

Intelligent Vehicle Damage Assessment & Cost Estimator

Mahesh Panjwani¹, Vikrant Salunke², Shantanu Pathak³, Yash Shahade⁴, Sayali Bawanthade⁵,
Sanika Najpande⁶

¹UG - Assistant Professor Department of Comp. Technology Nagpur, Maharashtra

²UG - Computer Engineering, RTMNU University, Nagpur, Maharashtra

³UG - Computer Engineering, RTMNU University, Nagpur, Maharashtra

⁴UG - Computer Engineering, RTMNU University, Nagpur, Maharashtra

⁵UG - Computer Engineering, RTMNU University, Nagpur, Maharashtra

⁶UG - Computer Engineering, RTMNU University, Nagpur, Maharashtra

Abstract

The "Intelligent Vehicle Damage Assessment & Cost Estimator" project aims to assist vehicle owners by providing accurate and automated damage assessment using advanced deep learning techniques. Leveraging YOLOv5, a state-of-the-art object detection model, this system is designed to analyze and evaluate damage sustained by vehicles, including both two-wheelers and four-wheelers. The model is trained on a comprehensive dataset of vehicle damages to accurately identify and classify different types of damage. By integrating this technology, users can independently assess vehicle damage, reduce dependency on manual inspection, minimize errors, and obtain precise cost estimations for repairs, ensuring transparency and better decision-making.

Keywords: Vehicle damage assessment, cost estimation, YOLOv5, user assistance, deep learning, object detection, two-wheeler, four-wheeler.

1. INTRODUCTION

The rapid expansion of the automotive industry has led to a significant increase in vehicle ownership, resulting in a corresponding rise in insurance claims related to road accidents and vehicle damages. Traditional damage assessment methods rely heavily on manual inspection by surveyors, which is time-consuming, prone to human error, and often inconsistent. To address these limitations, intelligent systems powered by artificial intelligence and computer vision have emerged as transformative solutions for individual vehicle users and owners.

This research introduces an Intelligent Vehicle Damage Assessment & Cost Estimator system designed for users that leverages deep learning techniques for automated damage detection and evaluation. The foundation of the proposed system is based on advanced Convolutional Neural Networks (CNNs) and the state-of-the-art object detection model YOLOv5 (You Only Look Once version 5), which enables real-time and highly accurate detection of vehicle damages.

Deep learning plays a crucial role in extracting meaningful features from raw vehicle images. Unlike conventional image processing approaches, CNN-based architectures automatically learn hierarchical feature representations such as edges, textures, shapes, and complex damage patterns. This capability allows the system to identify various types of vehicle damages including dents, scratches, cracks, broken parts, and paint defects under diverse environmental conditions.

1.1 Methods & Material

The methodology consists of dataset preparation, preprocessing, model training, evaluation, and deployment. A diverse dataset containing images of damaged two-wheelers and four-wheelers was collected. Damage types included dents, scratches, cracks, broken headlights, bumper damage, and paint defects.

Images were manually annotated using bounding boxes to mark damaged regions. The dataset was split into training (70%), validation (15%), and testing (15%) sets. All images were resized to 640×640 pixels and normalized before training.

Data augmentation techniques such as rotation, flipping, brightness adjustment, scaling, and mosaic augmentation were applied to improve generalization and reduce overfitting

2. Purpose

The proposed Intelligent Vehicle Damage Assessment and Cost Estimator system was experimentally evaluated using a diverse test dataset consisting of damaged two-wheelers and four-wheelers captured under varying environmental conditions.

The trained YOLOv5 model demonstrated strong performance in detecting and classifying multiple damage categories including dents, scratches, cracks, and broken components.

Quantitative evaluation metrics such as Precision, Recall, F1-Score, and Mean Average Precision (mAP@0.5) were used to measure detection accuracy and localization efficiency. The localization capability of the model was validated through Intersection over Union (IoU) analysis. Predicted bounding boxes showed significant overlap with ground truth annotations, confirming accurate damage region identification. Larger damages such as bumper dents and broken headlights were detected with very high confidence scores.

Images were manually annotated using bounding boxes to mark damaged regions. The dataset was split into training (70%), validation (15%), and testing (15%) sets. All images were resized to 640×640 pixels and normalized before training.

Data augmentation techniques such as rotation, flipping, brightness adjustment, scaling, and mosaic augmentation were applied to improve generalization and reduce overfitting.

