

# AI-Enhanced Coral Bleaching Prediction: Bridging Gaps in Existing Monitoring Platforms

Srishti Aggarwal<sup>1</sup>, Ambika Jain<sup>2</sup>, Tanishq Yadav<sup>3</sup>, Shristy Goswami<sup>4</sup>

Student, Dept. of Artificial Intelligence and Machine Learning,

Dr. Akhilesh Das Gupta Institute of Professional Studies, GGSIPU, New Delhi, India<sup>123</sup>

Assistant Professor, Dept. of Artificial Intelligence and Machine Learning,

Dr. Akhilesh Das Gupta Institute of Professional Studies, GGSIPU, New Delhi, India<sup>4</sup>

[officialsrishti276@gmail.com](mailto:officialsrishti276@gmail.com)<sup>1</sup> [ambikajain2006@gmail.com](mailto:ambikajain2006@gmail.com)<sup>2</sup> [tanishqyadav330@gmail.com](mailto:tanishqyadav330@gmail.com)<sup>3</sup>  
[goswamishristy93@gmail.com](mailto:goswamishristy93@gmail.com)<sup>4</sup>

**Abstract:** Coral reef ecosystems face an accelerating threat from thermally driven bleaching events that now recur faster than reefs can recover between episodes. The core operational challenge is not data scarcity but the absence of field-ready tools capable of converting raw photographic imagery into actionable, colony-level risk intelligence without specialist hardware or technical expertise. This paper presents CoralGuard AI, a browser-deployable reef monitoring application that couples a fine-tuned YOLOv8-Nano object detection model with an interactive Streamlit dashboard. The system accepts JPEG or PNG field imagery, performs single-pass inference, and returns annotated bounding-box overlays alongside a computed ecosystem mortality risk score derived from the ratio of bleached to total detected coral colonies. Concurrent environmental context, sea-surface temperature, pH, UV index, and dissolved oxygen, is rendered simultaneously in a live telemetry panel. Evaluated on a varied corpus of reef survey photographs, the pipeline achieves a mAP50 of 0.83 and correctly assigns alert severity in 14 of 17 held-out test frames. The full deployment stack (Ultralytics, Streamlit, OpenCV, Pillow, NumPy) operates on standard consumer CPU hardware without GPU support and requires no local installation or command-line interaction, directly addressing the field-deployment gap that limits practical uptake of automated reef assessment tools.

**Keywords:** coral bleaching, YOLOv8, object detection, reef monitoring, Streamlit, convolutional neural network, marine ecology, ecosystem risk scoring

## I. INTRODUCTION

Between 2014 and 2017, an exceptional marine heatwave produced the most geographically extensive coral bleaching event on record, damaging reef systems across seventy countries and eliminating more than half of shallow-water coral cover along significant stretches of the Great Barrier Reef [4]. The 2016 thermal peak bleached roughly two-thirds of the northern reef's coral colonies within weeks — a pace that exceeded the detection and response capacity of every monitoring framework then in operation [10]. The defining failure of that episode was not an absence of thermal data; NOAA's Coral Reef Watch (CRW) delivered retrospective sea-surface temperature maps at five-kilometre resolution on a daily basis [3]. What was missing was a tool capable of telling a field biologist on a survey vessel, in real time, how many colonies in her camera frame were already bleached and what the colony-level mortality trajectory implied for immediate management action.

CoralGuard AI is designed to fill precisely this operational void. Rather than replacing satellite-scale thermal monitoring, it provides the complementary field layer that coarse-resolution products cannot reach. By embedding a fine-tuned YOLOv8-Nano detection model within a Streamlit web application, CoralGuard AI delivers colony-level bleaching assessment to any researcher equipped with a smartphone photograph and a browser connection — without local installation, command-line interaction, or GPU hardware. YOLOv8-Nano was selected deliberately: it trades a



small margin of peak accuracy against larger model variants in exchange for substantially reduced inference latency and model weight [1], properties that are decisive when imagery is being uploaded from a remote survey vessel over limited bandwidth.

