate a globally consistent coral reef probability map—a probabilistic estimate of the geospatial extent of reef ecosystems—to facilitate scientific, conservation, and management efforts. We combined a global mosaic of high spatial resolution Planet Dove satellite imagery with regional Millennium Coral Reef Mapping Project reef extents to build training, validation, and application datasets. These datasets trained our reef extent prediction model, a neural network with a dense-unet architecture followed by a random forest classifier, which was used to produce a global coral reef probability map. Based on this probability map, we generated a global coral reef extent map from a 60% threshold of reef probability (reef: probability $\geq 60\%$, non-reef: probability $< 60\%$). Our findings provide a proof-of-concept method for global reef extent estimates using a consistent and readily updateable methodology that leverages modern deep learning approaches to support downstream users. These maps are openly-available through the Allen Coral Atlas.

Introduction

Accurate and reliable maps are a prerequisite for quantifying and analyzing geospatial patterns and the processes that underpin those patterns. With coral reefs experiencing unprecedented change (Hughes et al. 2018; Eakin et al. 2019), reef management and monitoring agencies, as well as the science community, require information on the location and extent of these environments from local (km) to global scales, to understand and manage these biodiverse and valuable ecosystems. Global reef maps are fundamental to the valuation of reef ecosystem services (Pendleton et al. 2016; Spalding et al. 2017), understanding the past (Heron et al. 2016), present (Burke et al. 2017), and future threats to reefs (Van Hooidonk et al. 2016), supporting more effective conservation (Beyer et al. 2018) and reef restoration strategies (Foo and Asner 2019), and facilitating scientific collaborations and research outcomes

Topic Editor Morgan S. Pratchett

✉ Gregory P. Asner
gregasner@asu.edu

1 Center for Global Discovery and Conservation Science, Arizona State University, Arizona 85281, USA

2 Remote Sensing Research Centre, School of Earth and Environmental Sciences, The University of Queensland, Queensland 4072, Australia

3 Vulcan Inc, Washington 98104, USA

4 College of Science and Engineering, James Cook University, Queensland 4811, Australia

(McManus 1994). Now, increasingly sophisticated Earth observation data and analytical tools provide an opportunity to improve the spatial and thematic resolution of reef maps, both over large spatial extents and at increasing temporal frequencies (Roelfsema et al. 2020).

A single coral reef map has been the basis of most modern scientific and management efforts. In 1994, the United Nations Environment Program (UNEP) World Conservation Monitoring Centre (UNEP-WCMC) began compiling and digitizing 150 years of reef maps into a single layer "World Atlas of Coral Reefs" (Spalding et al. 2001). Concurrently, Landsat 7 imagery made it possible to begin estimating reef extents from space (Millennium Coral Reef Mapping Project; hereafter, MCRMP), especially in remote or inaccessible locations (Andrefouet et al. 2006). The combination of these efforts and others has resulted in the UNEP-WCMC *Global Distribution of Coral Reefs* data product (Spalding et al. 2001; IMaRS-USF 2005; UNEP-WCMC et al. 2018), leading to many of the outcomes already cited and innumerable others. At the same time, its accuracy can be inconsistent because the underlying data are derived from multiple sources with different methodologies, and regions have different levels of sampling intensity.

There have been efforts to iterate on the UNEP-WCMC reef map with direct observations, remote sensing, and modeling approaches, often to produce maps with detailed geomorphic and ecological features. The NOAA National Ocean Service has developed their mapping effort, covering 43,000 km^2 from 0 to 150 m depth within US waters in the Pacific and Caribbean (Monaco et al. 2012). More recently, the Allen Coral Atlas is developing a global reef map using field and satellite techniques carried out on a region-by-region basis (Lyons et al. 2020; http://allencoralatlas.org), while the Living Oceans Foundation has mapped, by their estimates, 95,000 km^2 (5% of the global reef area) using seafloor observations, depth sounding, and WorldView satellite imagery (Purkis et al. 2019). These various approaches are promising and may lead to synergistic outputs, i.e., methodological differences between regions in a single map can be problematic, whereas map "ensembles" can be created and analyzed when maps are internally consistent but differ from one another in terms of goals, assumptions, and methodologies.

Convolutional neural networks (CNNs) are commonly used for computer vision and image processing tasks. CNNs have been used extensively in remote sensing (Cheng et al. 2017; Ma et al. 2019) and are rapidly being adopted for ecological applications of remote sensing data (Brodrick et al. 2019; Christin et al. 2019; Li et al. 2020b). Similar to other conventional approaches, CNNs can incorporate spectral or color information, but they have the advantage of incorporating spatial context, local or global texture features,

and can self-generate features rather than rely on hand-crafted features from experts. At the same time, CNNs also have greater programming and computational overhead, and they require large training datasets. Even so, CNNs are a natural choice for developing geospatial data resources from high-resolution satellite imagery.

We present a global coral reef probability map and a global coral reef extent map generated by convolutional neural networks and Planet Dove satellite imagery, with the ultimate goal of linking coral reef ecology and monitoring groups. Our maps are a useful resource because they: (1) use a single methodology and are therefore globally consistent and updateable, (2) leverage modern deep learning methods to build on previous approaches, while incorporating new ones, and (3) provide a unique comparison with other mapping resources. Here we describe the development of the satellite imagery and reef extent inputs, the model architecture, and training strategy. We then report on the results from the global coral reef probability map and global coral reef extent map at global and regional scales. Finally, we discuss opportunities for improving the approach and future maps.

Materials and methods

Planet Dove satellite imagery

The features used to train the model came from a global natural color composite mosaic generated from the Planet Dove satellite constellation. The mosaic was generated by the Planet team (www.planet.com) using 554,663 individual scenes collected between 1 Oct. 2017 and 1 Sept. 2018. It contains three spectral bands corresponding to red, green, and blue portions of the solar-reflected spectrum. The global visual mosaic is split into sub-regional images called "quads", with each image being 4096 × 4096 pixels at approximately 4.5 m resolution, or almost 20,000 × 20,000 m. We filtered the visual mosaic down to any quads in tropical and subtropical regions which may have any amount of reef area, for a total of 50,084 quads. The filtering was achieved by making a global mask of waters less than 20 m deep using the GEBCO bathymetry layer (GEBCO Group 2019), and large turbid or rocky embayments (non-reef) were additionally filtered out using a normalized difference index threshold from a 2017 global median MODIS mosaic ([band 1 − band 3]/[band 1 + band 3] > − 0.8). This mask was then buffered by 10 km, and we then confirmed that all existing reef polygons from the UNEP-WCMC mapping were covered.

