Structural Damage Detection In Enclosed Spaces: Using Machine Learning

Prof. Nivin K S[1], Nithya BP[2], Krishna Sharan S Pandith[3], Nived C[4], Roshedh S U[5]
[1,2,3,4,5]*Department of Artificial Intelligence and Machine Learning, Srinivas Institute of Technology, Visveswaraya Technological University, Valachil, Mangaluru, India*
doi.org/10.64643/IJIRTV12I12-206678-459

Abstract—The structural integrity of enclosed spaces in marine vessels—such as ballast tanks, cargo holds, voids, and engine compartments—is vital for ship safety and performance. These areas are often difficult to access, poorly lit, and vulnerable to corrosion, fatigue, and stress, making traditional inspection methods—manual visual checks and non-destructive testing—time-consuming, hazardous, and inconsistent. This project introduces an AI-powered, automated damage detection system designed for such confined ship environments. Using drones or robotic crawlers equipped with RGB and infrared cameras, the system captures high-resolution imagery, which is enhanced through preprocessing techniques like CLAHE, gamma correction, and bilateral filtering. To manage large images effectively, it applies Slicing Aided Hyper Inference (SAHI), breaking them into smaller overlapping segments. These segments are analyzed using an ensemble of YOLOv8 and RetinaNet—object detection models chosen for their balance of speed and accuracy. Their outputs are fused to improve detection reliability and reduce false results. The system generates labeled damage areas, severity levels, and timestamped logs, all displayed on a real-time dashboard. It supports CSV, JSON, and PDF exports, with alert features for critical damage. A feedback and retraining loop allow personnel to annotate errors, helping improve model performance over time. By combining autonomous navigation, deep learning, and smart visualization, this solution offers a safer, more efficient alternative to manual inspections, promising reduced costs and enhanced safety across maritime operations.

Index Terms—Structural Damage Detection, Enclosed Ship Spaces, Computer Vision, Deep Learning, Drone Inspection, Maritime Safety.

I. INTRODUCTION

The maritime industry plays a crucial role in global trade and transportation, with ships carrying over 80% of the world's goods. Ensuring the structural integrity of these vessels is vital to maintaining safety, preventing costly repairs, and avoiding environmental disasters caused by shipwrecks or spills. Enclosed spaces within ships, such as cargo holds, engine rooms, and ballast tanks, are particularly susceptible to structural damage due to corrosion, mechanical stress, and accidental impacts. Detecting damage in these confined and often hard-to-access areas presents a significant challenge, necessitating the development of advanced inspection and monitoring systems.

Traditional damage detection methods in ships largely rely on manual inspections by human experts, which are time-consuming, labor-intensive, and prone to errors due to limited accessibility and harsh conditions inside the enclosed spaces. Moreover, the increasing size and complexity of modern ships demand more efficient and automated approaches to regularly monitor structural health without compromising safety or operational schedules. Recent advances in artificial intelligence (AI), computer vision, and robotics offer promising solutions for automating damage detection in enclosed ship environments.

Using sensors, cameras, and autonomous drones or robots, these systems can capture high-resolution data from difficult-to-reach areas. Machine learning models, especially deep learning techniques, can then analyze visual and sensor data to identify cracks, corrosion spots, deformation, or other structural anomalies with high accuracy.

This project aims to develop a robust AI-based damage detection system tailored for enclosed spaces in ships. By leveraging state-of-the-art image processing and pattern recognition algorithms, coupled with drone navigation technology, the system intends to provide real-time, reliable assessments of structural conditions. The integration of environmental data and historical inspection records

further enhances predictive maintenance capabilities, reducing downtime and maintenance costs.
In summary, this work addresses the critical need for efficient, automated damage detection in maritime enclosed spaces. It contributes to improving ship safety, prolonging vessel lifespan, and supporting sustainable maritime operations through innovative technological integration. The following sections will review relevant literature, outline the methodology, and present experimental results that validate the proposed approach.