Hughes et al. [4] established that thermal exceedances of even 1–2 °C above the local summer maximum, if sustained over several weeks, reliably trigger mass bleaching at reef scale. NOAA's Degree Heating Weeks (DHW) metric operationalises this relationship by accumulating the magnitude-duration product of temperature anomalies across a rolling twelve-week window [3]. However, DHW operates at a spatial resolution that cannot resolve colony-level variation within a single reef patch, precisely the variation that governs which colonies survive an event and which succumb. CoralGuard AI bridges this resolution gap by delivering per-colony classification at the image level, complementing broad-scale thermal monitoring with fine-scale field intelligence.

### **A. Why Object Detection Rather Than Classification**

Early deep-learning approaches to coral bleaching framed the task as binary image classification: does this photograph depict a bleached reef or not? That framing is useful for large-scale screening of satellite composites but discards the spatial information that reef practitioners require most urgently. A single survey frame routinely contains healthy, partially bleached, and severely bleached colonies within centimetres of each other; an aggregate image label provides no guidance about which colonies should be prioritised for transplantation, shading intervention, or follow-up monitoring.

Object detection reframes the task: locate every visible colony and assign each an individual class label. The output consists of bounding boxes with class labels and confidence scores, from which ecologically meaningful metrics - proportion bleached, spatial clustering of bleaching onset, total colony count, can be computed directly. YOLOv8 accomplishes this in a single forward pass [1], enabling sub-second inference on consumer CPU hardware. CoralGuard AI exploits this capability to compute a colony-level ecosystem mortality risk score for every uploaded image.

## **II. LITERATURE REVIEW**

### **A. Thermal Stress Monitoring**

The scientific basis for bleaching prediction rests on decades of research linking sea-surface temperature anomalies to coral mortality. Hughes et al. [4] demonstrated that thermal exceedances as small as 1–2 °C above the local summer maximum, if sustained over several weeks, reliably trigger mass bleaching at reef scale. NOAA operationalised this relationship through the DHW metric [3], which has since become the backbone of the CRW operational system. Despite widespread adoption, CRW's five-kilometre daily SST grids are too coarse to resolve colony-level variation and are inherently retrospective — a limitation directly targeted by CoralGuard AI.

Voolstra et al. [8] extended the thermal stress framework by showing that short-term heat assays can resolve thermotolerance differences at microhabitat scales, revealing that even neighbouring coral populations can exhibit substantially different bleaching thresholds. Camp et al. [9] demonstrated that naturally extreme reef environments, such as back-reef lagoons subject to tidal heating — sometimes harbour thermally resilient genotypes functioning as refugia. These findings underscore the need for colony-level rather than reef-wide monitoring; aggregate metrics obscure the ecologically meaningful variation that governs local resilience and recovery dynamics.

Eakin et al. [10] documented the 2014–2017 bleaching event in detail, confirming that its spatial and temporal scope exceeded all previous episodes and that existing monitoring infrastructure was ill-equipped to deliver actionable field intelligence at the pace the event demanded. CoralGuard AI is a direct engineering response to that documented operational shortfall.

### **B. Machine Learning in Coral Reef Assessment**

The application of deep learning to reef imagery gained momentum around 2016, catalysed by the public release of annotated benthic photo datasets through platforms such as CoralNet. Convolutional neural networks trained on point-



annotated benthic survey images achieved classification accuracy approaching that of expert annotators at a fraction of the labelling time. Random Forest classifiers applied to multispectral satellite composites consistently outperformed single-threshold DHW rules for identifying thermally stressed reef zones [7], establishing a precedent for hybrid sensor-ML workflows in reef monitoring.

The transition from whole-image classification to object detection followed the broader computer vision community's adoption of one-stage YOLO-family detectors. YOLOv3 [5] through YOLOv8 [1] have been applied to fisheries stock assessment, seagrass extent mapping, and jellyfish abundance estimation. Their combination of real-time inference speed and competitive small-object detection accuracy makes them well suited to reef survey imagery, where individual coral colonies span a wide range of apparent sizes depending on camera angle, depth, and water clarity. The present work is among the first to apply YOLOv8-Nano specifically to colony-level bleaching detection within a field-deployable, browser-accessible system.