The Planet mosaicking process modified band values and inter-band relationships to create a "Dove visual mosaic" that was meant to be a seamless representation of

the globe with limited cloud and cloud shadow. In addition to losing the inter-band relationships, this also meant that raw reflectance values and the near-infrared band are not available. As a result, this visual mosaic of red, green, and blue bands precluded an effective classification of land features, correction for sun glint or waves, or water depth estimation, among other challenges. While these significantly limit the use of the visual mosaic for some scientific applications, the curated removal of cloud and cloud shadow by Planet in creating the visual mosaic was considered invaluable in the often cloud-covered tropics.

Millennium Coral Reef Mapping Project (MCRMP) reef features

The reef data used to train the model were acquired from the open-access MCRMP maps, which cover a subset of the world's reefs (IMaRS-USF 2005). The MCRMP-classified reef features were based on a supervised classification method and manual post-processing. The features were classified into several hierarchical categories, from coarse to detailed, and covered a representative sample of global reef diversity. We chose to train the model with only this dataset to ensure that reef features were created with consistent definitions and methodologies. The MCRMP maps were manually reviewed to assess consistency with the Dove visual mosaic. The MCRMP maps for Chagos, Sri Lanka, Tobago, and Vietnam were removed due to low correspondence between the visual mosaic and the MCRMP reef features, that is, sometimes reef features would not be visually apparent due to water depth or clarity and therefore would not be effective training examples for the model. This helped to ensure that the model would find more consistent relationships between the visual mosaic and reef features. The remaining MCRMP training data features were spread across 2518 of the 50,084 visual mosaic quads (2425 after reducing the training dataset further, see below).

The MCRMP reef classes were reduced to a subset that would be tractable to model and sufficient for predicting global reef extent. The MCRMP classifications have five hierarchical levels of resolution, from coarse ("L1") to detailed ("L5"). We chose the L4 attribute level as the starting point because the L4 level appeared to be adequate for our needs, that is, the L3 attributes were not fine enough to separate spectrally-distinct reef features into distinct classes, while the L5 attributes appeared to sometimes partition spectrally- and ecologically-similar reef features into distinct classes. We mapped the 62 L4 attribute classes to 11 combined classes for our model training process (Table 1): one for land, two for water, five for reef features, and three for non-reef features (submerged, but visible features which did not represent reef area). The reef classes

were separated into fore reefs, shallow reef flats, variable depth reef flats, pinnacles, and lagoons.

Several assumptions guided this mapping process as a result of our manual exploration of the data, with three critical assumptions driving the majority of our work. First, we focused on mapping reef and non-reef classes separately. Second, depth plays a significant role in the appearance (or lack thereof) of reef features in the imagery, so we largely kept shallow and variable depth classes separate, and removed most deep features completely, using the MCRMP depth attribute to group these classes. The two deep classes retained in the model were the "shelf hardground" and "shelf-slope", as they were more often visible in the Dove visual mosaic and represent a significant portion of the benthos in some regions, including the Caribbean. Third, we redefined lagoon classes from non-reef to reef because they are fundamental geological, ecological, and socio-economic components of reef ecosystems (Aswani and Vaccaro 2008; Montaggioni and Braithwaite 2009).

Other assumptions were made to best facilitate the modeling process (detailed information can be found in Table 1). First, we removed features with classes containing "with constructions" from the training data, including "terrace with constructions" or "lagoon with constructions". Each of these classes represents multiple spectral or geomorphic features in a single class—e.g., both terrace and constructions, or both lagoon and constructions—and we instead train the model to predict these classes separately and at a finer spatial resolution. Second, we mapped the "aquatic land features" and "brackish atoll lagoon" classes to their own water class, as they were visually distinct from other water features in many cases and a separate class would likely improve model performance. Third, we removed features in the "reticulated fringing" class as their boundaries were not always precisely or accurately defined, at least relative to our visual mosaic. All modifications above did not have a large effect on the results because they impacted a small proportion of features, or focused on non-reef classes.

Supplemental training data

The MCRMP data alone are not sufficient for training a global reef model because they do not include classes for other non-reef geospatial features found in the Dove visual mosaic. Specifically, the MCRMP data have no classes for deep water or clouds, and our sampling process (described below) initially under-sampled land areas. We addressed these issues by manually creating supplementary training data with three additional classes—deep water, clouds, and supplementary land—which led to the final 14 classes used for the model (Table 1).

Table 1 Cross-connections between original Millennium Coral Reef Mapping Project reef feature classes, model training classes, and aggregated model classes

MCRMP classes	Model training classes	Aggregated model classes
Land on reef, main land	Land	Land
N/A	Supplemental land	Land
Channel, deep lagoon, double barrier lagoon, drowned lagoon, haa enclosed lagoon, pass	Water	Water
Aquatic land features, brackish atoll lagoon	Water other	Water
N/A	Supplemental deep water	Water
Faro forereef, forereef	Forereef	Reef
Enclosed basin, enclosed lagoon, enclosed lagoon or basin, faro enclosed lagoon, shallow lagoon	Lagoon	Reef
Bay exposed fringing, crest, faro reef flat, fractal reef flat, intermediate reef flat, linear reef flat, pass reef flat, reef flat, ridge and fossil crest, uplifted reef flat	Shallow reef flat	Reef
Immature reef flat, subtidal reef flat	Variable depth reef flat	Reef
Barrier reef pinnacle/patch, lagoon pinnacle, pinnacle	Pinnacle	Reef
Diffuse fringing	Shallow non-reef	Non-reef
Cross-shelf, inner slope, inner terrace, outer terrace, shallow lagoonal terrace, shallow terrace, shelf terrace	Variable depth non-reef	Non-reef
Shelf hardground, shelf slope	Deep non-reef	Non-reef
N/A	Supplemental clouds	Non-reef (clouds)
Bridge, deep drowned reef flat, deep terrace, drowned bank, drowned inner slope, deep lagoon with constructions, deep terrace with constructions, drowned pass, drowned patch, drowned rim, enclosed lagoon or basin with constructions, enclosed lagoon with constructions, forereef or terrace, haa subtidal reef flat, reticulated fringing, shallow lagoon with constructions, shallow lagoonal terrace with constructions, shallow terrace with constructions, shelf terrace with constructions, undetermined envelope	Removed	N/A

