

ORIGINAL RESEARCH

Deep learning for coastal resource conservation: automating detection of shellfish reefs

Justin T. Ridge¹ , Patrick C. Gray¹ , Anna E. Windle² & David W. Johnston¹ ¹Division of Marine Science and Conservation, Nicholas School of the Environment Duke University Marine Laboratory, 135 Duke Marine Lab Rd., Beaufort North Carolina 28516²University of Maryland Center for Environmental Science, Horn Point Laboratory, 2020 Horns Point Rd, Cambridge Maryland 21613**Keywords**

convolutional neural network, deep learning, drones, habitat classification, oyster reefs, semantic segmentation

Correspondence

Justin T. Ridge, Duke University Marine Laboratory, 135 Duke Marine Lab Rd, Beaufort, NC 28516. Tel: 812-322-6857; Fax: 252-504-7648; E-mail: justin.ridge@duke.edu

Editor: Ned Horning

Associate Editor: Dimitris Poursanidis

Received: 5 July 2019; Revised: 1 October 2019; Accepted: 4 November 2019

doi: 10.1002/rse2.134

Remote Sensing in Ecology and Conservation 2020; **6** (4):431–440**Introduction**

Coastal and estuarine landscapes are biogenically structured mosaics that include both vegetated (marshes, mangroves, seagrasses, etc.) and animal-derived (shellfish reefs, corals, etc.) habitats. These structured habitats provide a number of ecosystem services that include essential nursery habitat, carbon sequestration, improved water quality, and shoreline protection (Barbier et al. 2011; Spalding et al. 2014; Lefcheck et al. 2019). However, these landscapes are at risk from the combined stress of climatic and direct human-driven changes (IPCC, 2014), making effective management and conservation of estuarine and coastal environments increasingly crucial. Unfortunately, the spatial complexity and regular submergence of these areas makes it challenging to assess their vulnerability on a broadscale using remotely sensed datasets, which can be spectrally and radiometrically inconsistent across sensor platforms, and may not be captured at appropriate spatial or temporal resolutions (Jollineau and Howarth 2008; Corbane et al. 2015).

Abstract

It is increasingly important to understand the extent and health of coastal natural resources in the face of anthropogenic and climate-driven changes. Coastal ecosystems are difficult to efficiently monitor due to the inability of existing remotely sensed data to capture complex spatial habitat patterns. To help managers and researchers avoid inefficient traditional mapping efforts, we developed a deep learning tool (OysterNet) that uses unoccupied aircraft systems (UAS) imagery to automatically detect and delineate oyster reefs, an ecosystem that has proven problematic to monitor remotely. OysterNet is a convolutional neural network (CNN) that assesses intertidal oyster reef extent, yielding a difference in total area between manual and automated delineations of just 8%, attributable in part to OysterNet's ability to detect oysters overlooked during manual demarcation. Further training of OysterNet could enable assessments of oyster reef heights and densities, and incorporation of more coastal habitat types. Future iterations will be applied to high-resolution satellite data for effective management at larger scales.

Biogenic shellfish reefs (e.g. oysters) are one such coastal ecosystem that are ecologically and economically valuable globally (Grabowski et al. 2012) but have been imperiled by numerous natural and anthropogenic impacts over the last century (Beck et al. 2011). Habitat mapping and population assessments of these environments remain problematic due to the expansive estuarine conditions the oysters can occupy (e.g. brackish to saline, subtidal to intertidal, etc.) and the unreliability of delineating oyster habitats from aerial and satellite imagery (Grizzle et al. 2002; Schill et al. 2006). Oyster reefs in remotely sensed imagery often appear too similar to adjacent mudflats or salt marshes, which confound the application of standard segmentation and classification tools within GIS programs (Fig. 1). In some cases, this necessitates manually intensive, prolonged field sampling initiatives to effectively map entire coasts (Jensen et al. 2014). The use of advanced processing techniques or novel sensors (e.g. hyperspectral imagery) has offered some success for remote detection and measurement of oyster reefs (Schill et al. 2006; Choe et al. 2012; Le Bris et al. 2016;

Grizzle et al. 2018) but has not been effectively implemented in broadscale management practices due to the specialized nature of data collection, intensive data processing, and regional variations where oyster reefs exist within the tidal frame (fully intertidal to subtidal).

While shellfish reef ecosystems may not be easily discernible in satellite imagery, one method to solve this issue is the use of higher resolution imagery available through unoccupied aircraft systems (UAS; or drone systems). UAS encompasses the aircraft (unmanned aerial vehicle, UAV) and the components necessary to operate the aircraft as well as the sensor payload. Over the last two decades, UAS implementation in environmental research has experienced an exponential increase (Manfreda et al. 2018). Imagery from UAS technology rivals resolutions achieved from occupied aircraft while reducing costs and safety risks as well as providing more flexibility with flight planning (Joyce et al. 2018; Johnston 2019). With increased resolution, differences between similar environments become more pronounced and challenging classes become more easily separable. This added resolution has proven beneficial for seagrass mapping (Nahirnick et al. 2019) and effective as a source for refining habitat classification training data in lower resolution satellite data (Gray et al. 2018). The cryptic nature of shellfish reef habitat in remotely sensed imagery makes it an ideal candidate to test the limits of combining high-resolution imagery with deep learning-based habitat classification systems in conservation science.