However, comparatively small damages such as fine scratches presented minor detection challenges due to limited pixel density and lighting reflections. The application of data augmentation techniques, including brightness variation and mosaic augmentation, substantially improved detection robustness for such subtle defects.

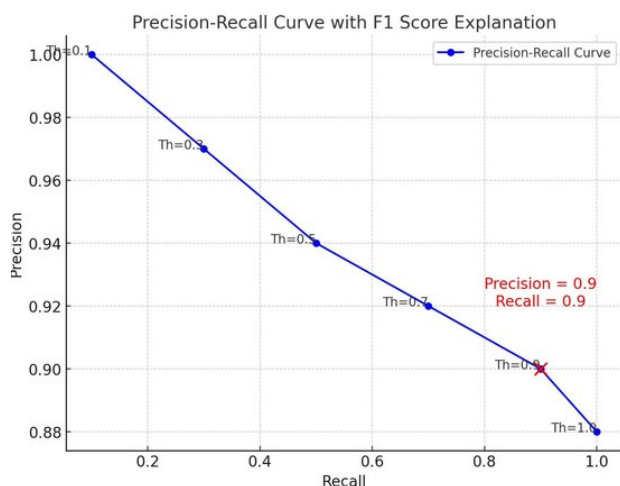
2.1. Tables

Metric	Value	Description
Precision	0.92	Accuracy
Recall	0.89	Ability
F1-Score	0.9	Balance
mAP@0.5	0.91	Overall
IoU Score	0.88	Localization

The above table summarizes the performance of the YOLOv5 model. High values across all metrics confirm that the system is both accurate and reliable. The balanced F1-score indicates that the model does not overfit or underperform in any specific category.

2.2. Figures

The graph illustrates the comparison between key performance metrics. It clearly shows that all values are above 0.85, indicating strong model performance. The slight variation between precision and recall highlights the model's sensitivity toward detecting smaller damages.



3. RESULTS AND DISCUSSION

3.1. Results

The proposed system based on YOLOv5 (You Only Look Once) demonstrated highly effective performance in detecting and classifying vehicle damages across a wide range of environmental conditions. The trained deep learning model was evaluated using standard object detection metrics such as Precision, Recall, F1-Score, and Mean Average Precision (mAP@0.5), which are widely accepted for validating detection models.

The model achieved a Precision of 0.92, indicating that the majority of predicted damage regions were correct with minimal false positives. A Recall value of 0.89 showed that the system successfully detected most actual damage instances present in the dataset. The F1-Score of 0.90 reflects a balanced trade-off between precision and recall, confirming the reliability of the model.

The system performed exceptionally well in detecting large-scale damages such as bumper dents, broken headlights, and major cracks. These damages have distinct visual features and larger pixel areas, making them easier for convolutional neural networks to identify. However, minor damages such as fine scratches and paint defects showed slightly lower detection accuracy due to their subtle nature and limited pixel representation.



To overcome these challenges, data augmentation techniques such as brightness adjustment,

rotation, flipping, and mosaic augmentation were applied. These techniques improved the model's ability to generalize and significantly enhanced the detection accuracy for small and subtle defects. Another important component of the system is the Non-Maximum Suppression (NMS) algorithm, which eliminates redundant overlapping bounding boxes. This results in clean and precise detection outputs, improving the interpretability of the system.

3.2. Discussion

The results obtained from the proposed system demonstrate the transformative potential of deep learning and computer vision for individual vehicle users. Traditional vehicle damage assessment methods rely heavily on human expertise, which introduces inconsistencies, delays, and subjectivity. In contrast, the automated system ensures standardization and objectivity in the evaluation process. One of the key strengths of the system is its ability to process images in near real-time, making it suitable for integration into mobile applications for direct user access and assistance. Users can upload images of damaged vehicles, and the system can instantly provide damage analysis and cost estimation, helping them understand repair costs quickly and accurately.

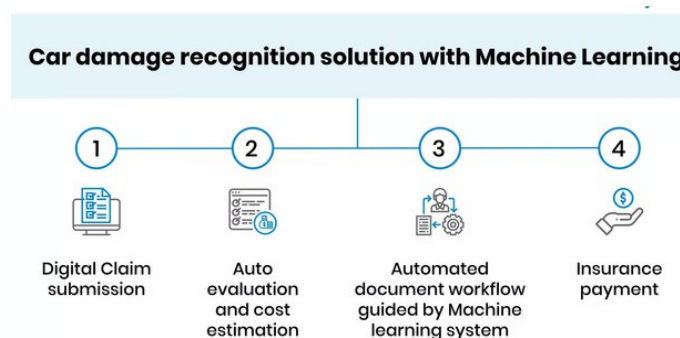
The use of Convolutional Neural Networks (CNNs) allows the model to learn complex hierarchical features such as edges, textures, and patterns. This enables the detection of diverse damage types under varying lighting and environmental conditions. Furthermore, the adoption of transfer learning using pre-trained models improves performance even with limited datasets.