We used early versions of the model to identify imagery with high-value supplemental training data. For deep water and land classes, the model would misclassify areas with unique spectral patterns or textures, and we were able to iteratively train the model, generate additional training data to improve performance, and retrain the model. For instance, in addition to deep water with a rich, sapphire color, we also collected samples from areas with distinct waves or sunglint patterns. For land, we selected additional samples from areas with grassland, forests, jungles, and cities, both in flat areas and in areas with topography that could cause shadows. Similarly, the supplementary cloud training data were selected by running earlier versions of the model and finding areas where the model had misclassified clouds, most often as reef features. Because the model output can be converted to vector format, we were able to easily convert the "outlines" that the model had drawn around clouds (and misclassified as reef) to a new cloud class.

Convolutional neural network

We used a CNN to model the relationship between the red, green, and blue band pixel values in the Dove visual mosaic and the 14 reef training classes. We tested several CNN architectures [i.e., u-net, dense u-net, Fully Convolutional Network (FCN)] and selected a dense u-net architecture because it showed the highest performance (Figs. 1, 2; Zhang et al. 2018; Guan et al. 2019). The dense u-net architecture is a combination of a u-net architecture (Iandola et al. 2014; Ronneberger et al. 2015) and a dense net architecture. The u-net structure is defined by the encoder-decoder pattern, through which the image features are downsampled to a smaller resolution and then upsampled to the original resolution for the final predictions (Fig. 1). This pattern drives the model to reduce the number of features to a subset that is most useful for predicting the responses, while also using pass-through layers to localize those predictions at the original image resolution during the upsampling process. The dense net component is that each convolutional layer additionally has
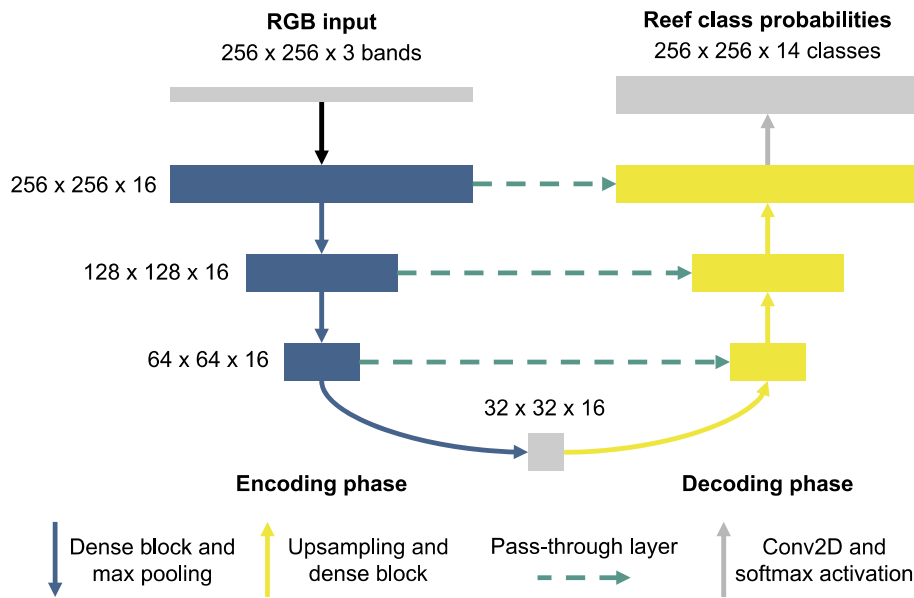
Fig. 1 Sample dense u-net architecture, with four levels of dense blocks and 16 filters. The RGB input is passed into the first layer of the network in the upper left. The first half of the network, or encoder, feeds the data through three levels of dense blocks and downsampling, i.e., max pooling, represented by the three dark blue rectangles and arrows. The bottom gray square represents the dense block and transition between the encoder on the right and decoder on the left. The second half of the network, or decoder, feeds the data through three levels of dense blocks and upsampling, represented by the light yellow rectangles and arrows, which return the data to the original spatial resolution. Finally, the features pass through one last convolution layer and a fully connected layer with softmax activations to calculate the final reef class probabilities. Note that features from the encoder are "passed-through" (dashed teal) to the decoder to be concatenated with the upsampled features prior to each dense block, helping to localize predictions
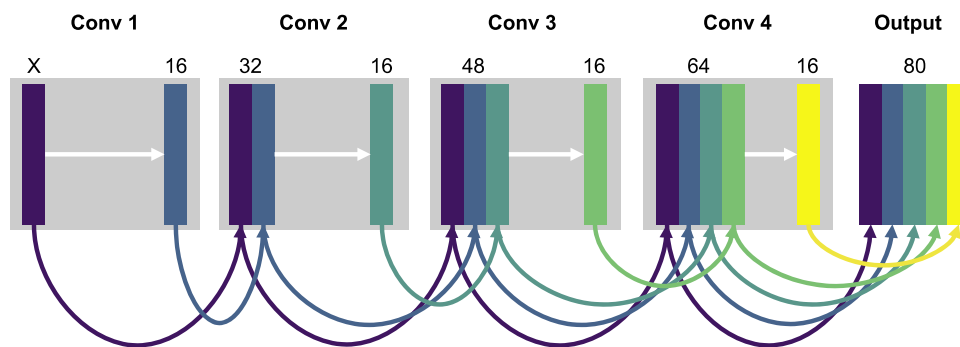
Fig. 2 Sample dense block, with four convolutions and 16 filters. The input data are passed to the first convolution on the left (gray block) and has a variable number of channels (X) depending on the network level and whether features from the pass-through layers are concatenated prior to the dense block. The input to each convolution layer is the concatenation of the input to the dense block (dark purple) and the output of each preceding convolution layer, e.g., the input to the third convolution is the concatenation of the dense block input (dark purple) and the output from convolution layers one (dark blue) and two (teal). The first operation in each convolution layer is a 1×1 convolution (white arrow) to reduce the number of channels to the specified filter number. The second operation in each convolution layer is a 3×3 convolution (gray arrow) to derive new features. Finally, the output of the dense block is the concatenation of the dense block input and all convolution layer outputs

access to the initial model inputs and inputs from each preceding layer (Fig. 2), rather than simply the outputs from the previous layer. This gives the model the opportunity to propagate information through the network across paths of varying lengths, rather than restricting information flow to a limited set of fixed-length paths, and ultimately increases model performance. The combination of these two architectures performed better than either architecture separately in early tests.