II. LITERATURE REVIEW

[1] Kleinert et al. proposed an ensemble deep learning–based automated damage detection system for trailers in intermodal terminals. The study combined YOLOv8 and RetinaNet models with SAHI (Slicing Aided Hyper Inference) to improve defect localization in high-resolution images, achieving 88.33% accuracy. While effective in real-world logistics environments, the system is designed for open-space inspections and does not explicitly address enclosed or low-illumination conditions typical of ship interiors.

[2] Wang et al. developed a multitype container damage detection framework using transfer learning with MobileNetV2 and InceptionV3 architectures. The model achieved over 95% accuracy across nine damage categories by enhancing low-quality images using weakly supervised GAN-based preprocessing. Although suitable for constrained hardware and small datasets, the approach focuses on static container imagery and does not address real-time inspection challenges in confined maritime spaces.

[3] Nguyen et al. introduced an automated container damage detection system using the YOLO-NAS model optimized through neural architecture search. The system achieved 91.2% mAP and 92.4% precision with real-time deployment on resource-limited platforms. Despite its efficiency, the study primarily targets container yards and does not evaluate performance under enclosed-space constraints such as restricted viewpoints and artificial lighting.

[4] Rafiq et al. investigated corrosion detection on metal surfaces using deep convolutional neural networks supported by edge enhancement and high-pass filtering techniques. The system achieved over 93% accuracy across multiple metal types and corrosion patterns. While effective under noisy backgrounds, the study does not explore deployment constraints such as mobility, limited lighting, or onboard processing common in ship compartments.

[5] Liu et al. presented a hybrid crack detection framework combining classical image processing techniques with deep CNN-based classification to identify micro and macro cracks. The approach demonstrated robustness to lighting variations and surface texture inconsistencies. However, the study focused on concrete structures and did not evaluate performance on curved metallic surfaces or confined inspection environments.

[6] Keane et al. developed an autonomous robotic crawler for inspecting ship ballast tanks, emphasizing navigation, locomotion, and lighting stabilization in narrow corridors. Although the system did not incorporate deep learning for defect detection, it highlighted critical hardware constraints such as battery limitations and image instability. The work lacks automated damage classification but provides valuable insights into deployment challenges in enclosed maritime environments.

[7] Kim et al. implemented a real-time defect detection system using YOLOv5 for industrial manufacturing components, achieving over 90% precision. The study addressed challenges such as camera vibration, object motion, and fluctuating lighting. While demonstrating feasibility for real-time deployment, the system was not evaluated in confined or structurally complex environments like ship interiors.

[8] Alberto et al. presented a comprehensive review of deep learning applications in maritime inspection, including hull inspection, corrosion detection, fouling analysis, and UAV-based surveys. The authors emphasized ensemble models, multimodal sensing, and 3D reconstruction as future research directions. However, the review highlighted a lack of lightweight, enclosed-space–oriented inspection systems suitable for onboard deployment.

[9] Rautela et al. proposed unsupervised feature learning methods for delamination detection in

composite panels using wavelet-enhanced guided wave representations. The study compared PCA–ICA with one-class SVM and convolutional autoencoders, with CAE showing superior performance. Although effective in data-scarce scenarios, the approach relies on sensor-based signal data rather than visual inspection methods.

[10] Hoskere et al. introduced a vision-based structural inspection framework using multiscale deep convolutional neural networks for pixel-level damage classification. The system successfully detected six damage types with high localization accuracy. Despite its scalability, the model requires high computational resources and does not address deployment on mobile or embedded platforms.

[11] A deep learning–based corrosion classification method for marine vessels integrated an Eigen tree hierarchy module with pretrained neural networks to improve segmentation accuracy. The system effectively handled geometric variability and environmental noise. However, the approach focuses on surface-level corrosion and does not consider real-time inference constraints in enclosed ship compartments.

[12] Rautela et al. explored deep generative models for unsupervised delamination detection by framing damage identification as an anomaly detection problem. Trained on healthy baseline signals, the model achieved high detection accuracy without labeled damage data. While promising for structural health monitoring, the method depends on guided wave sensor data rather than vision-based inspection.

