Environ Monit Assess (2020) 192:698
https://doi.org/10.1007/s10661-020-08653-z

Deep learning for automated analysis of fish abundance: the benefits of training across multiple habitats

Ellen M. Ditria · Michael Sievers · Sebastian Lopez-Marcano · Eric L. Jinks · Rod M. Connolly

Received: 29 April 2020 / Accepted: 30 September 2020

Abstract Environmental monitoring guides conservation and is particularly important for aquatic habitats which are heavily impacted by human activities. Underwater cameras and uncrewed devices monitor aquatic wildlife, but manual processing of footage is a significant bottleneck to rapid data processing and dissemination of results. Deep learning has emerged as a solution, but its ability to accurately detect animals across habitat types and locations is largely untested for coastal environments. Here, we produce five deep learning models using an object detection framework to detect an ecologically important fish, luderick (*Girella tricuspidata*). We trained two models on footage from single habitats (seagrass or reef) and three on footage from both habitats. All models were subjected to tests from both habitat types. Models performed well on test data from the same habitat type (object detection measure: mAP50: 91.7 and 86.9% performance for seagrass and reef, respectively) but poorly on test sets from a different habitat type (73.3 and 58.4%, respectively). The model trained on a combination of both habitats produced the highest object detection results for both tests (an average of 92.4 and 87.8%, respectively). The ability of the combination trained models to correctly estimate the ecological abundance metric, MaxN, showed similar patterns. The findings demonstrate that deep learning models extract ecologically useful information from video footage accurately and consistently and can perform across habitat types when trained on footage from the variety of habitat types.

Keywords Computer vision · Machine learning · MaxN · Monitoring · Reef · Seagrass

Electronic supplementary material The online version of this article (https://doi.org/10.1007/s10661-020-08653-z) contains supplementary material, which is available to authorized users.

E. M. Ditria (✉) · M. Sievers · S. Lopez-Marcano · E. L. Jinks · R. M. Connolly
Australian Rivers Institute – Coast & Estuaries, and School of Environment and Science, Griffith University, Gold Coast, QLD 4222, Australia
e-mail: ellen.ditria@griffithuni.edu.au

Introduction

People have been monitoring and counting wildlife for millennia, collecting invaluable data for several uses such as informing conservation, tracking population trends, and estimating abundance or biomass for fisheries stock assessments (Goldsmith 2012). As the world changes and ecosystems experience severe and sustained declines in extent and condition (Maxwell et al. 2016), monitoring wildlife has never been more important. The speed and scale at which the natural world is changing also mean that monitoring and analysing data quickly enough to be able to respond has become a global challenge. Aquatic coastal habitats are among the most severely affected by anthropogenic activities (Davidson 2014; Tulloch et al. 2020), despite being renowned for their roles in fisheries productivity, coastal protection, carbon sequestration, and biodiversity (Sievers et al. 2019; Silliman et al. 2019). The challenge of developing rapid, effective monitoring in coastal aquatic habitats is thus imperative.

Published online: 12 October 2020

The advent of cheap, high-resolution cameras has enabled large amounts of underwater data to be collected without many of the logistical issues encountered using manual methods of data collection. For example, cameras can be deployed in situ for hours to months without the need for human interaction (Podder et al. 2019). Additionally, the presence of humans and their equipment often causes animals to display avoidance behaviour and make data collection unreliable (Frid and Dill 2002). The ease with which data can now be collected, however, has exacerbated the challenge of being able to analyse data quickly, with manual analysis of photo and video footage laborious (Weinstein 2018). Scientists consequently need tools to analyse an enormous amount of data and quickly extract useful ecological information for management and conservation purposes.