We tested a variety of dense-unet network parameterizations to find a model variant with adequate performance. In general, our models took approximately 2 to 8 h to train on modern computer graphics processing units (GPU), limiting our ability to do an extensive search of all possible architectures. We found it sufficient to conduct a grid

search to identify reasonable model architectures and hyperparameters, even though random parameter searches, Bayesian optimization approaches, or network architecture searches can improve model performance (Dernoncourt and Lee 2016). We varied the number of dense blocks at different model resolutions from two to four blocks, varied the number of convolution layers within each dense block at four, six, and eight, and varied the number of filters in multiples of two from eight to 32. We split training data into 256px × 256px images, included no growth in filter number throughout the network, used 3 × 3 convolution kernels, and both upsampled and downsampled with 2 × 2 kernels.

Model training
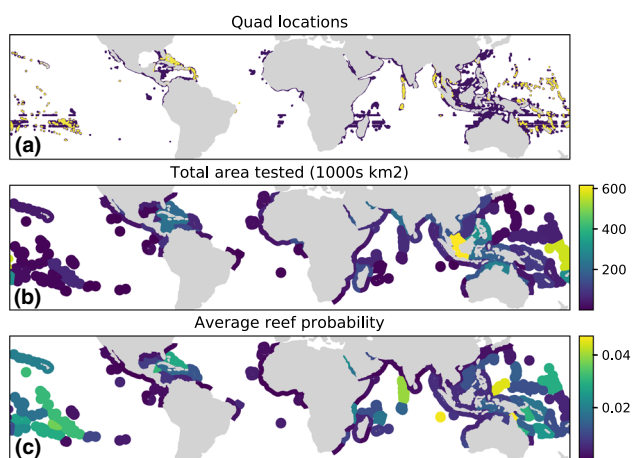
We sampled ∼ 5000 256 × 256 pixel images from the visual mosaic and response data. We restricted samples to areas within 64 pixels of reef classes, which was necessary to keep classes balanced and get a sufficient number of reef samples, given the massive amount of mainland available in the MCRMP data. In selecting model training data, we only used samples with at least 75% feature coverage (i.e., at least 75% pixels are in ocean region) and at least 10% response coverage. Asymmetric thresholds were appropriate because feature data were relatively complete, except near the boundaries of the visual mosaic, while the

Fig. 3 Global summaries of the data inputs and model outputs in the context of Marine Ecoregions of the World (MEOW). In panel a, training quads (yellow) come from a manually-selected subset of the MCRMP data, while testing quads come from all potential coral reef locations. Area is buffered for visibility and not proportional to actual coverage, as only 5% of the satellite imagery has associated training data. In panel b, the total area tested within each of the 130 MEOW ecoregions with quad coverage. In panel c, the average reef probability within each MEOW ecoregion. Please note that ecoregion summaries are not sufficient for data interpretation because values are dependent on ecoregion size, the proportion of ecoregion covered by satellite data, and the proportion of land and open water relative to reef area

MCRMP data were sparse in areas where reef segments were surrounded by unlabeled deep water.

We split samples into training and validation sets representing 90% and 10% of the data, respectively. We included 16 samples in each neural network training batch—i.e., the number of samples used for a single update of the model parameters—due to model size and augmented samples during the training process. Image augmentations included flips, rotations, and transpositions; crops and scalings; color shifts and brightenings; distortions, blurs, noise additions, and dropouts; and the addition of "fog-like" whitening. These augmentations improved model performance and are a standard practice for neural network model training (Christin et al. 2019; Ma et al. 2019). Response classes were weighted proportional to the inverse of their abundance, which over-weighted classes like reefs that appeared infrequently and under-weighted classes like land or deep water that were more common.

Global coral reef probability map through model classification

We generated a global coral reef probability map through model classification. The final layer in the CNN was configured to use softmax output activations, such that the model output is the relative probability that a pixel represents one of the 14 training classes. Because the model is based on only the three Dove satellite bands, it was limited in its ability to distinguish between spectrally-similar classes. As a result, we achieved greater accuracy by allowing the model to predict the 14 classes, and then aggregating those probabilities further into land, water, reef, and non-reef classes. We tested several methods of converting the 14 classes into the four aggregated classes, including directly in the CNN, using a maximum-likelihood approach on the CNN probabilities, and using a random forest classifier trained on the CNN probabilities. Here, we present the results for the random forest classifier because it outperformed the other two methods. We split the validation data into fivefold for training the random forest classifiers, ultimately settling on 200 trees with a max depth of 40 splits and weighted subsamples. The output of the random forest classifier is the probability that a pixel is one of the aggregated model classes (reef or non-reef (e.g., land, water, etc.); Table 1).

Global coral reef extent map and testing

We generated a global coral reef extent map from a 60% threshold of reef probability (reef: probability ≥ 60%, non-reef: probability < 60%). To test this map, we created a global coral reef field location dataset. We collected coral reef locations from public datasets, in particular, the Allen

Table 2 Comparison of Spalding et al. (2001) reef estimates to reef estimates generated from our reef probability maps

Region	Results, lower threshold	Spalding et al. (2001) (km^2)	Results, upper threshold
Global	301,110 km^2 (60%)	284,300	154,049 km^2 (65%)
Atlantic and Caribbean	25,958 km^2 (65%)	21,600	16,446 km^2 (70%)
Caribbean	24,981 km^2 (65%)	20,000	16,087 km^2 (70%)
Atlantic	2640 km^2 (60%)	1600	976 km^2 (65%)
Indo-Pacific	506,555 km^2 (55%)	261,200	257,323 km^2 (60%)
Red Sea and Gulf of Aden	18,088 km^2 (55%)	17,400	12,126 km^2 (60%)
Arabian Gulf and Sea	6327 km^2 (65%)	4200	2487 km^2 (70%)
Indian Ocean	35,630 km^2 (60%)	32,000	13,735 km^2 (65%)
Southeast Asia	180,065 km^2 (55%)	91,700	85,814 km^2 (60%)
Pacific	188,471 km^2 (55%)	115,900	104,554 km^2 (60%)
Eastern Pacific	2672 km^2 (55%)	1600	1312 km^2 (60%)

