Automated Coral Health Assessment in the Maldives: Addressing Dataset Limitations

Mohamed Rayyan Ameez¹, Houd Rizvee², Sujau Ahmed³ ¹School of Computer Science, Villa College, Male', Maldives ^{2,3}School of Computer Science, Villa College, Male', Maldives ¹<u>mohamed2.Ameez@live.uwe.ac.uk</u>, ²<u>ismail2.rizvee@live.uwe.ac.uk</u> ³ahmed36.mohamed@live.uwe.ac.uk

Abstract

Coral reefs support the Maldives' economy and environmental resilience by functioning as natural barriers that protect the country's low-lying islands from erosion and extreme weather. With increasing coral bleaching incidents due to climate change, accurate and scalable reef health monitoring has become critical; however, existing AI-based methodologies rely heavily on datasets from regions ecologically distinct from Maldivian reefs, limiting their applicability. This study aims to address this regional gap by fine-tuning YOLOv8 for coral health detection, specifically evaluating performance constraints posed by dataset quality. Experiments utilized a publicly available 16-class annotated dataset from Roboflow for training, alongside local images provided by Maldives Resilient Reefs for inference validation. Quantitative results indicated modest model performance (validation mAP50-95 = 0.253, mAP50 = 0.447; test mAP50-95 =0.225, mAP50 = 0.406), indicating significant challenges inherent in the dataset, such as overlapping,

1 INTRODUCTION

The Maldives, a nation of coral atolls, is particularly reliant on its reef ecosystems. With 2041 distinct reef structures larger than 0.01 km2, including 529 located on the rims 17 complex atolls, the Maldives possesses one of the most significant coral reef ecosystems globally (Naseer & Hatcher, 2004, as cited in Stevens and Froman, 2019). Over 80% of its land area lies less than one metre above sea level, making coral reefs essential natural breakwaters that protect against erosion and extreme weather events (Hilmi et al., 2023).

Additionally, the Maldives' economy is primarily

inconsistent labels, and annotation errors, which compromised model training efficacy. Qualitative inference on local Maldivian images demonstrated practical viability, although semantic accuracy remained limited due to training dataset deficiencies. The study identifies dataset quality as the primary obstacle for AI-driven coral reef assessments in the Maldives. Its primary contribution is to clearly outline necessary improvements for future dataset development, including the establishment of an ecologically meaningful and mutually exclusive taxonomy through expert collaboration, rigorous curation and annotation of region-specific data, and subsequent comprehensive validation. Addressing these dataset difficulties is critical to properly applying AI's potential for effective and scalable coral reef monitoring, thereby enhancing conservation management strategies in the Maldives.

Key words: Coral reef; AI; YOLOv8; Maldives; object detection; dataset.

dependent on reef-based tourism and fisheries. (Hilmi et al., 2023). Reef fishing in the Maldives has historically been small-scale and limited to local requirements, however, the rapid development of tourism, particularly the construction of resorts and guesthouses, substantially increased the demand for reef fish (Stevens and Froman, 2019). This expansion has put further strain on the country's reefs, as both tourism and fisheries continue to rise; thus, maintaining a balance between economic development and reef conservation has become increasingly crucial (Stevens and Froman, 2019).

Recurrent bleaching of coral reefs due to global warming has now become a growing concern (Hughes

et al., 2017). Degraded coral reefs are not able to withstand the increasing effects of sea level rise and global warming compared to healthy and resilient reefs, which is especially concerning as the Maldives is the lowest-lying nation in the world (Bianchi et al., 2016; Hoegh-Guldberg et al., 2007; Hughes et al., 2003; Hughes et al., 2007, as cited in Stevens and Froman, 2019). Effective reef monitoring and conservation strategies are therefore essential for biodiversity

2 BACKGROUND RESEARCH

protection, climate resilience, and sustainable resource management (Obura et al., 2019).

This project addresses a critical gap in region-specific coral monitoring by fine-tuning Yolo with localized Maldivian data, collected in collaboration with the Maldives Resilient Reef, with the aim to deliver scalable coral reef conservation tools.