Machine learning technologies have emerged as elegant solutions for automating the analysis of video and image-based datasets. Machine learning is broadly categorised as algorithms that generate predictions based on pattern detection in data (Christin et al. 2019). Traditional machine learning models such as support vector machines (SVMs) have been used to classify marine species in underwater imagery (Spampinato et al. 2010; Beijbom et al. 2012). However, these models are limited in their ability to process raw, unaltered images as the data require manual feature extraction and the associated software and domain expertise to do so (LeCun et al. 2015). Deep learning is a technology under the umbrella of machine learning that can outperform traditional machine learning algorithms when presented with underwater imagery (Villon et al. 2016). Deep learning frameworks consist of computational layers that can process raw data images and automatically extract features (LeCun et al. 2015). Additionally, deep learning models improve with higher volumes of training data relative to standard machine learning algorithms (Alom et al. 2019). Scientists have recently utilised deep learning technology (i.e. classification algorithms) to identify fish in aquatic ecosystems (Villon et al. 2018; dos Santos and Goncalves 2019; Salman et al. 2019b; Sheaves et al. 2020). More recently, object detection algorithms have been used to count individual fish of a target species to provide estimates of abundance with greater accuracy and speed than humans (Ditria et al. 2020). Although this greatly increases efficiency, repeatability, and accuracy of image-based data analysis (Weinstein 2018), training algorithms takes time and imposes high initial costs (Christin et al. 2019). Flexibility and robustness in identifying species across large spatial and temporal scales are therefore key for deep learning to be a practical method to replace manual analysis.

The use of object detection to identify and count fish in coastal environments presents a unique set of challenges. For example, many factors may affect the model's ability to detect fish, such as water turbidity, lighting variation, occlusion due to schooling fish, and changes in fish orientation (Mandal et al. 2018). Further, fish often use different habitats, whether for daily migrations among habitats to feed and shelter or ontogenetic habitat shifts (Lecchini and Galzin 2005; Igulu et al. 2014). Inter-habitat differences might mean that a deep learning model trained on one habitat will not be reliable for others. For instance, structural complexity may influence model performance, as background confusion and foreground camouflage might compromise accuracy (Salman et al. 2019a).

To maximise efficacy and accuracy of monitoring and analysis, deep learning models must be robust to changes in image backgrounds (e.g. across habitats). However, quantitative tests of performance across habitats are not available in the literature, so we do not know how a model trained on one habitat will perform on another. This is a common challenge in computer vision known as domain shift, where the data the model was trained on is not an accurate representation of real-world data (Kalogeiton et al. 2016). Generally, the performance of the deep learning model will depend on the environment, or domain, it was trained on (Kalogeiton et al. 2016). Here, we test the potential for deep learning algorithms to work effectively across habitats using luderick (*Girella tricuspidata*) as a target speices. Luderick are found in multiple habitats including seagrass meadows and rocky reefs along the temperate waters of east coast Australia and northern New Zealand (Abrantes et al. 2015). They are a recreational and commercial fisheries species and are important herbivores that control algal growth on reefs and in seagrass meadows affecting plant growth either positively through removal of epiphytic algae (Ferguson et al. 2015) or adversely though direct grazing of seagrass (Wendländer et al. 2020). By assessing the capacity of deep learning algorithms to transcend habitats across spatial scales, we provide evidence of the applicability of this technology to assist in monitoring and conservation efforts.

Materials and method

Datasets

The training dataset was collected using submerged action cameras (Haldex Sports Action Cam HD 1080p and GoPro 8 Black 1080p) deployed in two dissimilar habitats frequented by luderick, seagrass meadows and rocky reefs, in the Tweed River estuary on the border of Queensland and New South Wales, Australia (−28.169438, 153.547594). Cameras were positioned to collect footage at multiple angles and backgrounds to ensure variety in the training data. Footage was also collected at several points in time to increase variability in other environmental factors such as lighting and water turbidity. Videos were trimmed to remove footage without fish and split into 5 frames per second. Polygonal segmentation masks were manually drawn around the region of interest (ROI), here individual luderick. The algorithm extracts features automatically and begins to recognise patterns which “train” the computer to associate these with the ROI (LeCun et al. 2015). Five datasets, each consisting of ~4700 annotated luderick, were used for training. These contained seagrass footage only, reef footage only, or a combination of both habitats using a randomised subset of the videos. Three combination training datasets were created to test the stability and reproducibility of this method (Fig. 1). The videos used for the two test datasets did not appear in the training data (Fig. 1). These comprised 62 videos with approximately 1500 luderick annotations in each test used as the ground truth to quantify the model’s ability to accurately detect and count fish (Fig. 1). We incorporated footage from different days in both the test and the training set, so the model was only tested for its ability to detect fish in different habitats (domains) not throughout time.

Convolutional neural network

The object detection framework we used is an implementation of Mask R-CNN developed by Massa and Girshick (2018). Model development was conducted using a ResNet50 configuration, pre-trained on the ImageNet-1k dataset. Model training, testing, and prediction tasks were conducted on a Microsoft Azure Data Science Virtual Machine powered by an NVIDIA V100 GPU. Overfitting was minimised by using the early-stopping technique (Prechelt 1998).