Reef extents estimates were generated for probability thresholds in 5% increments (i.e., 0%, 5%, 10%, … 95%). We report the results for the thresholds which cover the range inclusive of the Spalding estimates; e.g., the Spalding reef estimate falls between 65 and 70% for the Caribbean, but between 60 and 65% for the Atlantic. The selected regions and ordering were chosen to match the original reef estimate table from Spalding et al. (2001)
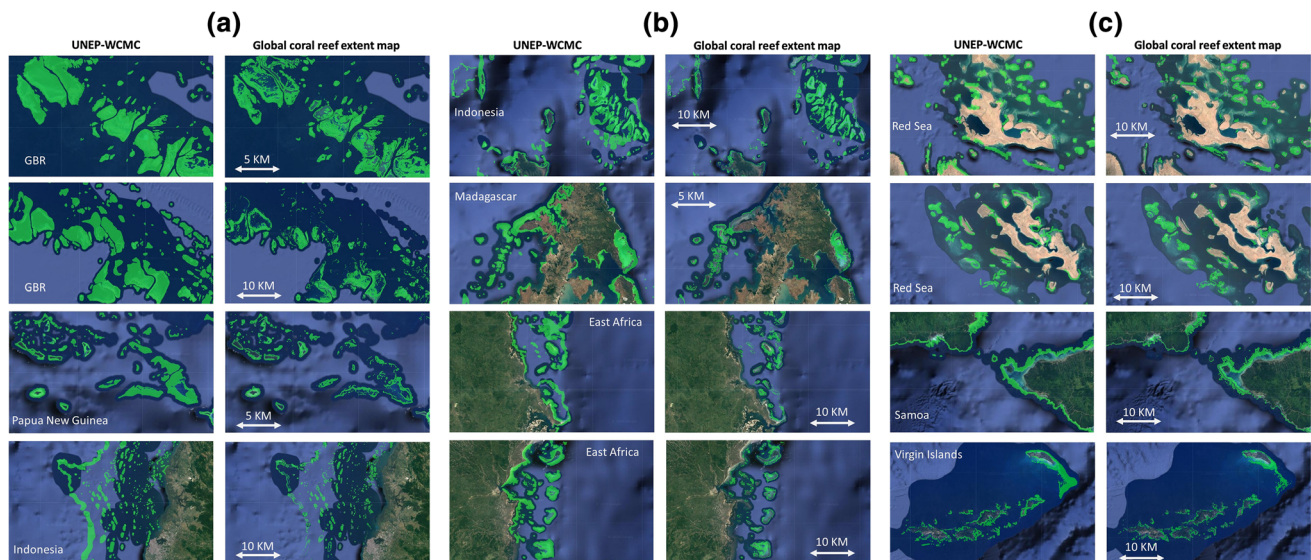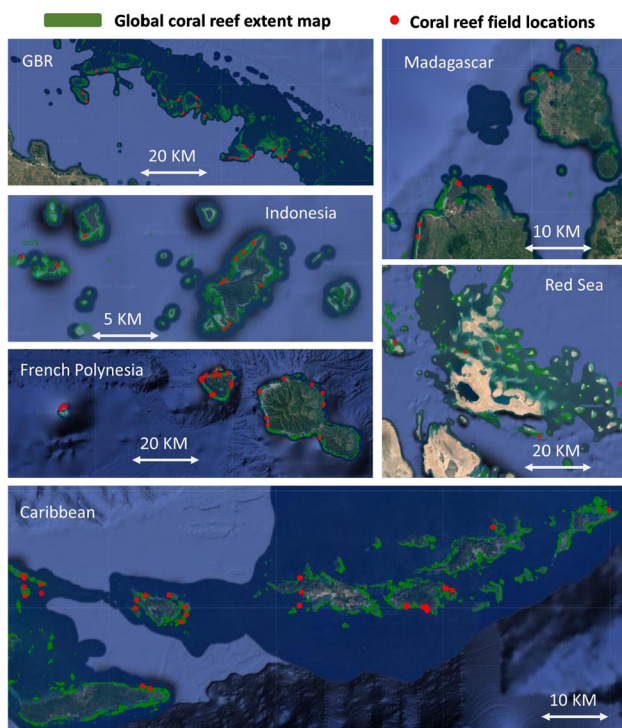

Fig. 4 Visual comparisons of the UNEP-WCMC reef map with our global coral reef extent map in different regions, including **a** Great Barrier Reef (GBR), Australia, Papua New Guinea, Indonesia; **b** Madagascar, East Africa; **c** Red Sea, Samoa, Virgin Islands

Coral Atlas (https://allencoralatlas.org/), Atlantic and Gulf Rapid Reef Assessment (https://www.agrra.org/), Khaled bin Sultan Living Oceans Foundation (https://www.livingoceansfoundation.org/), the National summary of NOAA's shallow-water benthic habitat mapping of U.S. coral reef ecosystems (https://repository.library.noaa.gov/view/noaa/748), the USGS Coral Reef Project, and Red Sea Biodiversity Surveys (http://redseabiodiversity.senckenberg.de/). We also used coral location records from published literature (Lang 2003; Rezai et al. 2004; Bertels et al. 2008; Solandt and Wood 2008; Bruckner 2011; Madduppa et al. 2012; Monaco et al. 2012; Riegl et al. 2012, 2012; Torres-Pulliza et al. 2013; Pramudya et al. 2014; Jadot et al. 2015; Hossain et al. 2016; Fujii 2017; Hafizt et al. 2017; Ampou et al. 2018; Edmunds and Kuffner 2018; Purkis et al. 2019). In total, we accessed 1952 coral reef location points (Fig. S1). Moreover, we also generated 1403 non-reef location points globally (e.g., land, water, cloud, etc.) to test our reef extent map using a confusion matrix.