Emerging methods in deep learning, a subfield of machine learning, have demonstrated potential to address environmentally and ecologically pressing issues in the coastal zone (Gray et al. 2019). Convolutional neural networks (CNNs) are a class of deep neural networks that scan over an image, building hierarchical representations of the input data by composing layers of nonlinear transformations, allowing it to learn features (e.g. lines, shapes, gradients, oyster reef attributes) for classification or regression from the input data automatically, rather than needing an expert image analyst to manually engineer these features. CNNs are proving capable of analyzing vast image catalogs and identifying or classifying objects of interest with near-human accuracy (Lecun et al. 2015). Providing a CNN with a training dataset of labeled images (thousands to millions of examples for most neural networks) allows the network to iteratively update the weights of adjustable parameters of its layers of transformations to find an error minimum in the predicted labels compared to the target labels. A major advantage CNNs have over other image classification algorithms is the ability to not just factor in spectral characteristics of the subjects but also learn




Figure 1. Oyster reef imagery comparison of RapidEye, WorldView-3, and unoccupied aircraft systems (UAS). Satellite imagery courtesy of DigitalGlobe Foundation and Planet, Inc.

complex spatial attributes and relationships among them, essentially “understanding” a subject’s context within an image (Maggiori et al. 2016; Kemker et al. 2018). This gives the CNN model some semblance of a landscape context to better delineate habitats that appear similar in color and texture but occur in different locations within a regional mosaic.

Here we demonstrate the application of a CNN to automate the assessment of a coastal ecosystem, intertidal oyster reefs, to rapidly determine habitat extent within an estuarine managed area. Specifically, we use high-resolution imagery obtained from a fixed-wing UAS to train a CNN for detecting and delineating intertidal reefs of eastern oyster (*Crassostrea virginica*) using one site of the North Carolina National Estuarine Research Reserve network – the Rachel Carson Reserve (RCR) adjacent to Beaufort, NC (Fig. 2). We also assess efficacy of the CNN model creation at different resolution scales as well as the accuracy of the final model output against manual oyster reef delineation.

Materials and Methods

Intertidal oyster reef surveys

We conducted UAS flights over the RCR during spring 2018 coincident with low tidal conditions, when most oyster reefs in the area are above water, and reduce glare with more favorable sun angles (solar elevations < 60 degrees) to maximize our reef detection potential. The RCR is a complex of islands (~10 km²) within the estuarine Back Sound including the gamut of coastal habitats (oyster reef, salt marsh, seagrass, tidal flats, dune grasses and shrubland, and maritime forest). High-resolution RGB (red, green, blue) imagery was collected with a senseFly eBee Plus fixed-wing aircraft equipped with a Sensor Optimized for Drone Applications (S.O.D.A.) camera, flown at approximately 100-m altitude for a ground sampling distance (GSD) of 2.2 cm/pixel with 75% and 80% lateral and longitudinal image overlap, respectively. The eBee Plus is equipped with a survey-grade GPS system that allows for real-time kinematic (RTK) or post-