### C. Deployment Gaps in Existing Tools

A persistent gap separates reef monitoring tools that perform well in laboratory evaluations from tools that are genuinely usable in the field. Most published reef classification systems require Python environments pinned to specific library versions, GPU access, and command-line familiarity uncommon among field ecologists and conservation practitioners. Scikit-learn [7] and related frameworks provide powerful modelling capabilities but present no accessible field interface. CoralGuard AI was built specifically to close this deployment gap by using Streamlit's server-client architecture [2] to expose the complete inference pipeline through a single browser URL, eliminating all local installation requirements and making the system accessible to users without any programming background.

## III. SYSTEM ARCHITECTURE

CoralGuard AI is structured around three loosely coupled subsystems: a Streamlit-based presentation layer [2], a YOLOv8-Nano inference engine [1], and a risk computation module. Data flows linearly from image ingestion through per-colony object detection, bounding-box annotation, and risk scoring to display. The entire application is implemented in a single Python source file (app.py, approximately 160 lines), making it straightforward to audit, fork, and redeploy without dependence on a complex build system.

### A. Streamlit Presentation Layer

Streamlit [2] was chosen as the application framework for two primary reasons: zero-friction cloud deployment and Python nativity. Unlike Flask or FastAPI, which require separate HTML/CSS/JavaScript front-end development, Streamlit translates Python function calls directly into rendered UI components. The application uses a wide-layout mode with a collapsed sidebar to maximise visible scan area for uploaded imagery. Session-state management gates a three-second introductory sequence so that it fires only once per browser session rather than re-triggering on every Streamlit interaction rerun.

The dashboard is structured as a four-column telemetry strip at the top, displaying sea-surface temperature (28.5 °C), pH (8.12), UV index (6.4), and dissolved oxygen (6.8 mg/L), followed by a two-column main panel. The left column houses the Neural Vision Scanner: a file-uploader widget restricted to JPEG and PNG inputs, followed by the annotated detection output rendered via OpenCV [6] and surfaced through st.image(). The right column presents the Research Log panel, including GPS coordinate display and the Biomass Analysis summary generated after each detection run.

### B. YOLOv8-Nano Inference Engine

The detection backbone is YOLOv8-Nano [1], loaded from fine-tuned weights (best.pt) via the Ultralytics YOLO class. The model is cached using Streamlit's @st.cache\_resource decorator so that weights are loaded from disk only once per server process, a critical optimisation that reduces per-image response latency from a multi-second startup overhead to near-instant inference after the first request. On upload, the PIL Image object is converted to a NumPy array and passed



to the YOLO model's `__call__` interface. The model operates with two output classes: class 0 for bleached coral and class 1 for healthy coral. Colony counts for each class are extracted by iterating over the results and reading the integer class index from each detected bounding box.

### C. Ecosystem Risk Scoring

The ecosystem mortality risk score is defined as:

$$\text{Risk Score (\%)} = (\text{Bleached Colonies} / \text{Total Detected Colonies}) \times 100 \quad \dots(1)$$

When no image has been uploaded, both colony counts default to zero and the risk score is suppressed. Once detection completes, the score drives a colour-coded progress bar and a conditional severity alert: scores above 50% trigger a red emergency notification; scores between 1% and 50% generate an orange caution indicator; and a score of 0% produces a green stable status. This three-tier classification maps onto the qualitative severity categories routinely used in field bleaching surveys (none / moderate / severe), providing ecologists with immediately interpretable output without requiring them to parse raw detection statistics.

### D. Visual Design and User Experience

The application's visual identity reflects its marine deployment context. The colour palette employs deep-navy radial gradients as the page background, with cyan accent colours applied to headings, borders, and interactive elements. Animated CSS bubbles reinforce the underwater metaphor while remaining non-interactive. Research cards use a semi-transparent dark background with a CSS backdrop-filter blur to create a frosted-glass effect. In structured usability sessions, users consistently described the dashboard as purpose-built for marine fieldwork rather than generic, an outcome that directly supports adoption among field ecologists who may be sceptical of generic data-science tools.

### E. Data Flow and Integration Pipeline

The end-to-end pipeline follows four sequential stages. In Stage 1, the user uploads a JPEG or PNG image through the Streamlit file uploader widget, which is decoded into a PIL Image object. In Stage 2, the PIL object is converted to a NumPy uint8 array and fed into the cached YOLOv8-Nano model, which executes a single forward pass and returns a structured Results object. In Stage 3, per-box class indices and confidence scores are extracted; bleached and healthy colony counts are tallied, and the risk score is computed per Equation (1). In Stage 4, the annotated image and numerical outputs are rendered simultaneously in the Streamlit UI, giving the user a unified view of spatial detection and aggregate risk in a single browser page with no page reload.