Results

Our global coral reef probability map and coral reef extent map is available for download and interactive viewing from Allen Coral Atlas (allencoralatlas.org). Our global Dove mosaic quads covered coral reefs in tropical and subtropical regions (Fig. 3). Within the Marine Ecoregions of the World boundaries (MEOW) framework, the quads spanned 9 of 12 realms, 39 of 62 provinces, and 130 of 232 ecoregions (Spalding et al. 2007). In total, these quads included 10.26 million km^2 of satellite imagery represented by 121.75 billion pixels. The total area mapped in the global reef probability map sufficiently covered coral reefs globally (Fig. 3).

We tested the feasibility of our reef probability map by using it to derive reef area estimates compared to other global and regional estimates (Spalding et al. 2001, Table 2). We derived reef area by setting a probability threshold and counting pixels as reef when their reef probability exceeded that threshold. We tested probability thresholds from 0 to 95% in increments of 5%. Depending on the region, the estimates derived from a threshold between 55 and 70% aligned well with Spalding's estimates, with most of the three regions or eight subregions having thresholds very close to 60% (Table 2). Therefore, we created a global coral reef extent map using a 60% threshold.

We compared our global coral reef extent map with the UNEP-WCMC reef map worldwide (UNEP-WCMC et al. 2018). In some regions of the UNEP-WCMC map (i.e., GBR), some reef contours did not match, polygons extended beyond reef structures, and there were gaps in reef area (Fig. 4a–c). In contrast, our reef extent map appeared to better track the actual reef extent (Fig. 4a–c). These suggest that our mapping approach is especially useful in delineating reef extent details in most regions, which cumulatively will generate more accurate uses for our map at jurisdictional scales such as country- or island-level applications.

Compared with a global coral reef field dataset, our global coral reef extent map had a producer's accuracy of 87.3% (Table S2). Most of the coral reef location points from the field dataset (1704) matched with reef regions of our extent map. We plotted coral reef field location points and the reef extent map (Fig. 5, Fig. S2 to Fig. S5) and found that coral reef regions were accurately illustrated in the reef extent map. Only several deep reef locations (e.g., Fig. S5, Caribbean) did not appear on our reef extent map.

As expected, neural network architecture and parameters had a large impact on model performance and results (Table 3, Table S1). The two best performing models had F-scores of 0.66, but one had slightly higher precision

Fig. 5 Visual comparisons of the coral reef field locations with our global coral reef extent map, including Great Barrier Reef (GBR), Australia, Indonesia, Madagascar, French Polynesia, Red Sea, and Caribbean

Table 3 Classification performance results (average values using 14 classes) for select model architectures

Dense block number	Convolution layers per block	Filters	Precision	Recall	F1-score	Accuracy
2	4	16	0.68	0.38	0.49	0.91
2	6	32	0.74	0.55	0.63	0.92
2	8	24	**0.77**	0.58	**0.66**	**0.93**
3	4	24	0.69	0.50	0.58	0.92
3	6	24	0.76	0.50	0.60	0.92
3	8	32	0.63	0.36	0.46	0.90
4	4	32	0.71	0.43	0.54	0.91
4	6	8	0.67	0.44	0.53	0.91
4	8	8	0.76	**0.59**	**0.66**	**0.93**

Bold means the best performances

(0.77 vs 0.76) while the other had slightly higher recall (0.59 vs 0.58). We chose to use the model with higher precision for global application after manual review indicated that the choice was not likely to lead to qualitative differences, and because the map generation process was too computationally expensive to generate multiple maps.

Discussion

Our global coral reef extent map and coral reef probability map add value to existing reef products, as both a methodological improvement and a new source of vitally needed data. Methodological advantages include high-resolution satellite imagery, modern deep learning approaches, and consistent global methods. Our coral reef extent map provides new high-resolution (i.e., 4.5 m) spatial information of coral reefs at a global scale. Our coral reef probability map can be used as a new data source by combining the reef probabilities with other global maps to create synergistic data products, or combining the reef probabilities with other types of data layers to create derived data or analytic products. Thus, our probability map could be used in concert with other maps or datasets.

Our methods were composed of high-value and low-cost approaches to modeling the global reef extent, but these methods are not exhaustive and can be improved over time, with corresponding improvements in the generated reef probabilities. Additional data sources could be incorporated, whether they include other satellite products like Sentinel-2 or Landsat-8, or derived data products like turbidity, bathymetry, or geomorphology (Hedley et al. 2016, 2018; Hafizt et al. 2017; Kerr and Purkis 2018; Purkis 2018; Roelfsema et al. 2018, 2020; Li et al. 2019a, b; Purkis et al. 2019; Lyons et al. 2020). These additional data sources could account for shortcomings in the Dove feature data or add additional power to separate response classes. Our model could be updated to include both spatial and temporal resolutions with new network architectures (e.g., Mou et al. 2018), again helping to distinguish between response classes. Supplemental response categories could be enhanced to better handle a variety of cloud patterns, turbid water, or other confounding spectral or textural patterns. Multiple models could be generated for different regions of the globe or different types of reef formations, or incorporated in model ensembles for more accurate predictions.

Our model performances are affected by the technical (e.g., mosaicking, misalignment) or environmental (e.g., clouds, turbid water) "noise" in the satellite data, and by inaccuracies in the extent or category of reef in the MCRMP training data. Moreover, reef features could simply not be visible due to excessive depth where bottom habitats are invisible (Brando et al. 2009; Thompson et al. 2017; Li et al. 2020a), even if the satellite data and reef extents are otherwise accurate. For these reasons and potentially others, we would expect a "perfect" model to have less than perfect performance when tested against validation data.

The growing quantity and quality of Earth observation data, as well as the increasing sophistication and performance of machine learning approaches, are double-edged swords. They provide new opportunities for coral reef research and applied outcomes, but present new barriers to entry and effective and efficient progress. In addition to coral reef domain knowledge and the ability to communicate and coordinate with policymakers and stakeholders, teams need members with robust data science and software engineering skills. This project required the capacity to store on the order of tens of terabytes of data and access to hundreds of computer GPU hours. Teams of organized collaborators can produce better outcomes more quickly than individual contributors, and continued collaborations can iterate on existing datasets, codebases, and models more quickly than teams starting from scratch. We encourage others to join us in continuing to develop this resource and others.