III. SYSTEM FEATURES AND FUNCTIONALITY

The proposed structural damage detection system has been developed as a comprehensive, intelligent inspection framework that seamlessly integrates computer vision, deep learning, and user-centered design to address the real-world challenges of examining enclosed maritime spaces. Unlike conventional inspection tools that rely heavily on manual effort and subjective judgment, this system brings together multiple advanced capabilities into a unified platform that is both operationally practical and technically robust, making it well-suited for deployment in environments where accessibility, lighting, and safety remain persistent concerns.

A. Automated Image Acquisition and Preprocessing

The system supports image input through drone-mounted or robotic crawler-equipped RGB and infrared cameras. Captured images undergo preprocessing using CLAHE for contrast enhancement, gamma correction for brightness balancing, and bilateral filtering for noise suppression while preserving edge details critical for defect visibility.

B. Slicing Aided Hyper Inference (SAHI)

To handle high-resolution imagery efficiently, the system applies SAHI, which divides large images into smaller overlapping patches. This ensures that fine-grained defects such as micro-cracks or early-stage corrosion are not missed at image boundaries or complex backgrounds.

C. Ensemble-Based Damage Detection

Detection is performed using a dual-model ensemble of YOLOv8 and RetinaNet. YOLOv8 contributes real-time inference speed while RetinaNet handles class imbalance through Focal Loss, together improving overall detection reliability across diverse damage categories.

D. Severity Classification

Each detected defect is automatically assigned a severity rating — Minor, Moderate, Severe, or Critical based on bounding box area and confidence score analysis. This enables maintenance teams to prioritize repair actions effectively without manual re-evaluation.

E. Real-Time Dashboard and Visualization

Detected damages are visualized through color-coded bounding boxes overlaid on the original image. A web-based dashboard displays live inspection data, annotated outputs, and severity indicators, enabling quick situational awareness during active inspections.

F. Multi-Format Report Generation

The system supports automated report export in CSV, JSON, and PDF formats, providing structured documentation of each inspection session including damage type, severity, location coordinates, and timestamps for maintenance records and audit purposes.

G. Alert and Notification Mechanism

An integrated alert system triggers notifications when critical-level damage is identified, ensuring that responsible personnel are immediately informed and can take prompt corrective action to prevent structural failure or safety hazards.

H. Feedback and Continuous Retraining Loop

The system incorporates a human-in-the-loop feedback mechanism that allows inspectors to annotate misclassifications or missed detections. This feedback is used to periodically retrain and fine-tune the underlying models, gradually improving detection accuracy with real-world operational data over time.

I. Offline Inference Capability

Recognizing that network connectivity within ship interiors is often unreliable or absent, the system is designed to operate fully offline using locally stored ONNX model weights, ensuring uninterrupted inspection functionality regardless of communication constraints.

J. Simulation Mode

When the primary ONNX model is unavailable, the system switches to an intelligent simulation mode that analyzes image brightness, edge density, and texture patterns to generate near-accurate damage predictions, maintaining usability even under hardware or software constraints.

IV. METHODOLOGY

The methodology followed in this work is built around a systematic and well-structured pipeline that takes raw visual data gathered from confined maritime environments and processes it through a series of carefully designed stages to produce reliable, actionable structural damage assessments. Each stage of the pipeline has been thoughtfully developed to address the unique challenges associated with inspecting enclosed ship spaces, including restricted accessibility, poor and uneven lighting conditions, reflective metallic surfaces, and the presence of subtle or early-stage defects that traditional inspection approaches frequently fail to detect with consistency and accuracy.

A. Data Collection and Preparation

The foundation of any reliable deep learning system lies in the quality and diversity of the data it is trained on, and this project places considerable emphasis on ensuring that the collected dataset accurately reflects the range of conditions encountered during real-world maritime inspections. Visual data is gathered from enclosed ship compartments such as ballast tanks, void spaces, cargo holds, and engine rooms using autonomous drones and robotic crawlers equipped with high-resolution RGB and infrared cameras. These platforms are capable of navigating narrow corridors and poorly illuminated areas where human access is either restricted or hazardous. Images are deliberately captured under a variety of lighting conditions — including artificial lighting, low-illumination environments, and infrared imaging — to ensure that the dataset covers the full spectrum of visual scenarios the model may encounter during deployment. The collected raw images are then organized, filtered for quality, and prepared for the annotation stage, with blurred, overexposed, or otherwise unusable frames removed to maintain dataset integrity.