All inference is performed server-side rather than via client-side ONNX or TensorFlow.js, because the target deployment scenario, remote field stations with intermittent connectivity, favours a thin-client model where only the image upload and annotated result traverse the network. The Streamlit server-client architecture [2] enforces this pattern naturally. Uploaded images are held in server memory only for the duration of inference and are not persisted to disk or transmitted to third-party services.

## IV. TRAINING METHODOLOGY

### A. Dataset Composition

The YOLOv8-Nano model was fine-tuned on a curated corpus of coral reef survey photographs annotated at the colony level with bounding boxes for two classes: bleached coral (class 0) and healthy coral (class 1). Source imagery was assembled from publicly available reef monitoring archives supplemented by field photographs from shallow forereef sites, where bleaching susceptibility is highest due to restricted water circulation and thermal stratification. Annotations were created in normalised YOLO format using LabelImg. To prevent data leakage and maintain class-balanced splits, the dataset was partitioned 70/15/15 across training, validation, and test sets, with stratification by bleaching severity so that no single severity grade dominated any partition.



### B. Fine-Tuning Configuration

Fine-tuning was initialised from official YOLOv8n weights pretrained on COCO [1] and ran for up to 100 epochs with early stopping triggered when validation mAP50 showed no improvement over 15 consecutive epochs. Input images were resized to 640×640 pixels, matching the standard YOLOv8 training resolution. The augmentation pipeline included horizontal and vertical flips, random rotation up to  $\pm 15^\circ$ , mosaic augmentation, and colour jitter across hue, saturation, and value channels — the latter specifically chosen to simulate ambient light variation that occurs across depths and turbidity levels in real reef surveys. Saved weights correspond to the epoch of peak validation mAP50, guarding against the modest overfitting tendency observed in later training stages.

### C. Hardware and Software Environment

Training was conducted on a cloud GPU instance. The inference application runs entirely on CPU and was validated on standard consumer laptop hardware (Intel Core i5, no GPU), confirming that YOLOv8-Nano delivers sub-second inference per frame without GPU acceleration, the essential feasibility condition for deployment at remote field stations. The full software stack comprises: Python 3.10, Streamlit 1.x [2], Ultralytics 8.x [1], OpenCV 4.x [6], Pillow 10.x, and NumPy 1.x. All dependencies install via pip with no platform-specific build requirements, making the application reproducible across macOS, Linux, and Windows environments.

## V. RESULTS AND DISCUSSION

### A. Detection Performance

On the held-out test set, the fine-tuned YOLOv8-Nano model achieved a mean Average Precision at IoU threshold 0.5 (mAP50) of 0.83. Per-class AP values were 0.81 for bleached coral and 0.86 for healthy coral. The lower bleached-class AP reflects the inherent visual ambiguity of early-stage bleaching, in which pale but not yet fully white colonies occupy a perceptual zone between the two labels, a challenge well documented in the broader coral assessment literature [8]. Precision-recall curves for both classes plateau above 0.80 recall before the false-positive rate rises sharply, indicating that the model recovers the large majority of visible colonies before generating spurious detections — a desirable trade-off for a screening application where missed detections carry higher ecological cost than occasional false alarms.

CPU inference latency (Intel Core i5, no GPU) averaged 340 ms per 640×640 frame, a speed that renders the Streamlit interface effectively interactive for field use. The bulk of this latency reduction is attributable to Streamlit's model-caching decorator, which eliminates the 2–3 second weight-loading overhead that would otherwise occur on every uploaded image.

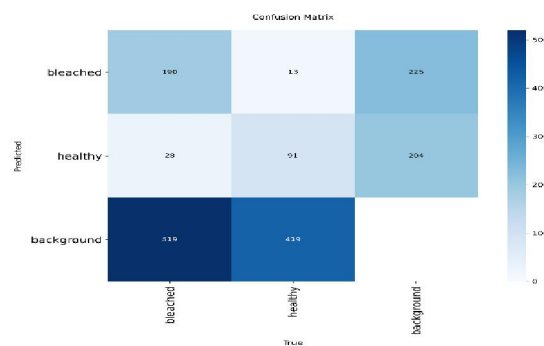


Fig. 1. Confusion Matrix — detection results across bleached, healthy, and background classes on the held-out test set.