References

Ampou EE, Ouillon S, Andréfouët S (2018) Challenges in rendering Coral Triangle habitat richness in remotely sensed habitat maps: The case of Bunaken Island (Indonesia). Marine pollution bulletin 131:72–82

Andrefouet S, Muller-Karger FE, Robinson JA, Kranenburg CJ, Torres-Pulliza D, Spraggins SA, Murch B (2006) Global assessment of modern coral reef extent and diversity for regional science and management applications: a view from space. 2:1732–1745

Aswani S, Vaccaro I (2008) Lagoon ecology and social strategies: habitat diversity and ethnobiology. Human Ecology 36:325–341

Bertels L, Vanderstraete T, Van Coillie S, Knaeps E, Sterckx S, Goossens R, Deronde B (2008) Mapping of coral reefs using hyperspectral CASI data; a case study: Fordata, Tanimbar, Indonesia. International journal of remote sensing 29:2359–2391

Beyer HL, Kennedy EV, Beger M, Chen CA, Cinner JE, Darling ES, Eakin CM, Gates RD, Heron SF, Knowlton N (2018) Risk-sensitive planning for conserving coral reefs under rapid climate change. Conservation Letters 11:e12587

Brando VE, Anstee JM, Wettle M, Dekker AG, Phinn SR, Roelfsema C (2009) A physics based retrieval and quality assessment of bathymetry from suboptimal hyperspectral data. Remote sensing of Environment 113:755–770

Brodrick PG, Davies AB, Asner GP (2019) Uncovering ecological patterns with convolutional neural networks. Trends in ecology & evolution 34:734–745

Bruckner AW (2011) The Global Reef Expedition: Science without Borders®. Journal of the Washington Academy of Sciences 19–32

Burke LM, Reytar K, Spalding M, Perry A (2017) Reefs at risk revisited: World Resources Institute

Cheng G, Han J, Lu X (2017) Remote sensing image scene classification: Benchmark and state of the art. Proceedings of the IEEE 105:1865–1883

Christin S, Hervet É, Lecomte N (2019) Applications for deep learning in ecology. Methods in Ecology and Evolution 10:1632–1644

Dernoncourt F, Lee JY (2016) Optimizing neural network hyperparameters with gaussian processes for dialog act classification. 406–413

Eakin CM, Sweatman HP, Brainard RE (2019) The 2014–2017 global-scale coral bleaching event: insights and impacts. Coral Reefs 38:539–545

Edmunds JRGPJ, Kuffner RDGIB (2018) Time-series coral-cover data from Hawaii, Florida, Mo'orea, and the Virgin Islands

Foo SA, Asner GP (2019) Scaling Up Coral Reef Restoration Using Remote Sensing Technology. Front Mar Sci 6:79

Fujii M (2017) Mapping the change of coral reefs using remote sensing and in situ measurements: a case study in Pangkajene and Kepulauan Regency, Spermonde Archipelago, Indonesia. Journal of oceanography 73:623–645

Group GC (2019) GEBCO 2019 grid

Guan S, Khan AA, Sikdar S, Chitnis PV (2019) Fully Dense UNet for 2-D Sparse Photoacoustic Tomography Artifact Removal. IEEE journal of biomedical and health informatics 24:568–576

Hafizt M, Manessa MDM, Adi NS, Prayudha B (2017) Benthic habitat mapping by combining lyzenga's optical model and relative water depth model in Lintea Island, Southeast Sulawesi. Earth and Environmental Sciences (98)

Hedley J, Roelfsema C, Chollett I, Harborne A, Heron S, Weeks S, Skirving W, Strong A, Eakin C, Christensen T (2016) Remote sensing of coral reefs for monitoring and management: a review. Remote Sensing 8:118

Hedley JD, Roelfsema C, Brando V, Giardino C, Kutser T, Phinn S, Mumby PJ, Barrilero O, Laporte J, Koetz B (2018) Coral reef applications of Sentinel-2: Coverage, characteristics, bathymetry and benthic mapping with comparison to Landsat 8. Remote sensing of environment 216:598–614

Heron SF, Maynard JA, Van Hooidonk R, Eakin CM (2016) Warming trends and bleaching stress of the world's coral reefs 1985–2012. Scientific reports 6:38402

Hossain MS, Bujang JS, Zakaria MH, Hashim M (2016) Marine and human habitat mapping for the Coral Triangle Initiative region of Sabah using Landsat and Google Earth imagery. Marine Policy 72:176–191

Hughes TP, Anderson KD, Connolly SR, Heron SF, Kerry JT, Lough JM, Baird AH, Baum JK, Berumen ML, Bridge TC (2018) Spatial and temporal patterns of mass bleaching of corals in the Anthropocene. Science 359:80–83

Iandola F, Moskewicz M, Karayev S, Girshick R, Darrell T, Keutzer K (2014) Densenet: Implementing efficient convnet descriptor pyramids. arXiv preprint arXiv:14041869

IMaRS-USF IRD (2005) Millennium coral reef mapping project. Validated maps Cambridge (UK): UNEP World Conservation Monitoring Centre

Jadot C, Darling ES, Brenier A (2015) MADAGASCAR: A Baseline Assessment of Coral Reef Fisheries. Wildlife Conservation Society

Kerr JM, Purkis S (2018) An algorithm for optically-deriving water depth from multispectral imagery in coral reef landscapes in the absence of ground-truth data. Remote Sensing of Environment 210:307–324

Lang JC (2003) Status of coral reefs in the Western Atlantic: results of initial surveys, Atlantic and Gulf Rapid Reef Assessment (AGRRA) program

Li J, Knapp DE, Schill SR, Roelfsema C, Phinn S, Silman M, Mascaro J, Asner GP (2019a) Adaptive bathymetry estimation for shallow coastal waters using Planet Dove satellites. Remote Sensing of Environment 232:111302

Li J, Schill SR, Knapp DE, Asner GP (2019b) Object-based mapping of coral reef habitats using planet dove satellites. Remote Sensing 11:1445

Li J, Fabina NS, Knapp DE, Asner GP (2020a) The Sensitivity of Multi-spectral Satellite Sensors to Benthic Habitat Change. Remote Sensing 12:532