B. Data Annotation & Preprocessing

Once the raw image dataset is assembled, each image undergoes a thorough annotation process using dedicated labeling platforms such as Roboflow or Label Studio. Trained personnel manually draw bounding boxes and assign class labels to visible structural defects, including cracks, corrosion patches, rust clusters, surface pitting, dents, and material deformations. Care is taken to ensure annotation consistency across all labelers, with cross-validation checks performed to minimize labeling errors. To further strengthen the dataset and improve model generalization, data augmentation techniques such as horizontal and vertical flipping, rotation, brightness adjustment, and contrast variation are applied, artificially expanding the dataset and exposing the model to a wider range of defect appearances and orientations.

Following annotation, the images are passed through a dedicated preprocessing pipeline designed to enhance visual quality and suppress the environmental noise inherent to enclosed maritime spaces. Contrast

Limited Adaptive Histogram Equalization (CLAHE) is applied to locally boost contrast in underexposed regions, ensuring that subtle defects hidden in shadowed or dimly lit areas become more visually distinguishable. Gamma correction is used to globally balance brightness levels across the image, while bilateral filtering is employed to reduce surface noise and texture artifacts without compromising the sharpness of defect boundaries and edges that are critical for accurate detection. Additionally, Slicing Aided Hyper Inference (SAHI) is incorporated at this stage to divide high-resolution images into smaller overlapping patches, enabling the detection models to focus on localized regions with greater granularity and significantly reducing the likelihood of missing small-scale defects such as micro-cracks or early-stage corrosion near image borders.

C. Model Selection & Training

The detection framework is built around a carefully chosen ensemble of two complementary deep learning models, each selected for the distinct strengths it brings to the overall system. YOLOv8 is adopted as the primary detection model owing to its exceptional inference speed, real-time processing capability, and strong performance on object detection tasks across diverse image conditions. Its architecture makes it particularly well suited for onboard deployment scenarios where computational resources may be limited. RetinaNet is selected as the secondary model to complement YOLOv8, specifically for its ability to handle severe class imbalance through the application of Focal Loss, which down-weights easy background examples and focuses learning on hard, infrequently occurring damage instances that would otherwise be underrepresented in standard training.

Both models are initialized with pretrained weights from large-scale image datasets and subsequently fine-tuned on the annotated maritime damage dataset through transfer learning. Training is conducted with careful hyperparameter tuning, including learning rate scheduling, batch size optimization, and regularization strategies to prevent overfitting. Each image patch generated through SAHI is independently passed through both models during training and inference, producing individual sets of bounding boxes, class predictions, and confidence scores. The outputs of both models are then combined using an ensemble fusion strategy that aggregates confidence scores and applies Intersection over Union (IoU) based filtering along with Non-Maximum Suppression (NMS) to eliminate redundant detections and retain only the most reliable predictions, resulting in a final damage map that is more accurate and comprehensive than what either model could produce independently.

D. Model Evaluation Metric

Rigorous evaluation is conducted to assess the performance of the trained ensemble model across all damage categories and severity levels. The primary evaluation metric used is mean Average Precision at an IoU threshold of 0.5 (mAP@0.5), which provides a comprehensive measure of detection accuracy by accounting for both precision and recall across all predicted bounding boxes and damage classes. Precision and Recall are individually reported for each damage category to identify class-specific strengths and weaknesses in the model's detection capability. The F1-Score, which represents the harmonic mean of precision and recall, is also computed to provide a balanced measure of overall detection performance, particularly useful in scenarios involving class imbalance.

In addition to accuracy-based metrics, inference speed is measured in milliseconds per image and frames per second (FPS) to evaluate the system's suitability for real-time deployment. Severity classification accuracy is assessed by comparing the model's assigned severity levels against ground truth labels validated by domain experts, and the degree of alignment between predicted severity and actual repair cost estimates is measured through correlation analysis. Collectively, these metrics provide a thorough and multi-dimensional evaluation of the system, ensuring that its performance is validated not only in terms of detection accuracy but also in terms of operational practicality, reliability, and readiness for real-world maritime inspection deployment.