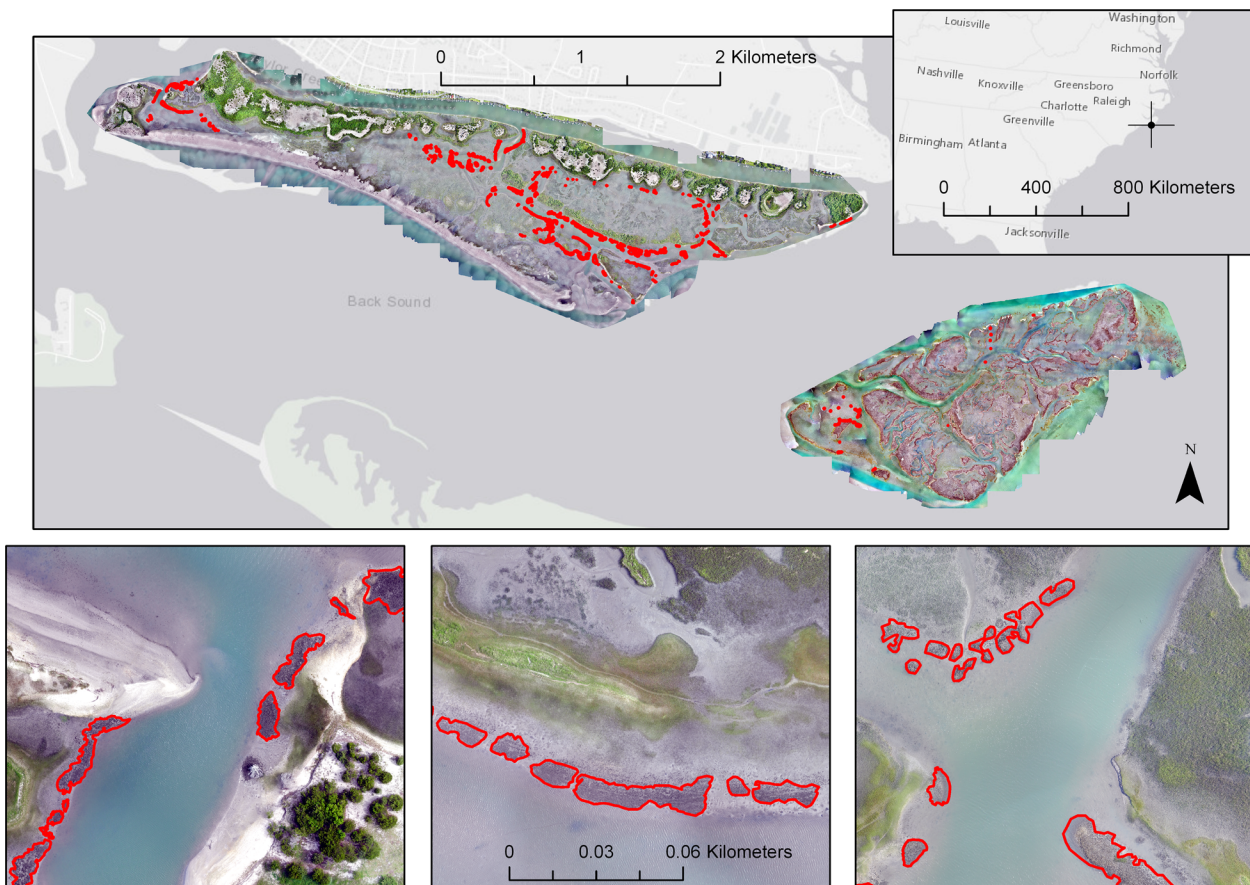


Figure 2. Study area map of the Rachel Carson Reserve from UAS orthoimagery with example areas where oyster reefs fit into the landscape mosaic.

processed kinematic (PPK) corrections to image geotags, providing subdecimeter horizontal accuracies (Seymour et al. 2018).

We processed images with Pix4D photogrammetry software to create georectified orthomosaics (Fig. 2). We then manually defined oyster reefs with individual polygons in ArcGIS v10.4 for the entire RCR (Fig. 3) where inside of a polygon was labeled “oyster” and outside was labeled “background.” This process required approximately 14 h of personnel time. To assess the CNN’s efficacy at different spatial scales, orthomosaics and respective polygon labels were tiled and parsed into image catalogs at three resolutions (1000 × 1000 pixels, 2000 × 2000 pixels, and 4000 × 4000 pixels with 826, 343, and 163 tiles, respectively, see Fig. 4 for examples of 2000 × 2000 tiles with label polygons). Each image catalog was split for model development (80% of images) and testing (20% of images) (Fig. 4). The model development dataset was further split for training (80% of images) and validation (20% of images). Using best practices to reduce model overfitting (Simard et al. 2003), we artificially increased our dataset size by applying random augmentations (e.g.

blurring, noise addition, adjusted contrast, translation, rotation, etc.) to images within the training dataset using the python tool *imgaug* (<https://github.com/aleju/imgaug>), effectively increasing our training dataset sixfold. Preparing data for training the OysterNet CNN required approximately 1.5 h of personnel time per dataset.

OysterNet architecture

We conducted all CNN development in the python programming language using the Keras framework and Tensorflow as a backend. Training and testing were done on a machine running Ubuntu 16.04 with an Intel i7-7800X CPU, 32GB of RAM, and an NVIDIA Tesla Titan Xp 12GB Graphics Processing Unit (GPU). All CNN code, analysis scripts, and a Docker container for replicating the development environment are available online (https://github.com/patrickcgray/oyster_net). OysterNet is based on Mask R-CNN (He et al. 2017), a flexible CNN architecture with excellent performance in instance segmentation. OysterNet was initially trained on the Microsoft Common Objects in Context (COCO; Lin et al.