Li X, Liu B, Zheng G, Ren Y, Zhang S, Liu Y, Gao L, Liu Y, Zhang B, Wang F (2020b) Deep learning-based information mining from ocean remote sensing imagery. National Science Review

Lyons MB, Roelfsema CM, Kennedy EV, Kovacs EM, Borrego-Acevedo R, Markey K, Roe M, Yuwono DM, Harris DL, Phinn SR, Asner, Gregory P., Li, Jiwei, Knapp, David, Fabina, Nicholas, Larsen, Kirk, Traganos, Dimosthenis, Murray, Nicholas (2020) Mapping the world's coral reefs using a global multiscale earth observation framework. Remote Sensing in Ecology and Conservation

Ma L, Liu Y, Zhang X, Ye Y, Yin G, Johnson BA (2019) Deep learning in remote sensing applications: A meta-analysis and review. ISPRS journal of photogrammetry and remote sensing 152:166–177

Madduppa HH, Ferse SC, Aktani U, Palm HW (2012) Seasonal trends and fish-habitat associations around Pari Island, Indonesia: setting a baseline for environmental monitoring. Environmental biology of fishes 95:383–398

McManus J (1994) ReefBase-a global database of coral reef systems and their resources. Naga, the ICLARM Quarterly 17:16

Monaco ME, Andersen SM, Battista TA, Kendall MS, Rohmann SO, Wedding LM, Clarke AM (2012) National summary of NOAA's shallow-water benthic habitat mapping of US coral reef ecosystems

Montaggioni LF, Braithwaite CJ (2009) Quaternary coral reef systems: history, development processes and controlling factors. Elsevier

Mou L, Bruzzone L, Zhu XX (2018) Learning spectral-spatial-temporal features via a recurrent convolutional neural network for change detection in multispectral imagery. IEEE Transactions on Geoscience and Remote Sensing 57:924–935

Pendleton L, Comte A, Langdon C, Ekstrom JA, Cooley SR, Suatoni L, Beck MW, Brander LM, Burke L, Cinner JE (2016) Coral reefs and people in a high-CO2 world: where can science make a difference to people? PloS one 11:e0164699

Pramudya FS, Wikantika K, Windupranata W (2014) Satellite-Based Benthic Habitat Mapping Using LANDSAT 8 in Nusa Lembongan and Nusa Ceningan Island. Faculty of Earth Science and Technology Institut Teknologi Bandung

Purkis SJ (2018) Remote Sensing Tropical Coral Reefs: The View from Above. Annual review of marine science 10:149–168

Purkis SJ, Gleason AC, Purkis CR, Dempsey AC, Renaud PG, Faisal M, Saul S, Kerr JM (2019) High-resolution habitat and bathymetry maps for 65,000 sq. km of Earth's remotest coral reefs. Coral Reefs 38:467–488

Rezai H, Wilson S, Claereboudt M, Riegl B (2004) Coral reef status in the ROPME sea area: Arabian/Persian Gulf, Gulf of Oman and Arabian Sea. Status of coral reefs of the world 1:155–170

Riegl BM, Bruckner AW, Rowlands GP, Purkis SJ, Renaud P (2012) Red Sea coral reef trajectories over 2 decades suggest increasing community homogenization and decline in coral size. PLoS One 7:e38396

Roelfsema C, Kovacs E, Ortiz JC, Wolff NH, Callaghan D, Wettle M, Ronan M, Hamylton SM, Mumby PJ, Phinn S (2018) Coral reef habitat mapping: A combination of object-based image analysis and ecological modelling. Remote Sensing of Environment 208:27–41

Roelfsema CM, Kovacs EM, Ortiz JC, Callaghan DP, Hock K, Mongin M, Johansen K, Mumby PJ, Wettle M, Ronan M (2020) Habitat maps to enhance monitoring and management of the Great Barrier Reef. Coral Reefs 1–16

Ronneberger O, Fischer P, Brox T (2015) U-net: Convolutional networks for biomedical image segmentation. 234–241

Solandt J-L, Wood C (2008) Maldives reef survey-June 13–30th, 2008. Marine Conservation Society, Ross on Wye, UK 1–18

Spalding M, Burke L, Wood SA, Ashpole J, Hutchison J, Zu Ermgassen P (2017) Mapping the global value and distribution of coral reef tourism. Marine Policy 82:104–113

Spalding M, Spalding MD, Ravilious C, Green EP (2001) World atlas of coral reefs. Univ of California Press

Spalding MD, Fox HE, Allen GR, Davidson N, Ferdaña ZA, Finlayson MAX, Halpern BS, Jorge MA, Lombana AL, Lourie SA (2007) Marine ecoregions of the world: a bioregionalization of coastal and shelf areas. BioScience 57:573–583

Thompson DR, Hochberg EJ, Asner GP, Green RO, Knapp DE, Gao B-C, Garcia R, Gierach M, Lee Z, Maritorena S (2017) Airborne mapping of benthic reflectance spectra with Bayesian linear mixtures. Remote sensing of Environment 200:18–30

Torres-Pulliza D, Wilson JR, Darmawan A, Campbell SJ, Andréfouët S (2013) Ecoregional scale seagrass mapping: A tool to support resilient MPA network design in the Coral Triangle. Ocean & coastal management 80:55–64

UNEP-WCMC, WorldFish Centre, WRI, TNC (2018) Global distribution of warm-water coral reefs, compiled from multiple sources including the Millennium Coral Reef Mapping Project. Version 4.0. Includes contributions from IMaRS-USF and IRD (2005), IMaRS-USF (2005) and Spalding et al. (2001). Cambridge (UK): UN Environment World Conservation Monitoring Centre. URL: http://data.unep-wcmc.org/datasets/1

Van Hooidonk R, Maynard J, Tamelander J, Gove J, Ahmadia G, Raymundo L, Williams G, Heron SF, Planes S (2016) Local-scale projections of coral reef futures and implications of the Paris Agreement. Scientific reports 6:1–8

Zhang Z, Liang X, Dong X, Xie Y, Cao G (2018) A sparse-view CT reconstruction method based on combination of DenseNet and deconvolution. IEEE transactions on medical imaging 37:1407–1417