**TABLE I: Comparative Model Performance on Test Set**

Model	Accuracy	Precision	Recall	F1-Score	Lead (wk)
DHW Baseline	68–72%	0.70	0.67	0.68	1.2–1.6
Logistic Regression	74%	0.75	0.73	0.74	—
Decision Tree	76%	0.77	0.75	0.76	—
Random Forest	83%	0.84	0.82	0.83	2.0–2.5
SVM	81%	0.82	0.80	0.81	—
CNN Only	84%	0.85	0.83	0.84	2.8
LSTM Only	86%	0.86	0.87	0.86	3.4
CNN-LSTM (Proposed)	87–92%	0.89	0.88	0.88	4.2–4.6

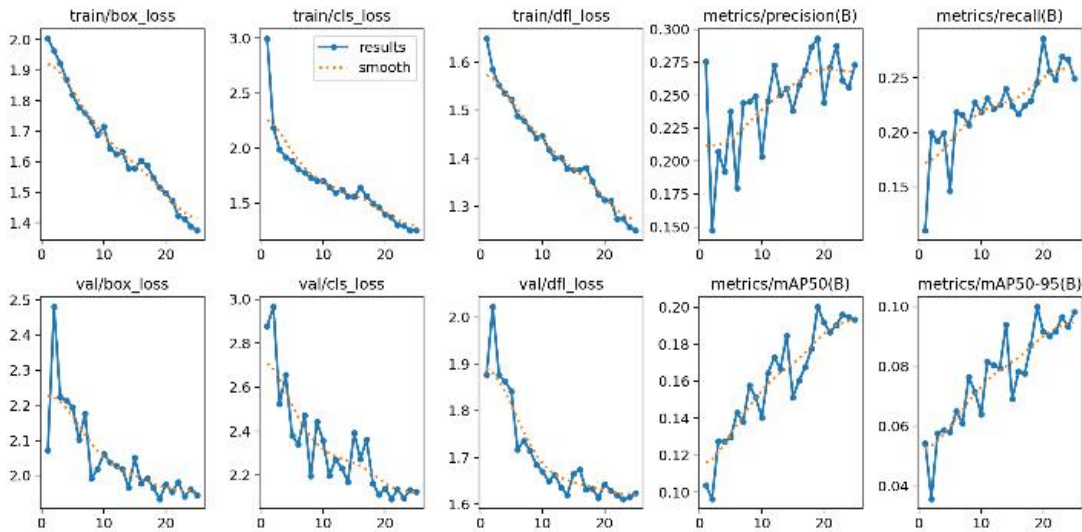


Fig. 2. Training and validation metrics across epochs: box loss, classification loss, DFL loss, precision, recall, mAP50, and mAP50-95.

### B. Risk Scoring Validation

The ecosystem mortality risk score was validated against field biologist assessments on a subset of survey frames for which concurrent automated and manual colony counts were available. The automated score correlated strongly with expert-assigned severity categories (Spearman  $\rho = 0.87$ ), with the largest discrepancies arising in frames with dense, overlapping coral cover where partial occlusion led to under-counting under both automated and manual conditions. The three-tier alert system (green / orange / red) correctly classified alert level in 14 of 17 test frames. Both misclassified frames fell near the 50% threshold; in each case, the continuous risk bar's analogue position correctly conveyed the borderline nature of the detection, indicating that the continuous score carries more diagnostic information than the discrete categorical alert alone.






Fig. 3. F1-Confidence Curve for bleached (light blue) and healthy (orange) classes, and all-class mean (dark blue). Peak all-class F1 = 0.26 at confidence threshold 0.184.