#	Study / Year	Task & Dataset (Scope)	Model & Mods	Key Metric(s) Reported	Pros	Caveats
1	(Beijbom et al., 2012)	Point-annotatio n of 1500 transect frames (Pacific reefs)	Random forest on hand-crafted colour-texture	83 % point-level accuracy (across sites)	First large-scale auto-annotation; reproducible	Still ~10 min/img manual QC; not real-time
2	(Beijbom et al., 2015)	4-reef benthic survey	CoralNet semi-auto CNN	53 % manual effort saved for \leq 5 % loss vs humans	Practical human-in-the-loo p workflow	Var. in expert labels limits ceiling accuracy
3	(Mahmood et al., 2016)	AUV mosaics, 11 classes	AlexNet-style CNN	86–92 % patch accuracy (5-fold)	Showed deep nets beat SVMs	Training patches, not object-level detections
4	(Santoso et al., 2022)	Chaetodontidae fish on video	YOLOv5 vs Faster-R-CNN vs SSD	YOLOv5 tops at 89 % AP, 0.06 s/img	Real-time; light weight	Fish only; not coral classes
5	(Luz et al., 2025)	Tubastraea invasion (Brazil)	Vanilla YOLOv8n	90 % precision / 90 % recall @ 16 ms (≈62 fps)	Handles class imbalance with aug.	Misses when reef heavily overgrown

6	(Rajan & Damodaran , 2023)	New benchmark, 5 health states	MAFFN-YOLOv 5 (attention neck)	mAP@0.5 = 90.72 %, 9 ms inference; +8.6 % over YOLOv5	Fastest on Jetson NX (0.5 ms/frame)	Custom dataset—gen eralisation unclear
7	(Li et al., 2024)	Live-coral-cove r from ROV video	Deeplab-v3+ × Unet ensemble	mean-pixel-acc = 94.47 %	First video-scale LCC estimator	Needs stabilised footage; no object IDs
8	(Gómez-Rí os et al., 2019)	Texture vs structure images (RSMAS)	ResNet-50 TL	97 % species accuracy	Near-expert taxonomy	Not real-time; high-res images only
9	(Thamarai & Aruna, 2023)	RSMAS + EILAT stressed-coral set	Ensemble of two fine-tuned CNNs	90 % overall accuracy	Simple architecture	Needs two separate nets; no localisation

Table 1: Literature Review

Manual reef surveys remain the best method for ecological assessments but have inherent limitations in scale, cost, and consistency (Chowdhury et al., 2024). Traditional OpenCV-based approaches, such as color-based segmentation and texture analysis using Gabor filters, provide partial automation but often require fine-tuning for specific datasets (Stokes & Deane, 2009). Early machine learning techniques used Support Vector Machines (SVMs) to classify benthic features based on hand-crafted feature extraction (Shihavuddin et al., 2013), but these approaches struggle with variation in underwater images (Gonzalez-Rivero et al., 2020).

Deep learning has revolutionized coral reef monitoring, with Convolutional Neural Networks (CNNs) achieving accuracy for benthic image classification comparable to that of human experts (Mahmood et al., 2016). The first large-scale deep learning study on reef monitoring by Beijbom et al. (2012) demonstrated 83% accuracy in automated coral cover estimation. Follow-up work by Beijbom et al. (2015) showed that even semi-automated annotation could reduce manual effort by 50% while maintaining high accuracy. Object detection in coral imagery presents unique challenges due to dense object distributions, occlusions, and variable lighting conditions (Zhong et al., 2022). Among deep learning models, Faster R-CNN (Ren et al., 2015) and SSD (Liu et al., 2016) have been applied to underwater detection, but their inference speeds limit real-time deployment (Santoso et al., 2022). The YOLO (You Only Look Once) model series has emerged as a leading choice for marine applications due to its high-speed, single-stage detection and robustness to cluttered environments (Redmon et al., 2016).

Recent research has validated YOLOv8 for coral reef and marine species detection. Luz et al. (2025) applied YOLOv8 to detect invasive *Tubastraea* coral in Brazil, achieving 90% precision and recall while operating at 62 FPS (frames per second). Rajan and Damodaran (2023) further optimized YOLOv5 by incorporating multi-scale attention mechanisms, improving mAP (mean average precision) by 8–18% for coral detection tasks.

Beyond detection, accurate coral health classification is

essential for tracking factors such as bleaching, disease, and recovery. ResNet-50, a 50-layer residual network, has demonstrated state-of-the-art accuracy in coral health classification (Thamarai & Aruna, 2023).

Recent studies confirm ResNet-50's suitability for coral classification (Gómez-Ríos et al., 2019; Bautista-Hernández et al. 2022). Gómez-Ríos et al. (2019) achieved 97% accuracy in distinguishing coral species using a ResNet-based approach. Bautista-Hernández et al., (2022) applied CNNs to

3 EXPERIMENTATION

This section details the experimental methodology employed to fine-tune and evaluate the YOLOv8 object detection model for potential application in Maldivian coral reef assessment. The primary focus was on establishing a reproducible workflow, evaluating performance on available annotated data, and demonstrating inference capabilities on localized imagery. It is important to note that while the introduction mentions exploration of ResNet50, the empirical results presented subsequently pertain exclusively to the YOLOv8 object detection experiments.