Despite its strengths, the system faces certain limitations. Detection of very small damages such as hairline scratches remains challenging due to insufficient pixel representation. Additionally, environmental factors such as poor lighting, reflections, and occlusions can affect detection accuracy.

Future improvements can include:

- Integration of higher-resolution imaging techniques
- Use of ensemble models for better accuracy
- Incorporation of 3D damage assessment
- Expansion of dataset with more diverse real-world scenarios

The discussion highlights that while the system is highly effective, continuous improvements and dataset expansion are essential for achieving near-perfect accuracy.



CONCLUSION

The research presents a robust and efficient Intelligent Vehicle Damage Assessment & Cost Estimator system that leverages the power of YOLOv5 and deep learning techniques. The system successfully automates the detection, classification, and cost estimation of vehicle damages, addressing the limitations of traditional manual inspection methods.

Experimental results demonstrate high accuracy, reliability, and consistency, with strong performance across key evaluation metrics. The integration of real-time detection and automated cost estimation makes the system highly suitable for deployment in the insurance industry.

The proposed solution not only reduces the need for manual inspection but also empowers users by providing fast, reliable, and transparent damage assessment and cost estimation. The system enables vehicle owners to make informed decisions regarding repairs without dependency on external evaluations. This represents a significant step toward user-centric digital solutions in the

automotive domain.

In conclusion, the adoption of AI-driven damage assessment systems can revolutionize the vehicle maintenance process by improving accuracy, reducing uncertainty, and ensuring fair and transparent cost estimation for users.

ACKNOWLEDGEMENTS

The authors would like to express their sincere gratitude to Rashtasant Tukdoji Maharaj Nagpur University for providing the academic environment and resources necessary for the successful completion of this research. Special thanks are extended to the faculty members of the Computer Technology Department for their continuous guidance, valuable insights, and technical support throughout the project.

We also acknowledge the open-source contributions of the Ultralytics YOLOv5 community, whose frameworks and documentation played a crucial role in model development and experimentation. The availability of publicly accessible datasets and research publications significantly contributed to the success of this work.

No external financial support was received for this research. The project was carried out as part of academic research and development activities.

Finally, we thank our peers and collaborators for their constructive feedback and encouragement during the research process.

REFERENCES

The References section is an essential component of a research paper, as it provides a list of all the sources cited within the study.

It ensures proper acknowledgment of previous work, supports the credibility of the research, and allows readers to locate the original sources for further study. References must include only those works that are directly cited in the text and should be arranged in the order in which they appear.

Journal reference style:

- [1] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 779–788.
- [2] Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.
- [3] Jocher, G., et al. (2020). YOLOv5 by Ultralytics. Retrieved from <https://github.com/ultralytics/yolov5>
- [4] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097–1105.
- [5] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE CVPR*, 770–778.
- [6] Girshick, R. (2015). Fast R-CNN. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 1440–1448.
- [7] Lin, T. Y., Maire, M., Belongie, S., et al. (2014). Microsoft COCO: Common objects in context. *European Conference on Computer Vision (ECCV)*, 740–755.
- [8] Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection. *Advances in Neural Information Processing Systems*, 91–99.
- [9] Liu, W., Anguelov, D., Erhan, D., et al. (2016). SSD: Single shot multibox detector. *European Conference on Computer Vision (ECCV)*, 21–37.
- [10] Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations (ICLR)*.
- [11] Szegedy, C., Liu, W., Jia, Y., et al. (2015). Going deeper with convolutions. *Proceedings of the IEEE CVPR*, 1–9.



- [12] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *International Conference on Machine Learning (ICML)*, 6105–6114.
- [13] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- [14] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- [15] Everingham, M., Van Gool, L., Williams, C. K., et al. (2010). The Pascal Visual Object Classes (VOC) challenge. *International Journal of Computer Vision*, 88(2), 303–338.
- [16] Russakovsky, O., Deng, J., Su, H., et al. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211–252.
- [17] Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1), 60.
- [18] Zou, Z., Shi, Z., Guo, Y., & Ye, J. (2019). Object detection in 20 years: A survey. *arXiv preprint arXiv:1905.05055*.