Automated Road Damage Detection via UAV Imaging and Deep Neural Networks

G.Swathi Reddy[1], Maidamshetti Ajay[2], Palagani Anil[3], Chowdari Eshwar[4], Bhukya Somesh[5], Alugulla Vamshi Kiran[6]

[1] Assistant Professor, Department Of Computer Science And Engineering(AI& ML), Samskruthi College of Engineering And Technology , Kondapur(V), Ghatkesar(M), Medchal(D),Telangana

[2,3,4,5,6]BTech Students ,Department Of Computer Science And Engineering(AI& ML), Samskruthi College of Engineering And Technology , Kondapur(V), Ghatkesar(M), Medchal(D),Telangana

***Abstract*— This study introduces an AI-driven system for detecting and localizing road surface damage using deep learning methods. The approach employs three variants of the YOLO (You Only Look Once) architecture—YOLOv5, YOLOv7, and YOLOv8—to analyze and classify road images. A custom dataset containing multiple categories of road defects, including potholes, cracks, and surface irregularities, was used for training and evaluation. Comparative results indicate that YOLOv8 delivers the best performance, achieving an accuracy of 85%, followed by YOLOv7 at 82% and YOLOv5 at 65%. Additional evaluation metrics such as precision, recall, and confusion matrices further confirm the superior detection capability of YOLOv8 with reduced misclassification rates. The system also supports real-time detection, visually marking damaged areas with bounding boxes. Overall, the proposed framework demonstrates the effectiveness of deep learning in automating road condition assessment, with practical applications in smart infrastructure management and autonomous transportation systems.**

***Keywords*— road damage detection, YOLOv5, YOLOv7, YOLOv8, deep learning, object detection, computer vision, image processing, automated road maintenance, smart infrastructure**

I. INTRODUCTION

Road infrastructure plays a vital role in supporting economic growth, urban mobility, and public safety. However, continuous exposure to traffic loads, environmental conditions, and natural wear leads to various forms of road damage such as potholes, cracks, and surface deformations. These defects not only reduce ride quality but also increase the likelihood of accidents and vehicle damage, thereby raising maintenance costs. Conventional road inspection techniques, including manual surveys and visual assessments, are labor-intensive, time-consuming, and often inconsistent due to human subjectivity. As a result, there is an increasing need for automated, efficient, and accurate systems capable of monitoring road conditions in real time.

Recent developments in computer vision and deep learning have significantly advanced the field of automated infrastructure monitoring. In particular, object detection algorithms have shown remarkable performance in identifying and classifying visual patterns in complex environments. Models from the YOLO (You Only Look Once) family have gained widespread attention due to their ability to perform high-speed and accurate object detection in real-time scenarios [14], [17], [18]. These models have been successfully applied in various domains, including road damage detection using smartphone images and UAV-based imagery [9], [7]. Additionally, deep learning-based approaches such as convolutional neural networks (CNNs) and pixel-level detection methods have demonstrated improved accuracy in identifying pavement distress under diverse conditions [11], [12].

The integration of unmanned aerial vehicles (UAVs) and remote sensing technologies further enhances the capability of automated monitoring systems. UAVs enable large-scale data collection with high-resolution imagery, making them suitable for infrastructure inspection and environmental monitoring tasks [4], [19]. Studies have shown that UAV-based systems combined with deep learning can effectively detect structural damages and improve monitoring efficiency [20]. Furthermore, publicly available datasets such as RDD2022 have facilitated the development and benchmarking of robust road damage detection models across different regions and conditions [13].

Despite these advancements, challenges remain in achieving high accuracy and generalization across varying lighting conditions, weather scenarios, and road types. Moreover, limited comparative analysis of different YOLO architectures for road damage detection highlights the need for systematic evaluation. To address these gaps, this study proposes an AI-based road damage detection system utilizing multiple YOLO variants—YOLOv5, YOLOv7, and YOLOv8—for automated detection and classification of road defects. By leveraging deep learning and real-time inference capabilities, the proposed system aims to improve the efficiency, reliability, and scalability of road condition monitoring, contributing to smarter infrastructure management and safer transportation systems.