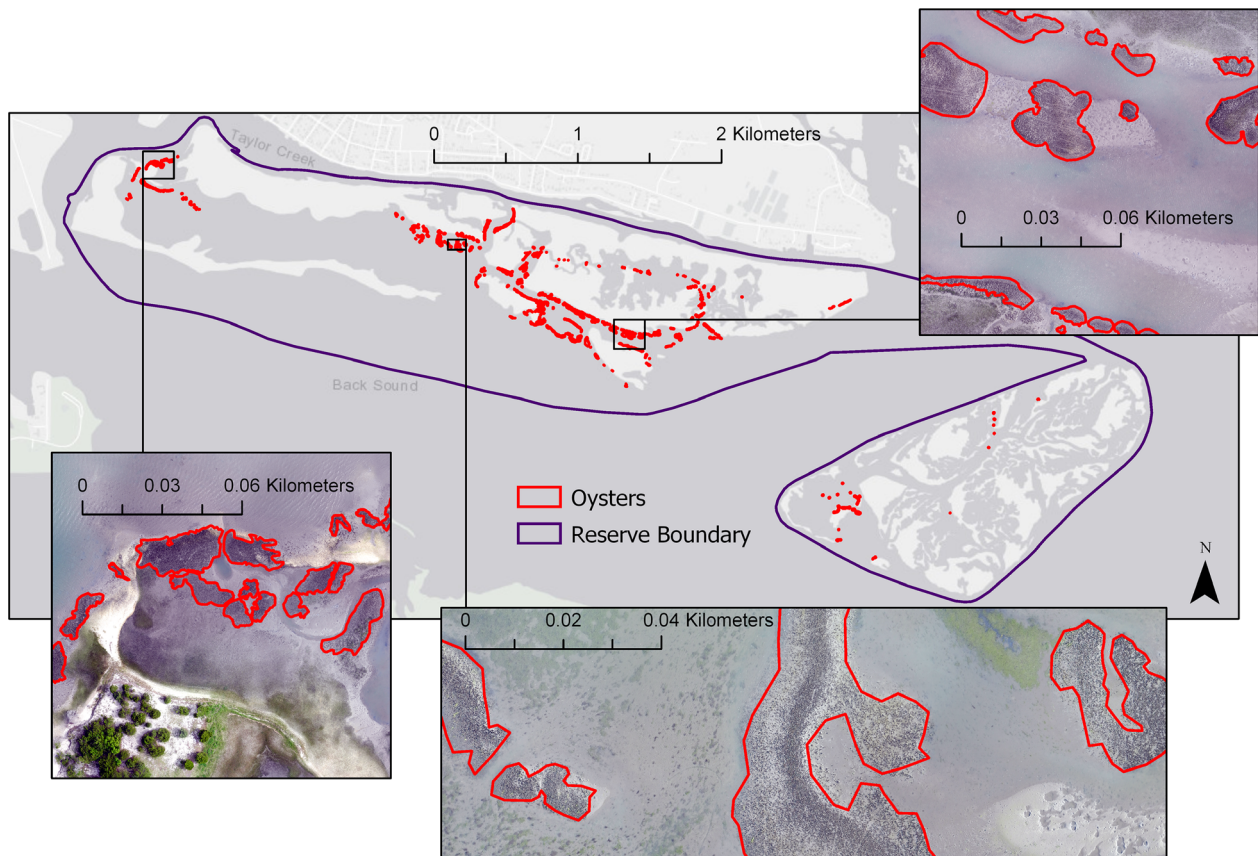


Figure 3. Extent of the Rachel Carson Reserve manually delineated oyster reef polygons. Insets provide examples of oyster reef polygons overlaying UAS imagery.

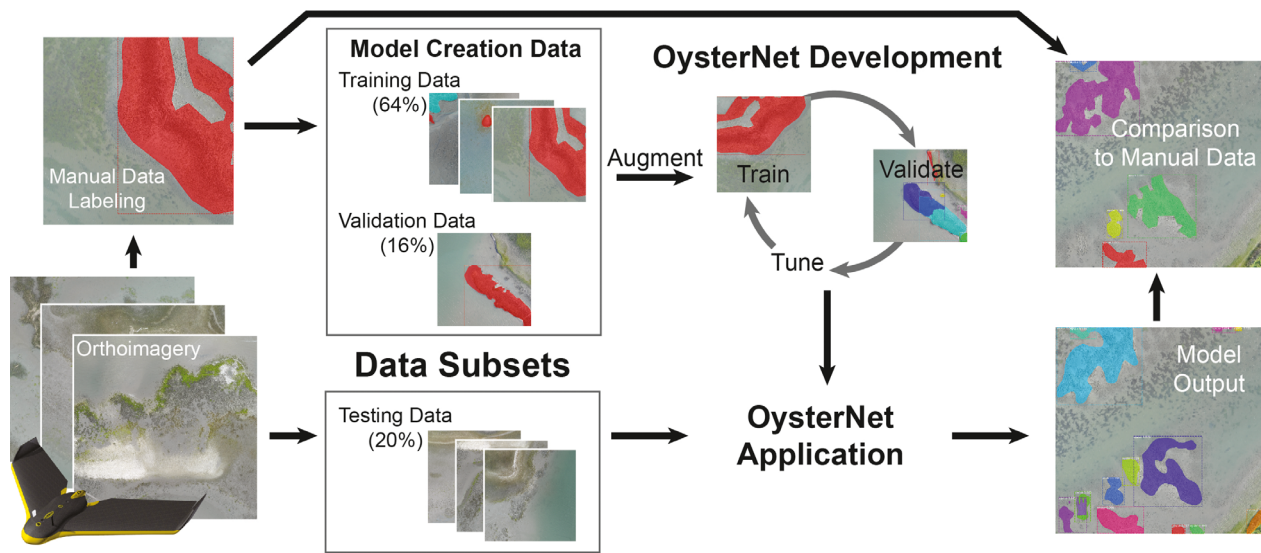


Figure 4. Conceptual model of workflow from tiled orthoimagery through final OysterNet output analysis. Fixed-wing aircraft pictured is a senseFly eBee.

2014) dataset before being “fine-tuned” on our oyster imagery. COCO, an open-access dataset of over 200 000 generic labeled images, allowed us to leverage transfer learning, the ability of neural networks to use features learned on one dataset on an entirely different task (Razavian et al. 2014; Yosinski et al. 2014). Based on the tiled image catalogs we created, each training sample for the CNN was a x by y by 3 tensor (R,G,B) with an accompanying polygon designation if an oyster reef was present within the image.

We independently trained OysterNet on each of the three image size datasets. During model training, the validation dataset is used to assess training efficacy during each update to the parameter weights (these update points are called epochs) by measuring both how well the model is minimizing loss (how well the model predictions fit the true labels) and how the loss on the training data compares to the loss on validation data to prevent overfitting (overly memorizing the training data inputs and not generalizing to new data). The CNN architecture has dozens of hyperparameters (or “knobs”) that can be tweaked to alter how the CNN views data, learns, and makes predictions. We only diverged from the hyperparameters reported in the initial Mask R-CNN paper where necessary to better suit the spatial structure of aerial imagery (loss function, training and validation loss during each training run, and hyperparameters available in Appendix A). Training for each network took approximately 6 h and was run for just over 150 epochs. To reduce overfitting, this training was done in three rounds with minor changes made during each iteration to increase the number of trainable layers and

decrease learning rate (please refer to Data S1 for more details).

