

A System for Automated Vehicle Damage Localization and Severity Estimation Using Deep Learning

Chaitra KA ¹, Dr. Geetha M ²

¹ Student, 4th Semester MCA, Department of MCA, BIET, Davanagere

² Associate Professor, Department of MCA, BIET, Davanagere

ABSTRACT

This paper introduces a sophisticated and comprehensive system for automated vehicle damage detection and severity estimation by leveraging the advanced YoloV8 deep learning model. The proposed solution operates through a web-based interface developed using Python, HTML, CSS, and JavaScript, hosted on the Flask framework. It supports three operational modes: image-based, video-based, and webcam-based predictions, thereby enabling real-time as well as post-incident evaluations. The model is trained on a custom-labeled dataset of 778 images that cover eight distinct vehicle damage categories. Achieving an accuracy rate of 91%, the system enhances reliability in assessing vehicle damages. Designed for scalability and versatility, this system offers valuable applications in domains such as automotive insurance claims, vehicle repair cost assessments, accident investigation, and fleet management.

Keywords:

Vehicle damage detection, YoloV8, deep learning, severity estimation, automated assessment, image classification, real-time analysis, web-based interface, Flask framework.

I. INTRODUCTION

Traditional vehicle damage assessment processes involve manual inspection, photographic evidence, and extensive documentation, often resulting in inefficiencies and inconsistencies. These methods are time-intensive and prone to human error, delaying critical decisionmaking in insurance and repair workflows. To address these limitations, this paper proposes a deep learning-based automated system that utilizes the YoloV8 object detection model to identify and classify vehicle damage with high precision.

seen significant advancements in recent

transitioning from manual visual inspections to intelligent, automated solutions powered by computer vision and deep learning technologies.

Integrated into a web application via the, recorded videos, significantly improving accessibility and usability. The interface caters to professionals (e.g., insurers, auto repair technicians) as well as end-users such as vehicle owners. Through the use of a robust model architecture and a diverse training dataset, the system offers detailed insights into specific damage types and severity levels, thereby enabling faster, more accurate, and cost-efficient assessments.

II. RELEATED WORK

Traditionally, vehicle damage assessments were carried out by human inspectors, often leading to inconsistent and time-consuming results.

Khan et al. proposed a UAV-based system for automated pothole detection and severity estimation using deep learning. High-resolution drone imagery was used to collect visual data from urban and rural roads. The system employed a modified Faster R-CNN architecture integrated with a pothole severity ranking module to classify the extent of road damage. The model was trained on a custom UAV dataset and validated against manually labeled ground truth data.

Zhang et al. developed a deep learning framework based on Residual Networks (ResNet) to recognize various road surface conditions from image data. Their dataset consisted of road images captured under different environmental scenarios, such as dry, wet, icy, and snowy roads. By applying deep residual learning, the system effectively mitigated the vanishing gradient problem common in deep CNNs.

Kim and Sohn presented an advanced crack detection method that combines CNNs with high-dimensional texture descriptors to identify structural cracks in road and bridge infrastructure. The model leverages statistical texture measures including Local Binary Patterns (LBP) and Gray-Level Co-occurrence Matrices (GLCM) as feature vectors. These features are fed into a CNN to enhance the differentiation of cracks from background noise.

Alaskar et al. designed an end-to-end vehicle damage assessment system utilizing deep learning-based...The system combines YOLOv3 for vehicle multilayer perceptron (MLP) to classify road irregularities.

Li et al. introduced a real-time vehicle accident detection framework using the YOLO object detection model integrated with a video stream analysis pipeline. The system monitors real-time road surveillance or dashcam footage to identify vehicle collisions based on abnormal motion patterns, sudden stops, and object deformations. It combines object tracking with spatiotemporal behavior modeling to differentiate normal traffic from crash events.

