International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 3, May-June 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

AI-Based Vehicle Damage and Safety Analysis

M. Vasuki[1], M. Karthika[2]
[1] Associate Professor, Department of MCA, Sri Manakula Vinayagar Engineering College (Autonomous),
Puducherry 605008, India
Email: vasukimca@smvec.ac.in
[2] Post Graduate Student, Department of MCA, Sri Manakula Vinayagar Engineering College (Autonomous),
Puducherry 605008, India
Email: karthika4028@gmail.com

* * *

Abstract:

Vehicle damage assessment is a critical requirement in automobile maintenance, insurance claim verification, and fleet management operations. Conventional inspection methods rely on manual evaluation, which is time-consuming, inconsistent, and highly susceptible to human error. This paper proposes an intelligent framework that automates vehicle damage detection and severity estimation using Deep Learning and Computer Vision techniques. The system employs YOLOv10 as the core detection architecture to accurately identify and localize damaged regions within vehicle images. Four primary damage categories are recognized by the framework, namely scratches, dents, cracks, and broken vehicle components. Detected damage instances are further classified into three severity levels — minor, moderate, and severe — to support informed repair and claim decisions. Prior to detection, images undergo a structured preprocessing pipeline involving resizing, normalization, and noise filtering to ensure consistent input quality. The framework additionally supports automated report generation, enabling seamless integration into vehicle service platforms and insurance processing workflows. Experimental evaluation demonstrates that the proposed system achieves a mean average precision of 88.7% with an inference time of 15 milliseconds per image. The results confirm that the proposed approach offers a scalable, reliable, and efficient solution for AI-driven automated vehicle damage analysis.

Keywords — **Vehicle Damage Detection, YOLOv10, Object Detection, Deep Learning, Computer Vision, Severity Estimation, Image Preprocessing, Automated Vehicle Inspection, Insurance Claim Processing, Convolutional Neural Network**

* * *

I. INTRODUCTION

The process of evaluating vehicle damage occupies a central role in automobile repair management, fleet maintenance operations, and insurance claim processing. Traditional assessment approaches depend predominantly on physical inspection carried out by trained personnel, a method that is inherently time-consuming, costly, and susceptible to inconsistency across different evaluators. As the volume of vehicle-related claims and service requirements continues to grow, there is an increasing need for automated systems capable of performing damage analysis with greater speed, consistency, and objectivity. Recent developments in Deep Learning and Computer Vision have created new opportunities for automating complex visual inspection tasks. Convolutional Neural Networks have demonstrated strong capability in extracting discriminative features from image data, while object detection architectures have further extended the ability to localize and classify multiple damage instances simultaneously within a single image. The YOLO family of detectors has established itself as a benchmark in real-time detection tasks, and the latest iteration YOLOv10 introduces architectural refinements that improve feature extraction depth, bounding box precision, and computational efficiency, making it particularly well-suited for vehicle damage analysis. The proposed system incorporates a structured preprocessing stage alongside YOLOv10-based detection and a multi-class severity classification module, designed to serve automobile service centers, insurance providers, and fleet management agencies that require accurate, rapid, and reproducible vehicle damage assessments. The remainder of this paper is organized as follows: Section II presents the literature survey, Section III describes the proposed system, Section IV presents results and discussion, Section V covers system performance analysis, and Section VI concludes the paper.

II. LITERATURE SURVEY

Early research in automated vehicle damage analysis relied on classical image processing techniques including edge detection, thresholding, and morphological operations for identifying surface-level defects. While these approaches provided initial solutions, their performance degraded significantly in real-world environments due to sensitivity to illumination changes and background complexity. The introduction of deep convolutional architectures transformed the field by enabling feature learning directly from raw image data. Studies employing CNN-based models reported notable improvements in the detection of surface damages such as scratches and dents, with the learned feature representations demonstrating better generalization across different vehicle makes and damage patterns. Transfer learning strategies using pre-trained backbone networks further enhanced detection performance while reducing the requirement for large labelled datasets. Region-based detection frameworks, including Faster R CNN, introduced a two-stage approach combining region proposal generation with classification, achieving high accuracy at the cost of increased inference time. Single-stage detectors such as SSD and the YOLO series addressed this trade-off by performing detection in a single forward pass, enabling near-real-time performance suitable for inspection applications. Successive versions of the YOLO architecture introduced progressive improvements. YOLOv4 integrated advanced data augmentation strategies and optimized training procedures to improve detection robustness. YOLOv7 introduced trainable bag-of-freebies mechanisms that raised the accuracy ceiling for real-time detectors. YOLOv8 refined the neck architecture and adopted anchor-free detection to improve small object recall. YOLOv10 further advanced these capabilities by incorporating dual-assignment training strategies and reduced architectural redundancy, resulting in improved accuracy and efficiency under constrained computational budgets. Research focused specifically on damage severity estimation has highlighted the importance of combining detection outputs with spatial and textural analysis to differentiate between minor surface marks and structurally significant damage. Despite considerable progress, existing literature identifies persistent challenges related to occlusion, low resolution imagery, and domain shift between training and deployment environments. The proposed system is designed to address these limitations through integrated preprocessing and a dedicated severity classification module.