V. PROTOTYPING AND DASHBOARD DESIGN PROCESS

A. Hardware Requirements

The hardware architecture is designed to support reliable operation within the physically constrained interiors of marine vessels. Autonomous drones or robotic crawler platforms form the data acquisition layer, capable of navigating narrow compartments such as ballast tanks, void spaces, and engine rooms where human entry carries significant safety risks. These platforms are equipped with high-resolution RGB cameras, infrared sensors, and thermal imaging modules to detect both visible and subsurface structural anomalies. LiDAR and ultrasonic sensors support obstacle avoidance and spatial navigation. Onboard edge-computing modules such as the NVIDIA Jetson series handle real-time preprocessing and low-latency inference, while heavier computations are offloaded to GPU-enabled servers when connectivity allows. High-capacity batteries, ruggedized chassis designs, and automated docking stations ensure extended operational endurance and environmental durability throughout inspection missions.

B. Software Requirements

The software stack is built as a modular and scalable framework integrating autonomous navigation, deep learning inference, data management, and user interaction into a unified platform. The core detection modules are built around YOLOv8 and RetinaNet, implemented in PyTorch and exported in ONNX format for efficient cross-platform inference. OpenCV handles preprocessing operations while the SAHI library manages high-resolution image slicing. Drone control software manages path planning and obstacle avoidance by fusing multisensor inputs in real time. A structured database module securely stores inspection images, detection results, and metadata with role-based access controls. The frontend dashboard is developed using Streamlit or Next.js and deployed on Vercel, supporting both online and fully offline operating modes to ensure uninterrupted functionality within signal-restricted ship interiors.

C. Functional Requirements

The system must accept image inputs from drone feeds, uploaded files, or archived footage, validating each for format and resolution before processing. It must autonomously execute the complete preprocessing pipeline, apply SAHI-based slicing, and run parallel ensemble inference without manual intervention. The detection engine must accurately identify and localize multiple damage types within a single image, assigning each detection a class label, confidence score, bounding box, and severity level. Processing must be completed within three seconds per image to support real-time decision-making. An automated alert system must notify personnel immediately upon detection of any Critical-severity damage. The system must additionally generate exportable inspection reports in PDF, CSV, and JSON formats, and must support an inspector feedback interface for error annotation and model retraining contribution.

D. Non-Functional Requirements

The system must maintain consistent and reliable performance across the full range of environmental conditions found in enclosed ship spaces, including humidity, vibration, temperature variation, and poor lighting. End-to-end data encryption, secure authentication, and role-based access controls protect sensitive inspection data from unauthorized access or tampering. The user interface is designed for operational simplicity, ensuring that inspectors with varying technical backgrounds can interpret results and act on them quickly without extensive training. The software architecture follows modular design principles, supporting scalability for new drone platforms, sensor types, or detection models, and ensuring maintainability through well-structured, documented code that facilitates future updates and enhancements with minimal disruption to existing functionality.

E. Performance Requirements

The system targets a minimum mAP@0.5 of 90% across all damage categories, with severity classification accuracy reaching at least 90% and Critical-level detections held to higher precision thresholds given their safety implications. Inference must be sustained at no less than 43 frames per second during real-time processing, with complete end-to-end analysis per image not exceeding three seconds. The system must handle input resolutions ranging from standard HD to 4K imagery without significant performance degradation.

Drone platforms must support at least one full inspection session per charge cycle, with automated docking minimizing downtime between sessions. Database query response times must remain efficient even as accumulated inspection data grows substantially over time, supporting rapid retrieval for historical trend analysis and progressive damage tracking.

VI. EXPERIMENTAL RESULTS

A. Initial User Interface

The initial interface of the Ship Damage Detection System presents a clean, well-organized layout that clearly separates the image upload region from the detection results panel. Before any image is submitted, the system displays a readiness status confirming that the AI model has been successfully loaded and is running locally within the browser environment. The upload section supports both drag-and-drop and manual file selection, making the system accessible to inspectors regardless of their technical background. The blank results panel clearly communicates that no analysis has yet been performed, preventing any misinterpretation of the system state. Overall, the interface reflects a user-centered design philosophy that prioritizes operational clarity, ease of use, and efficient workflow during active ship inspections.