OysterNet testing

Final OysterNet CNNs were applied to the test datasets, which had no overlap with the training and validation datasets. We compared OysterNet model outputs on a pixel-to-pixel basis from the test imagery against our manually delineated reefs (drawn polygons in ArcGIS; Fig. 4) using several metrics to determine the efficacy of the trained models including accuracy, precision, and recall. Accuracy of a CNN is the proportion of true positive and negative predictions over the total possible observations. Precision is the ratio of true positive predictions over the total number of predictions made. Recall is the proportion of true positive predictions over the total number of true positives. While accuracy is often a good metric for general performance, in a detection problem, where the class of interest is only a small proportion of total area, it can be misleading. For example, if only one in a hundred pixels is your class of interest, the model could predict every pixel background and have 99% accuracy but essentially be useless. Recall and precision are more informative when we are concerned about false positives (precision) or false negatives (recall). Output polygons from the OysterNet CNN include a confidence value between 0.0 and 1.0 for each prediction, allowing us to set a threshold (or probability cutoff), which determines whether the predicted oyster reef is included in the final output maps. Deciding on this threshold changes how conservative the output is and impacts the overall efficacy

metrics (e.g. allowing a user to trade precision for recall with a lower threshold). Our threshold was determined empirically from the recall/precision curve on the 2000×2000 model (Fig. 5) to be 0.98. Individual and overall areas of oyster reefs were generated based on the overall pixel count within classified reef matrices referenced back to the GSD of the UAS imagery. In addition to quantitative metrics, false positives and negatives were manually examined to assess validity and potential explanations for errors.

Results

The best balance of recall and precision achieved by the OysterNet CNN was accomplished with the 2000×2000 model (Fig. 5) achieving a precision of 0.771 and recall of 0.772. While the area of manually delineated oyster reefs within the test dataset totaled 8950.50 m^2 (0.90 ha), the 2000×2000 OysterNet model generated a predicted oyster reef area of 9657.27 m^2 (0.96 ha), equivalent to an 8% difference in area (Fig. 6).

Processing imagery through the OysterNet CNN took 2 seconds per image regardless of image resolution because they were all resampled to 1024×1024 resolution. Thus, for the 2000×2000 model at 2.2 cm/pixel, the CNN processed 0.0968 hectare per second. At this rate, the trained OysterNet CNN could determine the oyster reef distribution within the entire RCR ($\sim 10 \text{ km}^2$) in under 3 h. The time it would require to manually




Figure 5. Precision vs. Recall is presented for the 4000×4000 pixel model (green), the 2000×2000 pixel model (orange), and the 1000×1000 model (blue) showing the 2000×2000 model as the clear top performer. The 2000×2000 model curve shows higher precisions and lower recall at higher confidence thresholds. Precision decreases and recall increases as the threshold decreases. Setting this threshold allows a user to tune a model for detecting more oyster reef at the cost of more false positives and likely analyst time vetting false positives, or minimize false positives at the risk of missing many true positives.

delineate reef areas with an RTK-GPS backpack system while transiting sandflats and tidal channels within the reserve is in the order of 8 h for 0.5 km^2 , which has been conducted for prior work within the RCR (Fodrie et al. 2014; Rodriguez et al. 2014). Considering that oyster reefs are interspersed through just under half the RCR ($\sim 4 \text{ km}^2$) when major stretches of supratidal land are omitted, it would require a minimum of 8 work days to manually delineate all the reefs within the RCR (Table 1).

Discussion

This present study provides the first example of conducting semantic image segmentation with a CNN to classify and measure an economically and ecologically valuable intertidal habitat. While there are some initial time costs involved in training and refining the model, the true benefit of OysterNet emerges from the combination of rapid and cost-effective UAS-based data collection with a pre-trained, accurate model that can automatically assess large areas in just a few hours compared with the multiple days to weeks required to manually delineate reefs from occupied aircraft or satellite imagery, or the months to years required when employing purely field-based approaches. Other research has also found that deep learning methods are dramatically reducing the time required for ecological monitoring (Wäldchen and Mäder 2018; Weinstein 2018; Chilson et al. 2019), and continual advancements in computational efficiency (e.g. more capable GPUs) will further decrease the time necessary for CNNs like OysterNet to analyze large and rapidly collected image datasets. Even without waiting for these advancements, OysterNet can easily be parallelized on multiple machines to rapidly detect oyster reefs across vast areas of orthoimagery.