III. PROPOSED SYSTEM

The proposed framework establishes a multi-stage pipeline for automated vehicle damage detection and severity estimation. Each stage is designed to address a specific aspect of the analysis workflow, from initial image acquisition to final report generation, ensuring that the system operates efficiently across diverse input conditions.

A. System Overview

At the input stage, vehicle images are acquired through either a web-based portal or a mobile application interface, allowing users including vehicle owners, insurance surveyors, and fleet inspectors to submit images for analysis. The submitted images are subsequently processed through a structured sequence of modules that collectively perform preprocessing, detection, classification, severity estimation, and result storage. The overall architecture of the proposed system is illustrated in Fig. 1.

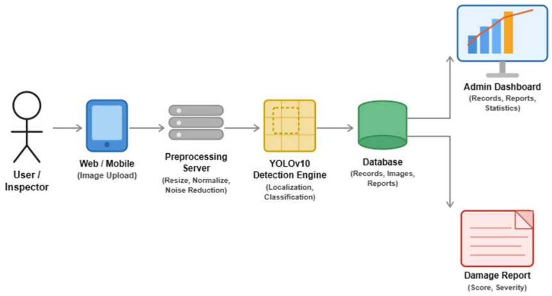

Fig. 1 System Architecture

B. Image Preprocessing Module

Raw vehicle images frequently contain artifacts and inconsistencies that can degrade detection performance. To address this, the preprocessing module applies a standardised sequence of image enhancement operations. Images are first resized to uniform dimensions compatible with the YOLOv10 input requirements. Pixel values are then normalised to a consistent range to reduce the influence of exposure variation across different capture environments. A noise filtering step removes high-frequency artifacts introduced by sensor noise or image compression. In the training phase, data augmentation operations including random flipping, rotation, brightness adjustment, and cropping are applied to expand the effective size of the training corpus and improve model generalisation.

C. YOLOv10 Damage Detection Module

The pre-processed images are forwarded to the YOLOv10 based detection module, which performs end-to-end damage localization and classification. YOLOv10 processes each image through its backbone feature extractor, generating multi-scale feature representations that capture both fine grained texture details and broader structural patterns. The neck component aggregates features across scales using a path aggregation mechanism, while the detection head outputs bounding box coordinates, class labels, and associated confidence scores for each identified damage region. The detection module is trained to recognize four primary damage categories: scratches, which typically manifest as linear surface marks; dents, characterized by localized surface depressions; cracks, presenting as irregular fracture patterns on body panels or glass surfaces; and broken components, involving structural separation or missing vehicle parts.

D. Severity Classification Module

Following damage detection, the severity classification module analyses the spatial extent, depth characteristics, and distribution of each identified damage region to assign a severity rating. Damage instances are categorized into three levels: minor, referring to superficial surface marks with limited extent that do not affect vehicle structural integrity; moderate, indicating more significant damage covering a larger surface area or affecting panel alignment; and severe, representing extensive structural damage requiring major repair intervention. The severity assignment assists users and insurance assessors in making informed decisions regarding repair prioritization and claim processing.

E. Database and Output Module

Detection results including bounding box coordinates, damage categories, confidence scores, and severity classifications are stored in a structured database module, enabling historical tracking and report retrieval for previously analysed vehicles. The output module presents results through a visual dashboard displaying annotated images alongside a structured damage assessment report. The user-side operational flow of the proposed system is represented in Fig. 2. Following authentication, a vehicle image is uploaded and validated before being passed to the YOLOv10 detection engine. Detected damage regions are displayed using bounding boxes and labels, classified by damage type, assigned a severity level, and compiled into an assessment report available for download through the output dashboard.