Zhang et al. proposed a severity classification model for traffic collisions using dashcam footage and

and damage localization with a severity classification module trained on annotated accident datasets. The architecture supports automatic extraction of damage regions and evaluates severity levels based on damage area, location, and image features. The system was deployed on a web interface for integration with insurance and traffic management platforms.

Wang and Wu proposed a hybrid method for road crack detection by combining deep learning with histogram-based image thresholding. The system uses a CNN-based segmentation model to identify candidate crack regions, which are then refined using adaptive histogram thresholding to improve edge clarity and reduce false positives. The approach was tested on both synthetic and real-world datasets of urban pavement images.

Maeda et al. introduced one of the early smartphone-based solutions for detecting road surface damage using classical image processing techniques. The application utilizes the smartphone's built-in camera to capture road images and applies edge detection and pattern matching to identify cracks and potholes.

Wang et al. developed a deep learning model integrated into a smartphone application for detecting and classifying road anomalies such as potholes, bumps, and cracks. Using accelerometer and gyroscope data in combination with GPS and camera feeds, the system employs a

deep transfer learning. The model applies a fine-tuned ResNet-101 to extract features from frames before, during, and after a crash event. These features are used to classify collisions into minor, moderate, and severe categories.

III. METHODOLOGY

1. Data Collection and Preprocessing

The proposed system uses a dataset of 778 labeled vehicle images, each belonging to one of eight categories of damage: damaged door, damaged window, damaged headlight, damaged mirror, dent, damaged hood, damaged bumper, and damaged

windshield. The dataset was gathered under varied conditions to ensure diversity, including different angles, lighting variations, and levels of severity. **Image Normalization:** Pixel values are adjusted to the range $[0, 1]$ to maintain uniformity and improve training efficiency. **Data Augmentation:** To increase robustness, transformations such as rotation, zooming, flipping, and adjustments in brightness and contrast are applied, enabling the model to handle a broader variety of real-world scenarios. **Dataset Splitting:** The dataset is divided into training (80%), validation (10%), and testing (10%) subsets to evaluate performance and ensure generalization on unseen data.

2. Model Selection and Training

The system is developed using the YOLOv8 model, a state-of-the-art object detection algorithm recognized for high accuracy and real-time processing. It was chosen for its efficiency even on devices with limited computational capacity. The training process includes: **Model Initialization:** Starting with pre-trained weights from a large dataset, making use of transfer learning to speed up training and improve accuracy. **Fine-Tuning:** Training continues on the labeled vehicle damage dataset so the model learns to detect and classify different damage categories. **Loss Function:** A balanced loss function is applied to optimize both bounding box localization and classification accuracy. **Hyperparameter Tuning:** Parameters like learning rate, batch size, and training epochs are adjusted to enhance performance.

3. Multi-modal Input Handling

The system is designed to support different input formats in order to provide flexibility for users. In image-based mode, static pictures of damaged vehicles are uploaded and analyzed to detect the affected parts and classify the damage. In video-based mode, the system examines video files frame by frame, making it possible to identify damage throughout the sequence. In webcam-based mode, live video streams are processed instantly, which

allows immediate inspection and reporting at workshops or service points. Each of these input types is processed in a way that best suits its application, ensuring both accuracy and efficiency.

4. Evaluation Metrics

The model's effectiveness is assessed using several measures. Accuracy indicates how many damage cases are correctly identified compared to the total tested. Precision and recall are used to measure the correctness of predictions and the system's ability to detect actual damages without missing them. Intersection over Union (IoU) is applied to evaluate how closely the predicted bounding boxes match with the ground truth boxes, which is vital for detection quality. The F1 score combines precision and recall to provide a balanced overall evaluation of the model.

5. Deployment and User Interface

The application is deployed as a web-based system built on the Flask framework. The interface is created using HTML, CSS, and JavaScript to ensure it works smoothly on desktops, tablets, and mobile devices. For deployment, the backend integrates the YOLOv8 model with Flask so that inputs can be processed and predictions delivered through RESTful APIs. The frontend is designed to be responsive and easy to use, giving users a simple way to interact with the system regardless of device.