Performance measurements

We tested how well the model performed at object detection and measuring fish abundance. Object detection performance was determined for each test as the mean average precision 50 value (mAP50, Everingham et al. 2010). This is the ability of the algorithm to accurately fit a segmentation mask to at least 50% of the ROI. Fish abundance performance was tested using MaxN, the maximum number of fish of the target species in any one frame, the most widely reported measure in ecological studies using video footage (Whitmarsh et al. 2017). How accurately MaxN was estimated by the model was calculated using an F1 score, the harmonic mean of precision and recall (Goutte and Gaussier 2005). True positives (correctly identified luderick), false negatives (missed luderick), and false positives (incorrect luderick identifications) are all considered when calculating precision and recall (Buckland and Gey 1994).

Results and discussion

For object detection, the seagrass and reefs models did not perform as well when trained on footage exclusively from the other habitat (Fig. 1). However, performance on tests by models trained on footage from the same habitat or from a combination of habitats was all high (> 87%) (Fig. 1, Appendix 2), indicating that the algorithm accurately fitted segmentation masks around luderick. The mAP50 test scores for the combination trained model were within 1% of the model for seagrass and within 3% for the reef habitat (Fig. 1, Appendix 2).

For estimating abundance, the overall pattern was similar; all three combination models gave high performance results (F1 range 87–92%), as did training singularly on the habitat being tested (Fig. 1). For the seagrass test, combination training was almost the same as the seagrass-only trained model (F1 91.1% vs 91.5%, respectively), whereas for the reef test, combination training was slightly lower than reef training (F1 87.1% vs 88.8%, respectively, Fig. 1). The combination training achieved the lowest number of false negatives for both tests and most closely enumerated the true number of luderick (Table 1, Appendix 1). In both tests, the model trained on the alternative habitat gave the worst performance.