Our results also reveal considerations for CNN-based semantic segmentation approaches in terms of the spatial context of training imagery and the overall fidelity of reef edges. Given the degraded model performance when using the 1000×1000 pixel image size for training, the added spatial context of the 2000×2000 pixel images appears to be more important than the addition of training data using smaller image sizes. While the 4000×4000 model also had considerably worse performance than the 2000×2000 version, this was potentially due to a lack of training examples. Because the same total area of UAS imagery was used to create the tiles, the 4000×4000 version had 1/4th the number of tiles, and this may have led to degraded performance. Additionally, due to memory requirements, all tiles used by our CNN were resized to 1024×1024 , so the downscaled 4000×4000 tiles (resulting GSD of $8.8 \text{ cm pixel}^{-1}$) may not have had the spatial resolution to accurately predict oyster habitat given the available training data and limited spectral




Figure 6. Comparison of orthoimagery tiles (top; A, D, G, J, M) with corresponding manually delineated oyster polygons (middle; B, E, H, K, N) and automatic polygon creation from the 2000 × 2000 OysterNet convolutional neural network (bottom; C, F, I, L, O). While overall polygons are similar, small variations can be seen. Automatic detections in (C) show an inclusion of sparser oyster that was not included in (B). Panels (I), (L) and (O) include small patches that were not included in manual delineations (H), (K) and (N) as well as a small section of false-positive polygon in the bottom right of (I) and top center of (O) that appear to have a similar texture to oyster reefs but is scrub-shrub in (I) and eroding marsh in (O). Colors indicate the order of polygon creation, which was arbitrary for both manual and CNN methods but helps differentiate detections.

bands. Further tests using the same number of tiles at each spatial scale will provide more insight into the trade-offs of resolution and spatial context. Doubtless though, increased training data will prove beneficial for future models.

Discrepancies in observed and predicted reef areas, presented in the recall and precision metrics, require careful consideration (Fig. 6) and may not be indicative only of poor performance of OysterNet. Manual bounding of reefs is subject to user error when determining

Table 1. Time required to delineate oyster reefs using manual GPS delineation, digital delineation, and automated delineation through OysterNet

Reef delineation method	Time required per 4 km ²
OysterNet (CNN)	13.15 h
Digital (GIS)	26 h
Manual (RTK-GPS)	64 h

OysterNet and digital delineation times incorporate a conservative 8 h required for acquiring imagery and a further 4 h for processing imagery. Timing is also based on the trained OysterNet.

reef edges, especially when analysts address sparse and disjointed portions of oyster habitat. Considering that the OysterNet CNN detected reefs that were missed during manual delineation, the 8% overestimate in reef area is likely closer to the true coverage of reefs within the RCR. It is a challenge for analysts to determine what lower bound of oyster density should still be classified as a discrete reef, and our training data excluded some portions of the RCR shoreline with sparsely clustered oysters. As future work and a potential solution to this issue, with sufficient examples of field measured density, the present system could be expanded to a multitask CNN that first isolates oyster habitat, and then regresses a continuous density value based on a density-labeled training dataset.

The results of the present study show great promise for broader application. The overall area parsed for training and testing was relatively small (<10 km²), yet yielded performance acceptable for management purposes at regional scales. Building the training dataset to encompass a larger area will enhance the CNN's efficacy. Even with a relatively small amount of training's data, the CNN was

able to isolate shallow subtidal oysters at a centimeter scale, which is promising for the applicability of a CNN over a range of oyster reef conditions spanning a greater geographic area and not necessarily relying on reefs being exposed during low tide. Given OysterNet's ability to define shallow submerged reefs, it is possible that flights done at higher tides may have still delineated reefs reasonably well; however, it is likely that we would lose definition of the deeper reefs. However, the regional variability in tidal growth dynamics and landscape composition may necessitate region-specific CNNs, perhaps relying on transfer learning from our initial instance of OysterNet if environmental conditions are not drastically different. In similar environments, like across the southeastern US, OysterNet may already be broadly accurate with some additional levels of accuracy achieved by supplementing our training pool with local data. It is possible that a trained CNN using data from a broader geographic span could accurately identify oyster reefs regardless of regional environmental differences. In the present study, labeling was limited to one type of intertidal species. Further training with imagery labeled with other habitat components (e.g. marsh grass species, submerged aquatic vegetation, etc.) would provide more detailed habitat assessments rather than the binary output currently provided.

Enhancing RGB training data with other synoptic datasets may increase OysterNet's performance. For example, digital surface models (DSMs) generated via Structure from Motion (SfM) processing of UAS data can provide a three-dimensional texture layer to help differentiate environments that exhibit similar spectral signatures in 2D imagery (Långkvist et al. 2016) and may assist with identifying oyster reefs. If the UAS orthoimagery is georectified within a vertical datum, the subsequent digital elevation model (DEM) could help OysterNet learn the elevation ranges that oyster reefs occupy within a landscape, enhancing detection. There are some caveats with using these data layers on submerged oysters, considering that SfM processing struggles when water is present (Joyce et al. 2018).

OysterNet may also be extremely useful for application to very high resolution (VHR) multispectral satellite imagery, potentially further reducing time requirements for data collection and vastly expanding the footprint. Indeed, data fusion techniques using UAS-derived imagery, where it is easier to clearly delineate oyster reefs from surrounding habitats (Gray et al. 2018), could generate more accurate and precise labels for VHR satellite imagery training datasets.

OysterNet results support the ongoing successful movement toward automated classification of remotely sensed coastal environments, particularly with cryptic intertidal

habitats, including other ecologically or economically valuable habitat components. Deep learning methods such as this will soon become powerful instruments for broad-scale coastal zone management. Future iterations of OysterNet will be a valuable tool for oyster fisheries management and broad estuarine ecosystem monitoring through rapid population assessments and change detection.