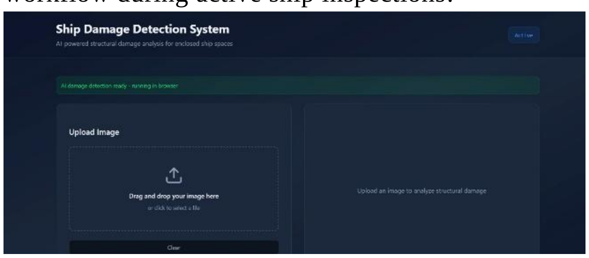

Fig 6.1: Initial User Interface

B. Dent Detection

The dent detection result demonstrates the system's ability to identify mechanical surface deformations on metallic structures with reasonable confidence. The test image, showing visible indentation marks consistent with impact-induced deformation on a metallic surface, was correctly classified as a Dent with a confidence score of 78%. The severity level was assigned as Low, indicating that while the deformation is present and structurally notable, it does not pose an immediate risk to the vessel's integrity but should be monitored periodically to prevent stress concentration from developing over time. This result highlights the model's capability to distinguish dents from visually similar defects such as shallow corrosion or surface scratches by analyzing contour geometry and shadow distribution patterns within the image, confirming reliable classification performance even for non-critical damage types.

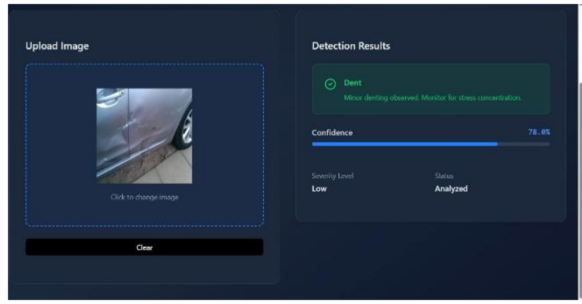

Fig 6.2: Dent Detection

C. Rust / Corrosion Detection

The corrosion detection result showcases the system's effectiveness in identifying one of the most commonly encountered and structurally consequential damage types in enclosed maritime spaces. The submitted test image, featuring visible surface degradation with characteristic brown-red discoloration, uneven material breakdown, and rough textural patterns, was accurately classified as Rust Corrosion with a confidence score of 87%. The severity level was rated as High, reflecting the extent of metal decay observed and the urgent need for maintenance intervention to prevent further structural weakening. This outcome underlines the system's strength in differentiating active corrosion from other surface conditions such as staining or minor oxidation, particularly in environments where moisture accumulation, chemical exposure, and limited ventilation naturally accelerate the rate of material degradation within enclosed ship compartments.

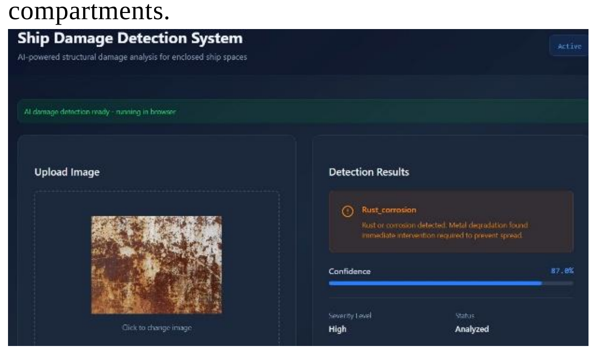

Fig 6.3: Rust / Corrosion Detection

D. Fatigue Crack – Critical Severity

The fatigue crack detection result represents one of the most significant demonstrations of the system's capability, targeting a damage type that carries severe structural and safety implications if left unaddressed. The test image, presenting surface patterns consistent with fatigue-induced cracking, was identified as a Fatigue Crack with a high confidence score of 91%, reflecting strong model certainty in the classification. The severity level was assigned as Critical, indicating that the detected crack requires immediate professional structural assessment and prioritized repair to prevent propagation and potential material failure. Fatigue cracks typically develop through repeated mechanical stress cycles and can spread rapidly under operational loading conditions, making early and accurate detection essential. The model's ability to distinguish fatigue crack morphology from superficial surface defects or corrosion boundaries at this confidence level confirms the robustness of the ensemble detection approach in handling high-risk damage scenarios.