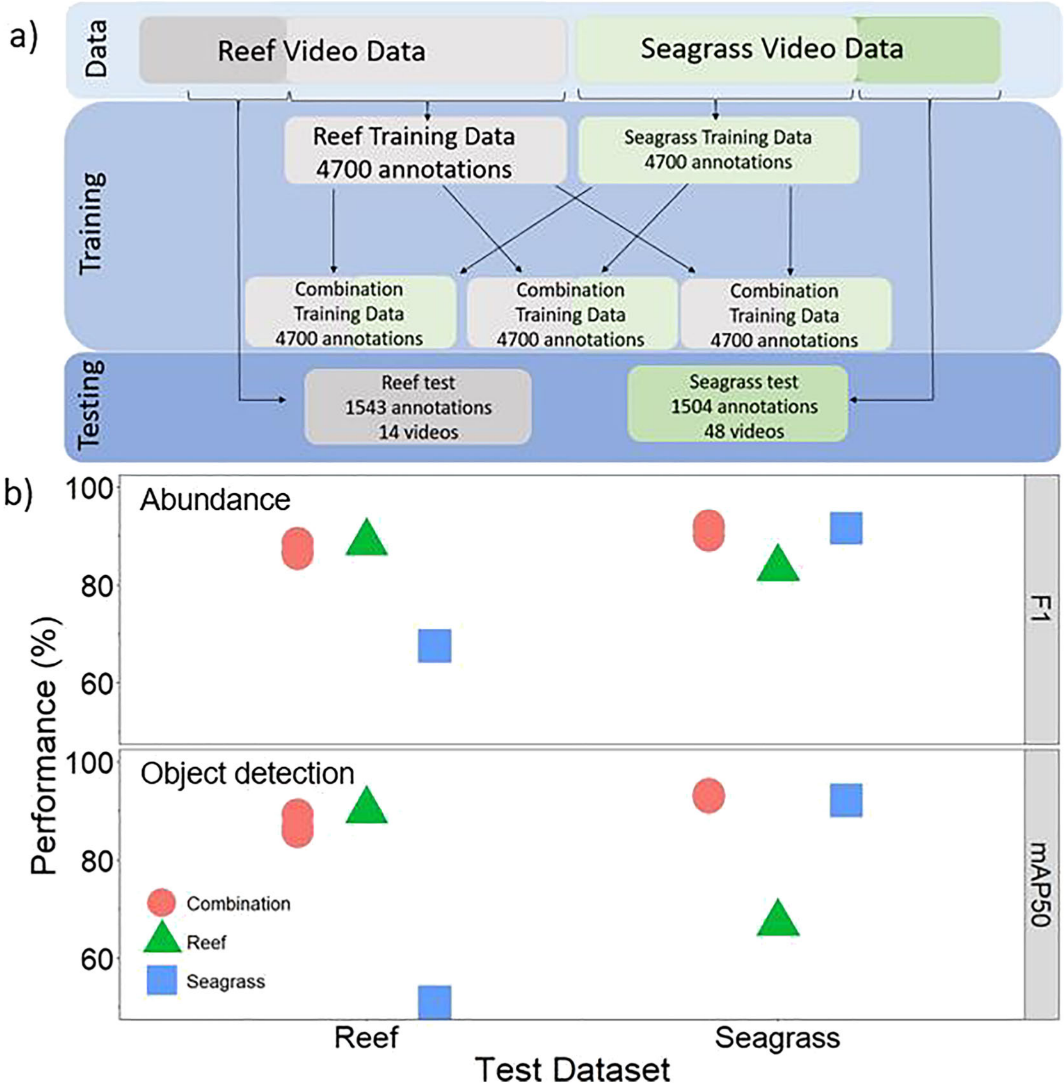

Fig. 1 **a** Trimmed videos from the original dataset from two habitats (seagrass and reef) were split into two training sets, with a random selection of the annotations from both training sets used to create the three combination training sets. Two separate test datasets were created from the remaining data from the seagrass and reef habitats. **b** Performance metrics of each model (seagrass trained, reef trained, and the three combination trained datasets) on predicting test data from a reef or seagrass habitat. Object detection performance is reported as mAP50. Abundance performance is reported as F1 score, denoting how well the model estimated MaxN per video

Deep learning models trained on a combination of habitats produced the best object detection (mAP50) and either the best or nearly as good abundance estimates (F1, Appendix 2). As expected, models also performed very well when tested on the same habitat they were trained on but were poor at object detection and estimating abundance when tested on the opposite habitat. Combination training had the lowest number of false negatives and most closely enumerated the true number of luderick. The trained reef dataset, however, had the lowest number of false positives for both tests, suggesting that models trained on this dataset were able to account for other environmental factors that could have been confused as luderick, such as other similar looking fish species. In general, the reef footage is comprised of a more complex background, with regularly changing lighting conditions on the substrate and greater fish species richness than in seagrass datasets. While the seagrass datasets contain images with comparatively low background complexity, the water was

Table 1 Summary of model performance. Ground truth is the number of fish in videos, and estimated fish is the number detected by the model (the sum of true positives and false positives)

Test	Training	Ground truth	Estimated fish	True positive	False positive	False negative
Seagrass						
	Seagrass	105	121	99	22	6
	Reef	105	81	75	6	30
	Combination 1	105	125	102	20	3
	Combination 2	105	114	96	9	9
	Combination 3	105	115	98	10	7
Reef						
	Seagrass	128	97	11	86	17
	Reef	128	137	120	17	8
	Combination 1	128	156	118	28	10
	Combination 2	128	159	119	31	9
	Combination 3	128	145	115	17	13

often more turbid, and picture clarity was thus affected (Fig. 2). Differences between habitats may explain why models trained on the singular opposite habitat performed poorly and why the seagrass-trained model performed particularly poorly when tested on reef footage. The close test results for all three combination-trained models suggest that they were robust in predicting across habitats. However, more tests from different reef and seagrass locations may be needed to determine the effectiveness of this method across spatial scales in novel environments.

To maximise the effectiveness of monitoring and the reliability of analyses, algorithms must prove robust across habitat types and often across distant locations. We have previously shown that deep learning models can have equally high performance in seagrass habitats from a different estuary than where the training footage was taken (Ditria et al. 2020). The transferability of algorithms from one habitat type and location to novel habitat types and locations strengthens this as an alternative option to manual analysis. However, this is not always the case. For example, Xu and Matzner (2018) found that a deep learning model for fish detection trained on two sites, and tested on the third, did not perform as well as those trained and tested on the same sites. This low transferability may have been due to variable water clarity making it difficult to detect fish in video footage (Xu and Matzner 2018). Further testing of fish detection algorithms across habitats and locations is required; training is improved when some variation in habitats is captured, but new training may not be required for every new habitat or location across a species distribution. An additional advantage of deep learning is that unlike traditional machine learning algorithms, deep learning algorithms are not saturated at higher volumes of data, so additional training data will generally improve the overall output performance of the existing dataset (Moniruzzaman et al. 2017; Sarwar et al. 2019; Tao et al. 2019). Collectively, this suggests that adding training from newly encountered habitats can continuously improve monitoring results.