Acknowledgments

We thank A. Seymour and E. Newton for data collection in the field and for initial quality control processing of UAS imagery, and thank J. Dale for logistical support and UAS testing and maintenance. We also thank B. Puckett, P. Gilikin, K. Dobroski, and C. Atkins-Davis for assistance during data collection. Funding for this work was provided by the North Carolina Division of Marine Fisheries (Marine Resources Fund, Grant #2017-H-068), North Carolina Sea Grant/Space Grant Fellowship program (Grant #2017-R/MG-1710), and the Duke Bass Connections program. Cloud-based computational portions of this project were funded by the Microsoft AI for Earth program and local computing was supported by the NVIDIA Corporation through the donation of a Titan Xp GPU.

Data Accessibility

All imagery and associated data are stored on Duke Digital Repository at <https://doi.org/10.7924/r4cv4gx0h>. All CNN code, analysis scripts, and a Docker container for replicating the development environment are available online at <https://doi.org/10.5281/zenodo.3464119> https://github.com/patrickcgray/oyster_net.

References

- Barbier, E. B., S. D. Hacker, C. Kennedy, E. W. Koch, A. C. Stier, and B. R. Silliman. 2011. The value of estuarine and coastal ecosystem services. *Ecol. Monogr.* **81**, 169–193. <https://doi.org/10.1890/0012-9615-81.2.169>
- Beck, M. W., R. D. Brumbaugh, L. Airoidi, A. Carranza, L. D. Coen, C. Crawford, et al. 2011. Oyster reefs at risk and recommendations for conservation, restoration, and management. *Bioscience* **61**, 107–116. <https://doi.org/10.1525/bio.2011.61.2.5>
- Chilson, C., K. Avery, A. McGovern, E. Bridge, D. Sheldon, and J. Kelly. 2019. Automated detection of bird roosts using NEXRAD radar data and Convolutional Neural Networks. *Remote Sens. Ecol. Conserv.* **5**, 20–32. <https://doi.org/10.1002/rse2.92>
- Choe, B. H., D. J. Kim, J. H. Hwang, Y. Oh, and W. M. Moon. 2012. Detection of oyster habitat in tidal flats using