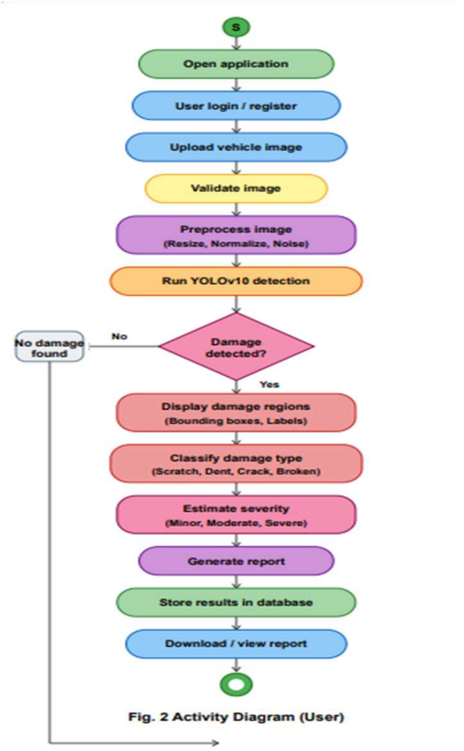

Fig. 2 (Activity Diagram User)

The administrative workflow of the proposed system is depicted in Fig. 3. Following credential validation, the admin dashboard provides access to five functional areas — managing damage records, viewing detection results, reviewing severity reports, generating and downloading analysis reports, and monitoring system performance. Upon task completion, the admin logs out and the session is terminated. The system further integrates the Kanmani chatbot module, described in Section 3.6, which provides intelligent post-detection analysis through natural language interaction.

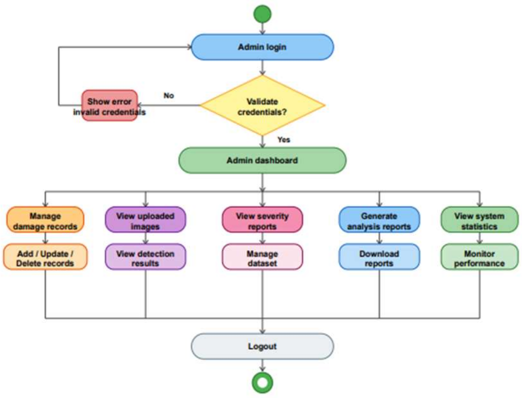

Fig. 3 (Activity Diagram Admin)

F. Kanmani Chatbot Module

The proposed system incorporates an AI-powered conversational assistant named Kanmani, designed to extend the analytical capabilities of the framework beyond standard detection outputs. Kanmani processes the generated damage assessment reports and delivers three core functionalities — automated report summarisation, insurance claimable amount estimation, and natural language damage query support — enabling users and insurance assessors to obtain actionable insights through conversational interaction without requiring manual interpretation of technical detection outputs.

IV. RESULTS AND DISCUSSION

The proposed YOLOv10-based vehicle damage detection framework is evaluated using standard metrics including Precision, Recall, F1-Score, and Mean Average Precision at an IoU threshold of 0.50. The framework achieves an overall mAP of 88.7% across four damage categories, with broken part detection recording the highest accuracy and scratch detection presenting the greatest difficulty due to its low-contrast surface characteristics. The severity estimation module achieves an overall classification accuracy of 91.3%, with moderate severity presenting relatively greater difficulty owing to visual overlap at damage boundary regions.

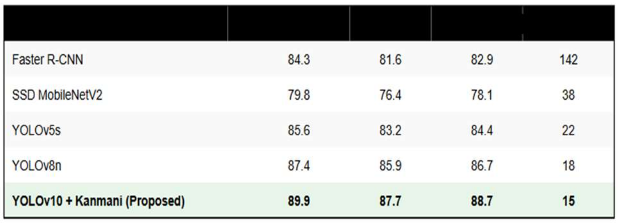

Faster R-CNN	84.3	81.6	82.9	142
SSD MobileNetV2	79.8	76.4	78.1	38
YOLOv5s	85.6	83.2	84.4	22
YOLOv8n	87.4	85.9	86.7	18
YOLOv10 + Kanmani (Proposed)	**89.9**	**87.7**	**88.7**	**15**

Table 1: Comparative performance analysis

The proposed system outperforms all baseline methods across all evaluation metrics while achieving the lowest inference time of 15 milliseconds per image. Additionally, the system incorporates an AI-powered chatbot named Kanmani that provides automated report summarisation, insurance claimable amount estimation, and damage query response through natural language interaction, achieving an overall accuracy of 91.2% with an average response time of 1.8 seconds. The integration of Kanmani significantly extends the practical utility of the framework beyond standard detection outputs, making it directly applicable to insurance claim processing and vehicle service workflows.