Although environmental issues such as different backgrounds, turbidity, lighting, and colour hue are not dissimilar to those faced by humans when identifying fish from videos, deep learning algorithms can outperform humans when faced with ambiguous images (Villon et al. 2018; Ditria et al. 2020). Mask R-CNN can "learn" that the unselected confounding background pixels are not the region of interest and do not require complex pre-processing of images for background subtraction (Massa and Girshick 2018). Our method is, however, only as good as the training data it receives. Supervised deep learning models require manually labelled data to learn the input categories (here luderick). If the quality of manually annotated data is poor, for example, if other fish species are misidentified and labelled as luderick, performance may drop (Rawat and Wang 2017). Comparing model performance against humans and finding robust statistical methods to account for manual training error rate would be beneficial to further understand the potential of deep learning as a monitoring method.

Fig. 2 Examples of luderick in test footage from reef (**a**, **b**) and seagrass (**c**, **d**) habitats, highlighting the diversity of environmental complexities and picture clarity

Given our rapidly changing world, using robust and flexible deep learning algorithms to monitor and track changes in species occurrence and abundance across entire spatial distributions is important. Efficient monitoring of luderick, for example, could benefit coastal ecological science as luderick make up a significant component of the total fish biomass on temperate reefs on the east coast of Australia and are important algal grazers (Ferguson et al. 2015). Abundance data for luderick are currently limited and geographically patchy (Abrantes et al. 2015), and there is a suggestion that their numbers have declined substantially at the northern (warmer) end of their range in southeast Queensland in recent decades (Pollock 2017). Waters along the south east coast of Australia are experiencing a warming rate over three times the global average (Ridgway 2007), leading to the tropicalisation of historically temperate reefs (Hobday and Pecl 2014; Vergés et al. 2018), and southward range shifts have been documented for several fish species here (Townhill et al. 2019). A southward shift in luderick distribution will reduce their role as a key grazer at the current northern limit of their distribution. Monitoring changing range shifts has become a necessary task for management and conservation of functional habitats, and implementing deep learning solutions to analyse the large amounts of data available is promising. Deep learning algorithms, when trained across a variety of habitat types, could assist in tracking distribution shifts and changes in population sizes across a species range.

Conclusions

Deep learning is emerging as an accessible and alternative method to manage and extract information from large volumes of raw video footage. The use of a diverse training dataset consisting of different habitats and a range of environmental conditions proved to be the most robust and flexible model when analysing footage from different habitats. These models can continually be added to without any adverse effects on performance. Deep learning has the potential to offer rapid data analysis for monitoring species across locations with high efficiency and a high level of accuracy and consistency.

Acknowledgements We thank the volunteers who assisted in manual training of the deep learning algorithm, M. Turner and A. Shand.

Authors' contributions ED and RC designed the study. ED and SL conducted the fieldwork. ED and EJ developed the deep learning architecture and user interface. RC provided resources. All authors helped interpret results. ED led the writing of the manuscript, with input from all authors.

Funding RC was supported by a Discovery Project from the Australian Research Council (DP180103124). All authors were supported by the Global Wetlands Project, with support by a charitable organisation which neither seeks nor permits publicity for its efforts.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

Abrantes, K. G., Barnett, A., Baker, R., & Sheaves, M. (2015). Habitat-specific food webs and trophic interactions supporting coastal-dependent fishery species: An Australian case study. *Reviews in Fish Biology and Fisheries, 25*(2), 337–363.

Alom, M. Z., Taha, T. M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M. S., Hasan, M., van Essen, B. C., Awwal, A. A. S., & Asari, V. K. (2019). A state-of-the-art survey on deep learning theory and architectures. *Electronics, 8*(3), 292.

Beijbom, O., Edmunds, P. J., Kline, D. I., Mitchell, B. G., & Kriegman, D. (2012). Automated annotation of coral reef survey images. *IEEE Conference on Computer Vision and Pattern Recognition, 2012*, 1170–1177.