- multi-frequency polarimetric SAR data. *Estuar. Coast. Shelf Sci.* **97**, 28–37. <https://doi.org/10.1016/j.ecss.2011.11.007>
- Corbane, C., S. Lang, K. Pipkins, S. Alleaume, M. Deshayes, V. E. García Millán, et al. 2015. Remote sensing for mapping natural habitats and their conservation status - New opportunities and challenges. *Int. J. Appl. Earth Obs. Geoinf.* **37**, 7–16. <https://doi.org/10.1016/j.jag.2014.11.005>
- Fodrie, F. J., A. B. Rodriguez, C. J. Baillie, M. C. Brodeur, S. E. Coleman, R. K. Gittman, et al. 2014. Classic paradigms in a novel environment: inserting food web and productivity lessons from rocky shores and saltmarshes into biogenic reef restoration. *J. Appl. Ecol.* **51**, 1314–1325. <https://doi.org/10.1111/1365-2664.12276>
- Grabowski, J. H., R. D. Brumbaugh, R. F. Conrad, A. G. Keeler, J. J. Opaluch, C. H. Peterson, et al. 2012. Economic valuation of ecosystem services provided by oyster reefs. *Bioscience* **62**, 900–909. <https://doi.org/10.1525/bio.2012.62.10.10>
- Gray, P. C., J. T. Ridge, S. K. Poulin, A. C. Seymour, A. M. Schwantes, J. J. Swenson, et al. 2018. Integrating drone imagery into high resolution satellite remote sensing assessments of estuarine environments. *Remote Sens.* **10**, <https://doi.org/10.3390/rs10081257>
- Gray, P. C., A. B. Fleishman, D. J. Klein, M. W. Mckown, V. S. Bézy, K. J. Lohmann, et al. 2019. A convolutional neural network for detecting sea turtles in drone imagery. *Methods Ecol. Evol.* **10**, 345–355. <https://doi.org/10.1111/2041-210X.13132>
- Grizzle, R. E., J. R. Adams, and L. J. Walters. 2002. Historical changes in intertidal oyster (*Crassostrea virginica*) reefs in a Florida lagoon potentially related to boating activities. *J. Shellfish Res.* **21**, 749–756.
- Grizzle, R., K. Ward, L. Geselbracht, and A. Birch. 2018. Distribution and condition of intertidal eastern oyster (*Crassostrea virginica*) reefs in apalachicola Bay Florida based on high-resolution satellite imagery. *J. Shellfish Res.* **37**, 1027. <https://doi.org/10.2983/035.037.0514>
- He, K., G. Gkioxari, P. Dollár, and R. Girshick. 2017. Mask r-cnn. Pp. 2961–2969.
- IPCC (Intergovernmental Panel on Climate Change). 2014. *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II, and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC, Geneva, Switzerland.
- Jensen, C. C., L. S. Smart, and A. S. Deaton. 2014. *Strategic Habitat Area Nominations for Region 3: The White Oak River Basin in North Carolina*. Final Report to Marine Fisheries Commission, North Carolina Division of Marine Fisheries.
- Johnston, D. W. 2019. Unoccupied aircraft systems in marine science and conservation. *Annu. Rev. Mar. Sci.* **11**, 439–465.
- Jollineau, M. Y., and P. J. Howarth. 2008. Mapping an inland wetland complex using hyperspectral imagery. *Int. J. Remote Sens.* **29**, 3609–3631. <https://doi.org/10.1080/01431160701469099>
- Joyce, K. E., S. Duce, S. M. Leahy, J. Leon, and S. W. Maier. 2018. Principles and practice of acquiring drone-based image data in marine environments. *Mar. Freshw. Res.* <https://doi.org/10.1071/MF17380>
- Kemker, R., C. Salvaggio, and C. Kanan. 2018. Algorithms for semantic segmentation of multispectral remote sensing imagery using deep learning. *ISPRS J. Photogramm. Remote Sens.* **145**, 60–77. <https://doi.org/10.1016/j.isprsjprs.2018.04.014>
- Långkvist, M., A. Kiselev, M. Alirezaie, and A. Loutfi. 2016. Classification and segmentation of satellite orthoimagery using convolutional neural networks. *Remote Sens.* **8**, <https://doi.org/10.3390/rs8040329>
- Le Bris, A., P. Rosa, A. Lerouxel, B. Cognie, P. Gernez, P. Launeau, et al. 2016. Hyperspectral remote sensing of wild oyster reefs. *Estuar. Coast. Shelf Sci.* **172**, 1–12. <https://doi.org/10.1016/j.ecss.2016.01.039>
- Lecun, Y., Y. Bengio, and G. Hinton. 2015. Deep learning. *Nature* **521**, 436–444. <https://doi.org/10.1038/nature14539>
- Lefcheck, J. S., B. B. Hughes, A. J. Johnson, B. W. Pfirman, D. B. Rasher, A. R. Smyth, et al. 2019. Are coastal habitats important nurseries? A meta-analysis. *Conserv. Lett.* **e12645**, <https://doi.org/10.1111/conl.12645>
- Lin, T., M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, et al. 2014. Microsoft COCO: Common Objects in Context, 740–755.
- Maggiori, E., Y. Tarabalka, G. Charpiat, and P. Alliez. 2016. Fully convolutional neural networks for remote sensing image classification. *IEEE Trans. Geosci. Remote Sens.* **55**, 645–657. <https://doi.org/10.1109/IGARSS.2016.7730322>
- Manfreda, S., M. F. McCabe, P. E. Miller, R. Lucas, V. P. Madrigal, G. Mallinis, et al. 2018. On the use of unmanned aerial systems for environmental monitoring. *Remote Sens.* **10**, <https://doi.org/10.3390/rs10040641>
- Nahirnick, N. K., L. Reshitnyk, M. Campbell, M. Hessing-Lewis, M. Costa, J. Yakimishyn, et al. 2019. Mapping with confidence; delineating seagrass habitats using Unoccupied Aerial Systems (UAS). *Remote Sens. Ecol. Conserv.* **5**, 121–135. <https://doi.org/10.1002/rse2.98>
- Razavian, A. S., H. Azizpour, J. Sullivan, and S. Carlsson. 2014. CNN features off-the-shelf: an astounding baseline for recognition. Pp. 806–813.
- Rodriguez, A. B., F. J. Fodrie, J. T. Ridge, N. L. Lindquist, E. J. Theuerkauf, S. E. Coleman, et al. 2014. Oyster reefs can outpace sea-level rise. *Nat. Clim. Chang.* **4**, 493–497. <https://doi.org/10.1038/NCLIMATE2216>
- Schill, S., D. E. Porter, L. D. Coen, D. Bushek, and J. Vincent. 2006. *Development of an Automated Mapping Technique for Monitoring and Managing Shellfish Distributions: A Final Report Submitted to The NOAA / UNH Cooperative Institute for Coastal and Estuarine Environmental Technology (CICEET)*. Port Norris, NJ.
- Seymour, A. C., J. T. Ridge, A. B. Rodriguez, E. Newton, J. Dale, and D. W. Johnston. 2018. Deploying Fixed Wing Unoccupied Aerial Systems (UAS) for coastal morphology

- assessment and management. *J. Coastal Res.* **34**(3), 704–717. <https://doi.org/10.2112/jcoastres-d-17-00088.1>
- Simard, P. Y., D. Steinkraus, and J. C. Platt. 2003. Best practices for convolutional neural networks applied to visual document analysis. *Proceedings of the Seventh International Conference on Document Analysis and Recognition* **2**, 958–962.
- Spalding, M. D., A. L. Mcivor, M. W. Beck, E. W. Koch, I. Möller, D. J. Reed, et al. 2014. Coastal ecosystems: A critical element of risk reduction. *Conserv. Lett.* **7**, 293–301. <https://doi.org/10.1111/conl.12074>
- Wäldchen, J., and P. Mäder. 2018. Machine learning for image based species identification. *Methods Ecol. Evol.* **9**, 2216–2225. <https://doi.org/10.1111/2041-210x.13075>
- Weinstein, B. G. 2018. A computer vision for animal ecology. *J. Anim. Ecol.* **87**, 533–545.
- Yosinski, J., J. Clune, Y. Bengio, and H. Lipson. 2014. How transferable are features in deep neural networks? *Adv. Neural. Inf. Process. Syst.* 1–9.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Data S1. Supplemental Materials.