3.1 Datasets

Two distinct datasets formed the basis of this investigation:

3.1.1 Training and Evaluation Dataset

A publicly available dataset hosted on Roboflow Universe, identified as coral-research/coral-pathology-vtwfg (version 1, CC BY 4.0 license), was utilized for model training, validation, and quantitative performance testing. Analysis of this dataset's accompanying data.yaml configuration file revealed critical structural characteristics that significantly influenced the experimental outcomes. The dataset specified 16 distinct class labels: ['Band disease'. 'Band-diseasemultiple', 'Bleached disease', 'Bleached disease Healthy Coral multiple', 'Bleached disease White Pox Disease'. 'Bleached-disease', 'Bleached-disease-White-Pox-Disea',

detect stony coral tissue loss disease, reaching 94.1% accuracy with limited training data.

Coral reefs exhibit significant regional variability, necessitating fine-tuning models with local datasets (Chowdhury et al., 2024). Existing reef AI models, such as those developed for the Great Barrier Reef or the Caribbean, may not generalize well to Maldivian coral assemblages (Hilmi et al., 2023).

'Bleached-diseasemultiple',

'Bleached-diseasemultiple-Health', 'Dead Coral', 'Healthy Coral', 'Healthy-Coral', 'White Pox Disease', 'White-Pox-Disea', 'bleached corals', 'bleachedcorals']. This taxonomy exhibits considerable semantic overlap (e.g., Healthy Coral vs. Healthy-Coral), inconsistent naming conventions (e.g., White Pox Disease vs. White-Pox-Disea), and includes compound labels attempting to represent multiple simultaneous conditions (e.g., Bleached-disease-White-Pox-Disea). Such a structure deviates from standard object detection practices requiring distinct, mutually exclusive classes and poses significant challenges for model training. Furthermore, during data loading, approximately 11-15% of annotated items across the train, validation, and test splits were automatically excluded due to technical annotation errors (e.g., out-of-bounds bounding box coordinates), reducing the effective dataset size and introducing noise.

3.1.2 Local Inference Demonstration Dataset

An independent set of 83 underwater images acquired at the Bodufinolhu reef site, Maldives, was provided by the Maldives Resilient Reefs project. This dataset served solely as a target for demonstrating the inference capabilities (Section 3.4) of the trained model on local imagery and was not used during model training or quantitative evaluation.

3.2 Model Architecture and Training Protocol

The YOLOv8m (medium) architecture, pre-trained on the COCO dataset, was selected for fine-tuning, leveraging its established balance of performance and computational efficiency in related marine applications (Luz et al., 2025; Santoso et al., 2022). All experiments were conducted using PyTorch v2.7.0 (CUDA 12.6 backend) and Ultralytics YOLO v8.3.118 on an NVIDIA GeForce RTX 4090 GPU. Reproducibility was maintained through a fixed random seed (42) and deterministic CUDA settings.

A limited hyperparameter sweep was performed using the Roboflow dataset to evaluate the influence of input image resolution and batch size. Models were trained for 40 epochs under three configurations: (1) 640x640 resolution, Batch Size 8; (2) 768x768 resolution, Batch Size 8; (3) 768x768 resolution, Batch Size 16. Standard YOLOv8 training defaults, including data augmentation and AdamW optimizer, were employed. The primary evaluation metric tracked during training was mean Average Precision averaged over IoU thresholds from 0.50 to 0.95 (mAP50-95) on the validation split. The model checkpoint (best.pt) corresponding to the configuration yielding the highest validation mAP50-95 was selected for subsequent analysis.

3.3 Quantitative Evaluation

The performance of the selected fine-tuned YOLOv8m model was quantitatively assessed using the validation and test splits of the 16-class Roboflow dataset. Standard object detection metrics, mAP50 and mAP50-95, were computed.

The hyperparameter sweep identified the configuration using 768x768 resolution and a batch size of 8 (Run 2) as producing the highest validation mAP50-95, although overall performance was modest across all runs (Table 2).