Buckland, M., & Gey, F. (1994). The relationship between recall and precision. *Journal of the American Society for Information Science, 45*(1), 12–19.

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

Davidson, N. C. (2014). How much wetland has the world lost? Long-term and recent trends in global wetland area. *Marine and Freshwater Research, 65*(10), 934–941.

Ditria, E. M., Lopez-Marcano, S., Sievers, M., Jinks, E. L., Brown, C. J., & Connolly, R. M. (2020). Automating the analysis of fish abundance using object detection: Optimizing animal ecology with deep learning. *Frontiers in Marine Science, 7*, 429.

dos Santos, A. A., & Goncalves, W. N. (2019). Improving Pantanal fish species recognition through taxonomic ranks in convolutional neural networks. *Ecological Informatics, 53*, 100977. https://doi.org/10.1016/j.ecoinf.2019.100977.

Everingham, M., Van Gool, L., Williams, C. K., Winn, J., & Zisserman, A. (2010). The pascal visual object classes (voc) challenge. *International Journal of Computer Vision, 88*(2), 303–338.

Ferguson, A., Harvey, E. S., Rees, M., & Knott, N. A. (2015). Does the abundance of girellids and kyphosids correlate with cover of the palatable green algae, Ulva spp.? A test on temperate rocky intertidal reefs. *Journal of Fish Biology, 86*(1), 375–384.

Frid, A., & Dill, L. (2002). Human-caused disturbance stimuli as a form of predation risk. *Conservation Ecology, 6*(1).

Goldsmith, F. B. (2012). *Monitoring for conservation and ecology* (Vol. 3). Springer Science & Business Media.

Goutte, C., & Gaussier, E. (2005). A probabilistic interpretation of precision, recall and F-score, with implication for evaluation. Paper presented at the European Conference on Information Retrieval.

Hobday, A. J., & Pecl, G. T. (2014). Identification of global marine hotspots: Sentinels for change and vanguards for adaptation action. *Reviews in Fish Biology and Fisheries, 24*(2), 415–425.

Igulu, M. M., Nagelkerken, I., Dorenbosch, M., Grol, M. G., Harborne, A. R., Kimirei, I. A., et al. (2014). Mangrove habitat use by juvenile reef fish: Meta-analysis reveals that tidal regime matters more than biogeographic region. *PLoS One, 9*(12), e114715.

Kalogeiton, V., Ferrari, V., & Schmid, C. (2016). Analysing domain shift factors between videos and images for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 38*(11), 2327–2334.

Lecchini, D., & Galzin, R. (2005). Spatial repartition and ontogenetic shifts in habitat use by coral reef fishes (Moorea, French Polynesia). *Marine Biology, 147*(1), 47–58.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444.

Mandal, R., Connolly, R. M., Schlacher, T. A., & Stantic, B. (2018). *Assessing fish abundance from underwater video using deep neural networks.* Paper presented at the 2018 international joint conference on neural networks (IJCNN).

Massa, F., & Girshick, R. (2018). Maskrcnn-benchmark: Fast, modular reference implementation of instance segmentation and object detection algorithms in PyTorch.

Maxwell, S. L., Fuller, R. A., Brooks, T. M., & Watson, J. E. (2016). Biodiversity: The ravages of guns, nets and bulldozers. *Nature News, 536*(7615), 143–145.

Moniruzzaman, M., Islam, S. M. S., Bennamoun, M., & Lavery, P. (2017). Deep learning on underwater marine object detection: A survey. Paper presented at the International Conference on Advanced Concepts for Intelligent Vision Systems.

Podder, T. K., Sibenac, M., & Bellingham, J. G. (2019). *Applications and challenges of AUV docking systems deployed for long-term science missions*. Monterey Bay Aquarium Research Institute.

Pollock, B. (2017). Latitudinal change in the distribution of luderick Girella tricuspidata (Pisces: Girellidae) associated with increasing coastal water temperature in eastern Australia. *Marine and Freshwater Research, 68*(6), 1187–1192.

Prechelt, L. (1998). Early stopping-but when? In *Neural Networks: Tricks of the trade* (pp. 55-69): Springer.

Rawat, W., & Wang, Z. (2017). Deep convolutional neural networks for image classification: A comprehensive review. *Neural Computation, 29*(9), 2352–2449.