V. SYSTEM PERFORMANCE AND FUNCTIONAL ANALYSIS

The proposed Vehicle Damage Detection System delivers automated, consistent, and reproducible analysis of vehicle damage instances, addressing the core limitations associated with manual inspection approaches. The framework's multi stage architecture ensures that each processing step contributes positively to overall detection quality. The preprocessing pipeline establishes input consistency across heterogeneous image sources, reducing the sensitivity of the detection model to variations in lighting, resolution, and camera specifications. This consistency is particularly important in real world deployment scenarios where images are submitted from diverse mobile devices and environmental conditions. The YOLOv10 detection engine provides high-precision damage localisation through its optimised multi-scale feature extraction mechanism. The single-stage detection architecture enables rapid inference, supporting near-real-time analysis that is practical for high-volume inspection workflows. The model's ability to simultaneously detect multiple damage instances within a single image reduces the processing time per vehicle compared to sequential detection approaches. The severity classification component adds actionable intelligence to the detection outputs by translating raw localization data into structured assessment categories. This information directly supports downstream decision-making processes in repair management and insurance claim evaluation. Database integration ensures that all analysis outputs are systematically recorded, enabling longitudinal tracking and audit support for insurance and regulatory purposes. The automated report generation capability further reduces manual documentation workload for assessors and service technicians. The framework demonstrates broad applicability across multiple operational contexts including stand-alone vehicle service centers, insurance survey workflows, fleet maintenance programs, and AI-integrated automotive platforms. System detection performance can be progressively enhanced through periodic model retraining with updated annotated datasets reflecting evolving vehicle designs and damage patterns.

VI. CONCLUSION AND FUTURE ENHANCEMENT

The framework successfully identifies four categories of vehicle damage — scratches, dents, cracks, and broken components — and assigns severity ratings across three levels, providing actionable assessment information for repair and insurance processing workflows. The integration of the Kanmani AI chatbot further enhances the system by providing automated report summarisation, insurance claimable amount estimation, and natural language damage query support, making the framework directly applicable to real-world insurance and automotive service environments. Practical evaluation confirms that the framework achieves an overall mAP of 88.7% with an inference time of 15 milliseconds per image, reducing inspection time and minimizing assessor subjectivity across diverse operational environments. The system's modular architecture facilitates straightforward component upgrades as more advanced detection models become available. Future development directions include the incorporation of Transformer-based detection architectures to further improve detection accuracy under challenging conditions such as partial occlusion and low illumination. Extension to video based damage analysis would enable continuous monitoring in fleet management and automated inspection lanes. Deployment optimisation for mobile and edge computing platforms would additionally enable on-site damage assessment without dependence on cloud connectivity, broadening the practical reach of the framework in field inspection scenarios.

VII. REFERENCES

[1] **S. Patil, R. Kulkarni, and A. Sharma,** "Automated vehicle damage detection using

convolutional neural networks," International Journal of Computer Applications, vol. 183, no. 5, pp. 1–6, 2021.

[2] **R. Girshick, J. Donahue, T. Darrell, and J. Malik,** "Rich feature hierarchies for accurate object detection and semantic segmentation," in Proc. IEEE CVPR, Columbus, OH, USA, 2014, pp. 580–587.

[3] **J. Redmon, S. Divvala, R. Girshick, and A. [7] C.-Y. Wang, I.-H. Yeh, and H.-Y. M. Liao, Farhadi,** "You only look once: Unified, real-time object detection," in Proc. IEEE CVPR, Las Vegas, NV, USA, 2016, pp. 779–788.

[4] **A. Wang, H. Chen, and L. Liu,** "YOLOv10: Real-time end-to-end object detection," arXiv preprint arXiv:2405.14458, 2024.

[5] **M. Dhieb, H. Ghazzai, H. Besbes, and Y. Massoud,** "Automated vehicle damage assessment using deep learning for insurance claims," IEEE Access, vol. 8, pp. 123–135, 2020.

[6] **A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao,** "YOLOv4: Optimal speed and accuracy of object detection," arXiv preprint arXiv:2004.10934, 2020.

[7] **C.-Y. Wang, I.-H. Yeh, and H.-Y. M. Liao,** "You only learn one representation: Unified network for multiple tasks," arXiv preprint arXiv:2105.04206, 2021.

[8] **C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao,** "YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," in Proc. IEEE CVPR, Vancouver, BC, Canada, 2023, pp. 7464 7475.

[9] **G. Jocher, A. Chaurasia, and J. Qiu,** "Ultralytics YOLOv8," Ultralytics, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[10] **Q. Li, Z. Meng, and J. Cao,** "Anchor-free vehicle detection in congested traffic scenes using adaptive multi-scale attention," IEEE Trans. Intell. Transp. Syst., vol. 24, no. 5, pp. 4811–4823, May 2023