Run Index	Img Size	Bat ch	Epo chs	Val mAP50-9 5	Val mAP50
2	768	8	40	0.253	0.447
1	640	8	40	0.241	0.439
3	768	16	40	0.219	0.412

Table 2: Hyperparameter Sweep Results (Validation SetPerformance on 16-Class Roboflow Data)

The promoted model from Run 2 yielded the following performance on the dataset splits:

- Validation Set: mAP50-95 = 0.253; mAP50 = 0.447
- Test Set: mAP50-95 = 0.225; mAP50 = 0.406

These mAP scores, particularly the mAP50-95 values, are substantially lower than typically desired for reliable ecological monitoring applications. This outcome is interpreted as being directly constrained by the fundamental structural flaws and annotation errors within the Roboflow dataset used for training. The significant difference between mAP50 and mAP50-95 suggests the model acquired a limited ability for coarse localization but struggled with precise boundary prediction and, critically, with discriminating between the numerous inconsistent and overlapping class definitions. Therefore, these metrics primarily quantify performance on the ill-defined task presented by the dataset, rather than representing the intrinsic potential of YOLOv8 for accurately assessing coral health states defined by a clean, curated dataset.

3.4 Inference Pipeline Demonstration

To verify the technical workflow's applicability to local data, the promoted model (trained on the 16-class Roboflow dataset) was utilized to perform inference on the 83 images from the Bodufinolhu site (Maldives Resilient Reefs). The pipeline successfully processed 82 images, generating annotated outputs. This confirmed the operational feasibility of deploying the trained model on previously unseen, geographically specific imagery. However, the ecological validity and semantic interpretation of the detections produced (labelled according to the 16 flawed classes) are inherently limited due to the training data's deficiencies and require qualitative assessment by domain experts.

4 CONCLUSIONS

This study successfully implemented and validated a reproducible computational pipeline for fine-tuning and deploying the YOLOv8 object detection model in the context of coral reef image analysis. The workflow, including automated hyperparameter exploration and inference on local Maldivian data, functioned reliably from a technical perspective.

The most critical finding, however, pertains to the determinative role of dataset quality. The quantitative performance evaluation was severely hampered by fundamental issues within the publicly sourced training dataset, specifically: (1) a highly inconsistent and overlapping 16-class taxonomy unsuitable for robust object detection, including redundant and compound labels, and (2) a significant rate of technical annotation bounding box errors (incorrect coordinates). Consequently, the low achieved mAP scores (Test mAP50-95 = 0.225) primarily reflect these data limitations rather than the inherent capabilities of the YOLOv8 architecture.

This pilot investigation highlights that the most immediate and crucial prerequisite for developing effective AI-based coral reef monitoring tools for the Maldives is the **establishment of a high-quality**, **curated dataset**. Future efforts must prioritize:

- Taxonomy Definition: Collaborating with local marine ecology experts (such as those at Maldives Resilient Reefs) to define a clear, consistent, ecologically meaningful, and mutually exclusive set of target classes for coral health and benthic composition relevant to Maldivian reefs.
- 2. **Data Curation:** Creating a new, large-scale annotated dataset using local imagery based on the refined taxonomy, or undertaking a rigorous cleaning and re-annotation effort of existing resources to conform to these new standards. This includes rectifying bounding box errors.
- 3. **Model Retraining and Validation:** Re-training YOLOv8 (or other suitable architectures) on the curated dataset and performing comprehensive validation across diverse Maldivian sites and conditions to establish reliable performance benchmarks.

In summary, while demonstrating a viable technical pipeline, this work underscores that meticulous dataset development is the foundational bottleneck that must be addressed to successfully leverage AI for robust and scalable coral reef conservation efforts in the Maldives.

REFERENCES

Beijbom, O., Edmunds, P.J., Kline, D.I., Mitchell, B.G., & Kriegman, D.J. (2012). Automated annotation of coral reef survey images. *Proc. IEEE CVPR 2012*, 1170–1177. DOI: 10.1109/CVPR.2012.6247798

Beijbom, O., Edmunds, P.J., Roelfsema, C., Smith, J., Kline, D.I., et al. (2015). Towards Automated Annotation of Benthic Survey Images: Variability of Human Experts and Operational Modes of Automation. *PLoS ONE*, *10*(7): e0130312. DOI: 10.1371/journal.pone.0130312

Chowdhury, A., Jahan, M., Kaisar, S., Khoda, M.E., Rajin, S.M.A.K., & Naha, R. (2024). Coral Reef Surveillance with Machine Learning: A Review of Datasets, Techniques, and Challenges. *Electronics*, *13*(24): 5027. DOI: 10.3390/electronics13245027