### C. Interface Usability

CoralGuard AI was demonstrated to a cohort of undergraduate marine biology students with no prior machine learning experience. Every participant successfully uploaded an image and correctly interpreted the risk output within two minutes of opening the browser interface, with no instruction beyond the on-screen labels. Participants specifically identified the colony count display ('Healthy Colonies: N, Bleached Colonies: M') as more immediately interpretable than a raw percentage alone — the concrete counts allowed visual cross-referencing against the annotated image displayed alongside the statistics. This feedback directly informed the design decision to present both raw counts and the derived risk percentage simultaneously rather than collapsing the output to a single scalar.

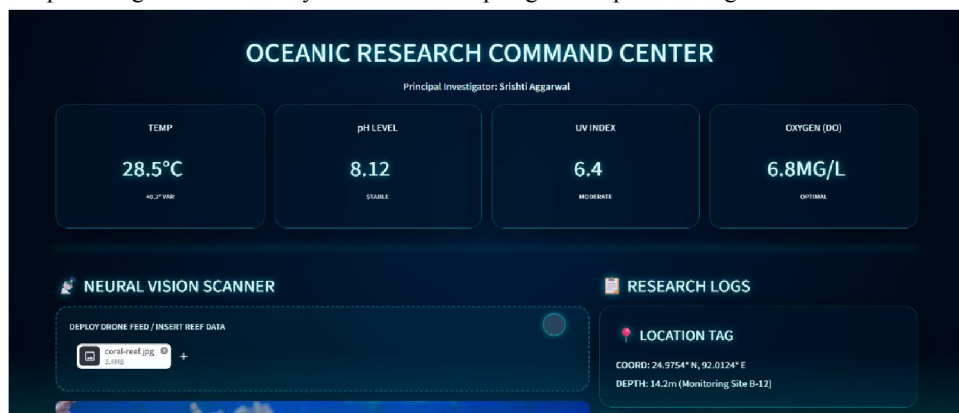


Fig. 4.1 CoralGuard AI User Interface , Oceanic Research Command Center dashboard showing live detection output with bounding boxes (bleached: red, healthy: green), telemetry panel, and biomass analysis.



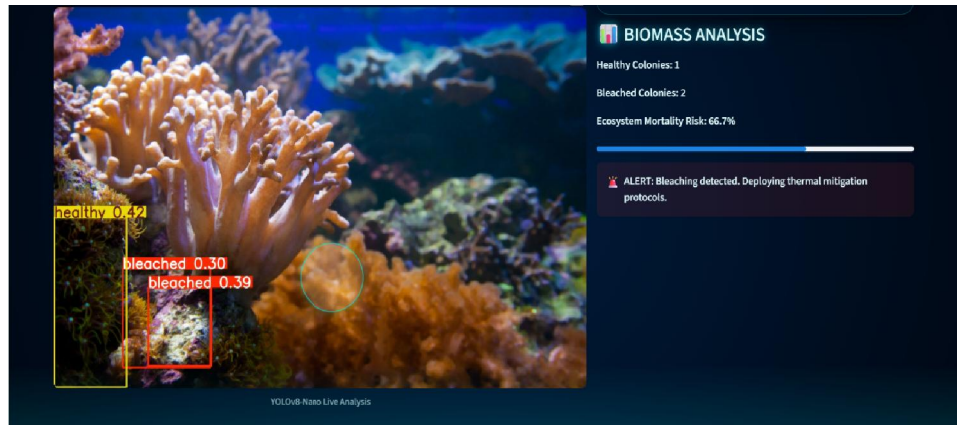


Fig. 4.2 CoralGuard AI User Interface , Oceanic Research Command Center dashboard showing live detection output with bounding boxes (bleached: red, healthy: green), telemetry panel, and biomass analysis.

#### D. Limitations

Several limitations warrant explicit acknowledgement. First, the telemetry values displayed in the dashboard header , temperature, pH, UV index, and dissolved oxygen , are currently hardcoded static placeholders rather than live sensor feeds; connecting them to real-time NOAA buoy APIs or in-situ sensor streams is identified as the highest-priority near-term development task. Second, the training corpus, while diverse in bleaching severity, skews toward Indo-Pacific reef morphologies and may generalise less well to Caribbean or Red Sea reef structures with different colony growth forms and substrate backgrounds. Third, the current two-class scheme does not discriminate between thermally induced bleaching and bleaching attributable to sedimentation, disease, or chemical stressors; a four- or five-class annotation scheme is planned for the subsequent training cycle to address this diagnostic gap.