Ridgway, K. (2007). Long-term trend and decadal variability of the southward penetration of the east Australian current. *Geophysical Research Letters, 34*(13).

Salman, A., Maqbool, S., Khan, A. H., Jalal, A., & Shafait, F. (2019a). Real-time fish detection in complex backgrounds using probabilistic background modelling. *Ecological Informatics, 51*, 44–51.

Salman, A., Siddiqui, S. A., Shafait, F., Mian, A., Shortis, M. R., Khurshid, K., et al. (2019b). Automatic fish detection in underwater videos by a deep neural network-based hybrid motion learning system. *ICES Journal of Marine Science.*

Sarwar, S. S., Ankit, A., & Roy, K. (2019). Incremental learning in deep convolutional neural networks using partial network sharing. *IEEE Access*.

Sheaves, M., Bradley, M., Herrera, C., Mattone, C., Lennard, C., Sheaves, J., & Konovalov, D. A. (2020). Optimizing video sampling for juvenile fish surveys: Using deep learning and evaluation of assumptions to produce critical fisheries parameters. *Fish and Fisheries.*

Sievers, M., Brown, C. J., Tulloch, V. J., Pearson, R. M., Haig, J. A., Turschwell, M. P., et al. (2019). The role of vegetated coastal wetlands for marine megafauna conservation. *Trends in Ecology & Evolution, 34*, 807–817.

Silliman, B. R., He, Q., Angelini, C., Smith, C. S., Kirwan, M. L., Daleo, P., et al. (2019). Field experiments and meta-analysis reveal wetland vegetation as a crucial element in the coastal protection paradigm. *Current Biology, 29*(11), 1800–1806. e1803.

Spampinato, C., Giordano, D., Di Salvo, R., Chen-Burger, Y.-H. J., Fisher, R. B., & Nadarajan, G. (2010). Automatic fish classification for underwater species behavior understanding. In *Proceedings of the first ACM international workshop on Analysis and retrieval of tracked events and motion in imagery streams* (pp. 45–50).

Tao, Y., Tu, Y., & Shyu, M.-L. (2019). Efficient Incremental Training for Deep Convolutional Neural Networks. In *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)* (pp. 286–291): IEEE.

Townhill, B. L., Radford, Z., Pecl, G., van Putten, I., Pinnegar, J. K., & Hyder, K. (2019). Marine recreational fishing and the implications of climate change. *Fish and Fisheries, 20*(5), 977–992.

Tulloch, V. J., Turschwell, M. P., Giffin, A. L., Halpern, B. S., Connolly, R., Griffiths, L., et al. (2020). Linking threat maps with management to guide conservation investment. *Biological Conservation, 245*, 108527.

Vergés, A., Doropoulos, C., Czarnik, R., McMahon, K., Llonch, N., & Poore, A. G. (2018). Latitudinal variation in seagrass herbivory: Global patterns and explanatory mechanisms. *Global Ecology and Biogeography, 27*(9), 1068–1079.

Villon, S., Chaumont, M., Subsol, G., Villéger, S., Claverie, T., & Mouillot, D. (2016). Coral reef fish detection and recognition in underwater videos by supervised machine learning: Comparison between Deep Learning and HOG+ SVM methods. In *International Conference on Advanced Concepts for Intelligent Vision Systems* (pp. 160–171). Springer.

Villon, S., Mouillot, D., Chaumont, M., Darling, E. S., Subsol, G., Claverie, T., & Villéger, S. (2018). A deep learning method for accurate and fast identification of coral reef fishes in underwater images. *Ecological Informatics, 48*, 238–244.

Weinstein, B. G. (2018). A computer vision for animal ecology. *Journal of Animal Ecology, 87*(3), 533–545.

Wendländer, N. S., Lange, T., Connolly, R. M., Kristensen, E., Pearson, R. M., Valdemarsen, T., & Flindt, M. R. (2020). Assessing methods for restoring seagrass (Zostera muelleri) in Australia's subtropical waters. *Marine and Freshwater Research, 71*(8), 996–1005.

Whitmarsh, S. K., Fairweather, P. G., & Huveneers, C. (2017). What is big BRUVver up to? Methods and uses of baited underwater video. *Reviews in Fish Biology and Fisheries, 27*(1), 53–73.

Xu, W., & Matzner, S. (2018). *Underwater fish detection using deep learning for water power applications.* Paper presented at the 2018 international conference on computational science and computational intelligence (CSCI),