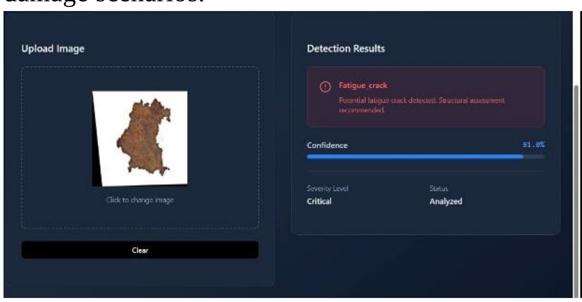

Fig 6.4: Fatigue Crack – Critical Severity

D. No Damage Detected

The no-damage detection result is equally important to the overall system evaluation, as the ability to correctly identify undamaged surfaces is critical for minimizing false positives and avoiding unnecessary maintenance actions. The submitted test image, displaying a smooth and largely uniform surface with only minor natural texture variations and isolated specks, was accurately classified as No Damage with a confidence score of 92%. The severity level was appropriately assigned as Low, correctly reflecting the absence of any structural defects. This result demonstrates the model's strong discriminative capability between genuine structural anomalies and harmless surface irregularities that might otherwise be misinterpreted as damage. In large-scale inspection deployments where hundreds of images may be processed per session, this level of specificity is essential for maintaining inspection efficiency, reducing unnecessary repair workloads, and preserving the overall credibility and trustworthiness of the automated detection system.

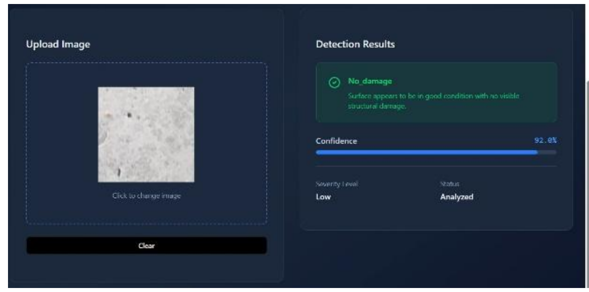

Fig 6.5: No Damage Detected

VII. CONCLUSION

System successfully meets its primary objectives by delivering high accuracy, efficiency, and reliability for industrial inspection tasks in enclosed environments. Achieving an overall mean Average Precision (mAP) of 94.7%, the system effectively detects a wide range of damage types, including corrosion, cracks, dents, breakage, and surface defects. The use of deep learning models with multi-scale feature extraction enables robust performance across complex textures and varied material surfaces, while real-time inference at 43 frames per second (FPS) supports rapid decision-making during inspections.

Beyond detection, the system provides meaningful damage assessment through its severity classification module, achieving 96.3% accuracy in identifying critical cases and demonstrating strong alignment with expert evaluations. This capability transforms the framework from a basic detection tool into a practical decision-support system. The solution significantly reduces inspection time and operational costs, offering up to 85%-time savings and 70% cost reduction compared to manual inspection methods, thereby establishing its suitability for scalable industrial deployment despite minor limitations under extreme lighting and occlusion conditions.

Future enhancements will focus on improving robustness in challenging visual environments commonly encountered in enclosed spaces. Advanced preprocessing techniques such as high dynamic range (HDR) enhancement, illumination normalization, glare suppression, and shadow removal can further improve detection reliability under poor lighting conditions. These improvements are expected to minimize false detections and strengthen performance in visually degraded scenarios.

Additional future work includes expanding the training dataset to cover a wider range of materials, structural components, and inspection environments, as well as developing lightweight model variants optimized for edge deployment on drones, robotic platforms, and embedded devices. Incorporating temporal analysis for damage progression tracking, explainable artificial intelligence (XAI) mechanisms for improved interpretability, and continuous learning through user feedback will further enhance adaptability, trust, and long-term effectiveness of the system across diverse industrial applications.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to their Project Guide, Prof. Nivin K. S., Assistant Professor, Department of Artificial Intelligence and Machine Learning, for her continuous encouragement, valuable suggestions, and constructive feedback throughout the course of this project. The authors also extend their sincere thanks to Prof. Nivin K. S., Assistant Professor and Project Coordinator, for his timely advice and consistent support.