Gonzalez-Rivero, M., Beijbom, O., Rodriguez-Ramirez, A., Bryant, D.E., Ganase, A., et al. (2020). Monitoring of coral reefs using artificial intelligence: A feasible and cost-effective approach. *Remote Sensing*, *12*(3): 489. DOI: 10.3390/rs12030489

Hilmi, N., Basu, R., Crisóstomo, M., Lebleu, L., Claudet, J., & Seveso, D. (2023). The pressures and opportunities for coral reef preservation and restoration in the Maldives. *Front. Environ. Econ.* 2: 1110214. DOI: 10.3389/frevc.2023.1110214

Hughes, T.P., Kerry, J.T., & Simpson, T. (2018). Large-scale bleaching of corals on the Great Barrier Reef. *Ecology*, *99*(2): 2092. DOI: 10.1002/ecy.2092

Hughes, T.P. et al. (2017). Global warming and recurrent mass bleaching of corals. *Nature*, *543*: 373–377. DOI: 10.1038/nature21707

Kohler, K.E. & Gill, S.M. (2006). Coral Point Count with Excel extensions (CPCe): A Visual Basic program for determination of coral and substrate cover in transect images. *Proc. 10th Int. Coral Reef Symp.*, 13: 593–598. *(software)*

Li, Z., Zhao, S., Lu, Y., Song, C., Huang, R., & Yu, K. (2024). Deep Learning-Based Automatic Estimation of Live Coral Cover from Underwater Video for Coral Reef Health Monitoring. *Journal of Marine Science and Engineering*, *12*(11): 1980. DOI: 10.3390/jmse12111980

Luz, A.C.N., Barroso, V.R., Batista, D., Lessa, A.A., Coutinho, R., & Xavier, F.C. (2025). Early detection of marine bioinvasion by sun corals using YOLOv8. *Intell. Marine Tech. and Systems, 3*(1): 2. DOI: 10.1007/s44295-024-00052-1

Mahmood, A., Bennamoun, M., Sohel, F., Boussaid, F., Hovey, R., Kendrick, G., & Fisher, R. (2016). Automated annotation of coral reef survey images using deep learning. *Proc. MTS/IEEE OCEANS 2016*, 1–7. DOI: 10.1109/OCEANS.2016.7761417

Naseer, A. & Hatcher, B.G. (2004). Inventory of the Maldives' coral reefs using morphometrics generated from

Landsat ETM+ imagery. *Atoll Research Bulletin*, 505: 1–25. DOI: 10.5479/si.00775630.505.1

Obura, D.O. et al. (2019) Coral Reef Monitoring, Reef Assessment Technologies, and Ecosystem-Based Management. *Frontiers in Marine Science* [online]. 6. DOI: 10.3389/fmars.2019.00580

Rajan, S.K.S. & Damodaran, N. (2023). MAFFN_YOLOv5: Multi-Scale Attention Feature Fusion Network on the YOLOv5 Model for the Health Detection of Coral Reefs. *Analytics*, 2(1): 77–104. DOI: 10.3390/analytics2010006

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. *Proc. IEEE CVPR 2016*, 779–788. DOI: 10.1109/CVPR.2016.91

Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *NeurIPS* 28, 91–99. arXiv:1506.01497

Santoso, S.A., Jaya, I., & Priandana, K. (2022). Optimizing Coral Fish Detection: Faster R-CNN, SSD MobileNet, YOLOv5 Comparison. *Indones. J. Comput. Cybern. Syst. 18*(2): 1–5. DOI: 10.22146/ijccs.95011

Stevens, G.M.W. and Froman, N. (2019) The Maldives Archipelago. *World Seas: an Environmental*

Evaluation [online]. 2, pp. 211-236.

Stokes, M.D. & Deane, G.B. (2009). Automated classification of underwater multispectral imagery for coral reef monitoring. *Proc. MTS/IEEE OCEANS 2009*, 1–7. DOI: 10.23919/OCEANS.2009.5422138

Thamarai, M. & Aruna, S.P. (2023). Stressed Coral Reef Identification Using Deep Learning CNN Techniques. *J. Electronic & Information Systems*, 5(2): Art.5808. DOI: 10.30564/jeis.v5i2.5808

Tran, S.V. (2020). Using Color Thresholding and Contouring to Understand Coral Reef Biodiversity. *Master's Projects.* 915, San Jose State Univ. DOI: 10.31979/etd.tdbq-k4u6

Appendix A

Dataset ,code and code outputs

https://drive.google.com/drive/folders/1MIVmxnByW8Xw5_i 05Fd7FleFi40q5vhS?usp=drive_link