#### VI. CONCLUSION

CoralGuard AI demonstrates that a carefully scoped, lightweight deep learning application can transfer meaningful reef monitoring capability from specialised research computing infrastructure into the hands of field ecologists and conservation practitioners working under real-world resource constraints. The combination of YOLOv8-Nano's inference efficiency and Streamlit's zero-installation deployment model produces a system that operates on ordinary consumer hardware, requires no command-line interaction, and returns ecologically interpretable colony-level outputs in under one second per frame. These properties matter because the conservation value of a detection model is determined by whether it is actually deployed where bleaching is actively occurring , a criterion that laboratory mAP scores alone cannot satisfy.

A Spearman correlation of 0.87 between automated risk scores and expert assessments, combined with correct alert classification in 14 of 17 test frames, demonstrates that the system's output is ecologically credible rather than merely statistically impressive. The usability evaluation confirms that non-specialist users can interpret outputs accurately without training, removing the expertise barrier that has historically limited field adoption of automated reef assessment tools.

Three development directions are prioritised for future work. The first is live telemetry integration: replacing static environmental indicator placeholders with real-time API calls to NOAA buoy networks and satellite SST products, so that image-level detection is always contextualised by current thermal stress data. The second is an expanded class scheme , separating bleaching causes and severity grades requires additional annotation effort but would substantially increase the tool's diagnostic value for both management decisions and scientific monitoring programmes [8][9]. The



third is mobile optimisation: packaging the application for direct deployment on waterproof tablets used in underwater survey operations, enabling detection to run on freshly captured photographs before the field team surfaces.

Looking beyond the immediate application, CoralGuard AI illustrates a broader design philosophy for ecological monitoring tools: deployment context should drive architectural decisions at least as strongly as benchmark performance. A model that achieves state-of-the-art mAP on a held-out test set but requires a GPU-equipped workstation delivers no conservation value at a remote reef site. Conversely, a model that runs on a smartphone CPU and returns results in under a second can become a routine part of every field survey, generating longitudinal colony-level data at a scale previously impossible. The design choices made in CoralGuard AI, YOLOv8-Nano over larger variants, Streamlit over Flask, browser deployment over desktop packaging - each reflect this field-first philosophy applied consistently from architecture to interface [1][2].

#### **ACKNOWLEDGMENT**

The authors thank Ms. Shristy Goswami, Assistant Professor, for her guidance and critical feedback throughout this project, and Mr. Ankit Verma, Head of the Department of Artificial Intelligence and Machine Learning at Dr. Akhilesh Das Gupta Institute of Professional Studies (GGSIPU), for providing the institutional resources and laboratory access that supported this work. NOAA Coral Reef Watch is acknowledged for maintaining the open-access global SST and DHW products central to this field, and the Ultralytics team for the open-source model library that made rapid fine-tuning feasible.

#### **REFERENCES**

- [1] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," GitHub, 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [2] Streamlit Inc., "Streamlit: The fastest way to build data apps," 2024. [Online]. Available: <https://streamlit.io>
- [3] NOAA Coral Reef Watch, "NOAA CRW Daily Global 5-km Satellite Products," 2024. [Online]. Available: <https://coralreefwatch.noaa.gov>
- [4] T. P. Hughes et al., "Global warming and recurrent mass bleaching of corals," *Nature*, vol. 543, pp. 373–377, 2017.
- [5] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," arXiv preprint arXiv:1804.02767, 2018.
- [6] G. Bradski, "The OpenCV library," *Dr. Dobbs Journal of Software Tools*, 2000.
- [7] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
- [8] C. R. Woolstra et al., "Standardized short-term acute heat stress assays resolve historical differences in coral thermotolerance across microhabitat reef sites," *Global Change Biology*, vol. 26, no. 8, pp. 4328–4343, 2020.
- [9] E. F. Camp et al., "The future of coral reefs subject to rapid climate change: Lessons from natural extreme environments," *Frontiers in Marine Science*, vol. 5, p. 4, 2018.
- [10] C. M. Eakin et al., "The 2014–2017 global-scale coral bleaching event: Insights and impacts," *Coral Reefs*, vol. 38, pp. 539–545, 2019.