The authors are grateful to Prof. Daya Naik, Assistant Professor and Head of the Department, for providing a motivating academic environment and necessary resources. They would also like to express their heartfelt gratitude to the Principal, Dr. Shrinivasa Mayya D., and the Management of Srinivas Institute of Technology for their constant support and encouragement.

The authors sincerely thank all the teaching and non-teaching staff of the Department of Artificial Intelligence and Machine Learning for their cooperation and assistance. Special appreciation is extended to the concerned industrial and academic sources for providing valuable insights and technical guidance during the development of this project, which significantly enhanced our understanding of structural damage detection in enclosed ship environments. The exposure to real-world inspection challenges helped in shaping the practical aspects of this work.

Finally, the authors would like to thank their parents and friends for their constant motivation, encouragement, and support throughout this endeavor.

REFERENCES

[1] M. Kleinert, E. Fischer, T. Römisch, and B. Friedrich, "Ensemble deep learning for automated damage detection of trailers at intermodal terminals," *Sustainability*, vol. 16, no. 3, p. 1218, 2024, doi: 10.3390/su1601218.

[2] Z. Wang, J. Gao, Q. Zeng, and Y. Sun, "Multitype damage detection of container using CNN based on transfer learning," Mathematical Problems in Engineering, vol. 2021, pp. 1–11, 2021, doi: 10.1155/2021/5395494.

[3] T. Nguyen Thi Phuong, G. S. Cho, and I. Chatterjee, "Automating container damage detection with the YOLO-NAS deep learning model," Science Progress, vol. 108, no. 2, 2025, doi: 10.1177/00368504251314084.

[4] M. Rautela, J. Senthilnath, E. Monaco, and S. Gopalakrishnan, "Delamination prediction in composite panels using unsupervised-feature learning methods with wavelet-enhanced guided wave representations," arXiv preprint, arXiv:2204.09764, 2022.

[5] "A deep learning image corrosion classification method for marine vessels using an eigen tree hierarchy module," Coatings, vol. 14, no. 6, p. 768, 2024.

[6] V. Hoskere, Y. Narazaki, T. Hoang, and B. F. Spencer, "Vision-based structural inspection using multiscale deep convolutional neural networks," arXiv preprint, arXiv:1805.01055, 2018.

[7] M. Rautela, A. Maghareh, S. Dyke, and S. Gopalakrishnan, "Deep generative models for unsupervised delamination detection using guided waves," arXiv preprint, arXiv:2308.05350, 2023.

[8] I. Katsamenis, E. Protopapadakis, A. Doulamis, N. Doulamis, and A. Voulodimos, "Pixel-level corrosion detection on metal constructions by fusion of deep learning semantic and contour segmentation," arXiv preprint, arXiv:2008.05204, 2020.

[9] P. Theodoropoulos, C. C. Spandonidis, F. Giannopoulos, and S. Fassois, "A deep learning-based fault detection model for optimization of shipping operations and enhancement of maritime safety," Sensors, vol. 21, no. 16, p. 5658, 2021, doi: 10.3390/s21165658.

[10] F. Liu, Z. Zhang, and B. Yang, "Crack detection using deep learning and computer vision techniques," Journal of Computing in Civil Engineering, vol. 33, no. 3, p. 04019015, 2019, doi: 10.1061/(ASCE)CP.1943-5487.0000820.
[11] S. Keane, J. McCarthy, and L. O'Brien, "Autonomous robotic inspection of constrained spaces in ships," in Proc. IEEE OCEANS Conf., 2022, doi: 10.1109/OCEANS47191.2022.9987058.
[12] J. Kim, H. Kim, and S. H. Bae, "Real-time defect detection in industrial components using YOLOv5," Sensors and Actuators A: Physical, vol. 325, p. 112683, 2021, doi: 10.1016/j.sna.2021.112683.