6. Real-Time Processing Integration

The system is capable of handling live video and webcam inputs, which enables real-time processing of vehicle damages. This is made possible through efficient inference where YOLOv8 processes each video frame with very little delay. Such real-time performance is especially useful for quick inspections, roadside checks, or immediate assessments at service locations.

System Architecture




Fig1. System Architecture

IV MODULE DESCRIPTION

Image Upload Module:

Purpose: Allows users to upload images for damage analysis.

Functionality: Supports various file formats and sizes, ensuring compatibility with the system.

Pre-processing Module:

Purpose: Prepares the input data for analysis.

Functionality: Includes resizing, normalization, and augmentation techniques to enhance model accuracy.

Damage Detection Module:

Purpose: Identifies and classifies vehicle damages.

Functionality: Utilizes the YoloV8 model to detect and categorize damages into predefined classes.

Severity Estimation Module:

Purpose: Assesses the severity of detected damages.

Functionality: Analyzes the size and impact of damages to estimate repair costs.

Reporting Module:

Purpose: Generates comprehensive damage reports.

Functionality: Compiles detection results, severity assessments, and recommendations into a user friendly format.

V. RESULT

The system was rigorously evaluated on a comprehensive dataset consisting of 778 meticulously labeled vehicle images encompassing eight distinct damage categories. Achieving an impressive accuracy rate of 91%, the model demonstrated strong reliability in correctly detecting and classifying various types of vehicle damages. Beyond static images, the system's capability to process multiple input modalities—including still images, recorded videos, and live webcam streams—showcases its adaptability and robustness in real-world applications. This multi-modal input support enables the system to function effectively in a wide range of environments, from offline claim assessments to real-time damage inspections at service centers or accident sites. The adoption of the advanced YoloV8 model markedly improved detection precision and speed compared to previous approaches, such as those relying on Mobile Net or earlier YOLO versions, thus offering a significant advancement in automated vehicle damage evaluation.




Fig2. Result Graph

VI. CONCLUSION

In summary, the developed system presents a **scalable, efficient, and highly dependable solution** for automating vehicle damage localization and severity estimation, addressing critical limitations found in traditional and earlier AI-based approaches. By harnessing the advanced deep learning capabilities of the YoloV8 model, the

system achieves a **notable accuracy of 91%**, ensuring precise identification and categorization of a diverse range of damage types. Its multimodal functionality, supporting images, videos, and live webcam streams, significantly broadens its applicability, enabling both **real-time inspections and post-incident analyses** across various contexts. This versatility allows stakeholders—ranging from insurance companies and automotive repair shops to law enforcement agencies—to estimation, the system not only accelerates claim validation and repair cost estimation

Overall, this system marks a significant step forward in intelligent vehicle damage assessment technology, offering a robust foundation for future enhancements such as integration with telematics data, augmented reality-based repair guidance, and

streamline their workflows, reduce processing times, and minimize human error.

Moreover, the user-friendly web interface built on the Flask framework facilitates **easy accessibility** for professionals and consumers alike, promoting widespread adoption. The system's ability to handle high-resolution inputs and large datasets ensures that it can be deployed effectively

expanded damage classification categories. Its deployment has the potential to transform insurance operations, improve customer experience, and optimize automotive maintenance processes globally.

VII. REFERENCES

1. Y. Ma and colleagues proposed an automated framework for identifying the location and severity of vehicle damage through deep learning. Their study was published in the *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 6, pp. 5627–5639, June 2024. DOI: 10.1109/TITS.2023.3334616.
- 2.K. Doshi and Y. Yilmaz presented an ensemble deep learning approach for detecting road damage. The work appeared in *arXiv preprint* arXiv:2011.00728, 2020, and is also indexed with DOI: 10.1109/ACCESS.2024.3366990.
3. M. Roy and J. Bhaduri introduced a vision-based model for damage detection that improves the YOLOv5 architecture.