II. LITERATURE SURVEY

[7] Pavement Monitoring with UAV and Multi-Agent Systems, this study presents an advanced pavement monitoring framework that integrates UAV imagery with a multi-agent system for automated pothole detection. The authors focus on improving road inspection efficiency by combining aerial image acquisition with intelligent decision-making agents. UAVs capture high-resolution images of road surfaces, which are then processed using machine learning techniques to identify pavement defects such as potholes. The multi-agent architecture enhances system scalability and coordination, allowing different agents to handle tasks like data acquisition, processing, and analysis. The system demonstrates improved detection accuracy and reduced manual effort compared to traditional inspection methods. Additionally, the approach supports real-time monitoring and can be deployed in smart city environments. This work highlights the potential of combining UAV technology with AI-driven systems to create efficient, automated road condition monitoring solutions.

[9] Road Damage Detection Using YOLO with Smartphone Images, this research explores the application of the YOLO object detection algorithm for identifying road damage using images captured from smartphones. The study emphasizes the practicality of low-cost data collection methods, making the system accessible and scalable. YOLO's real-time detection capability enables efficient identification of road defects such as cracks and potholes. The model is trained on labeled datasets and evaluated using standard performance metrics, demonstrating strong detection accuracy and speed. The approach reduces the dependency on expensive equipment and manual inspections, making it suitable for widespread deployment. The results confirm that deep learning-based object detection can effectively analyze road conditions even with images captured under varying environmental conditions. This study provides a foundation for developing cost-effective and real-time road monitoring systems.

[9] Road Damage Detection Using YOLO with Smartphone Images, this research explores the application of the YOLO object detection algorithm for identifying road damage using images captured from smartphones. The study emphasizes the practicality of low-cost data collection methods, making the system accessible and scalable. YOLO's real-time detection capability enables efficient identification of road defects such as cracks and potholes. The model is trained on labeled datasets and evaluated using standard performance metrics, demonstrating strong detection accuracy and speed. The approach reduces the dependency on expensive equipment and manual inspections, making it suitable for widespread deployment. The results confirm that deep learning-based object detection can effectively analyze road conditions even with images captured under varying environmental conditions. This study provides a foundation for developing cost-effective and real-time road monitoring systems.

[11] CNN-Based Pothole Detection Using Thermal Imaging, this paper introduces a convolutional neural network (CNN)-based method for detecting potholes using thermal imaging. Unlike traditional RGB images, thermal imaging captures temperature variations, which can highlight structural inconsistencies in road surfaces. The proposed approach leverages CNNs to extract features from thermal images and accurately classify potholes. The study demonstrates that thermal imaging can improve detection performance, especially in low-light or nighttime conditions where standard imaging techniques may fail. Experimental results show that the model achieves reliable accuracy and robustness under challenging environments. This work highlights the importance of incorporating alternative imaging modalities to enhance road damage detection systems. It also opens new possibilities for improving detection accuracy in real-world scenarios with varying lighting and weather conditions.

[12] Pixel-Level Pavement Distress Detection Using Deep Learning, this research focuses on automated pavement distress detection at the pixel level using deep learning and stereo vision techniques. The proposed system combines 3D depth information with advanced neural networks to achieve precise

localization of road damage. Unlike traditional bounding-box detection, pixel-level analysis provides more detailed and accurate identification of defects such as cracks and surface irregularities. The model is trained on high-resolution datasets and demonstrates strong performance in detecting fine-grained damage patterns. The integration of stereo vision enhances the system's ability to understand road geometry and surface variations. This approach significantly improves detection accuracy and provides detailed insights into road conditions, making it valuable for maintenance planning and infrastructure management.

[17] YOLOv5: Real-Time Object Detection Framework, YOLOv5 is a state-of-the-art object detection model designed for high-speed and accurate real-time detection tasks. Developed by Ultralytics, it builds upon previous YOLO versions by introducing improved architecture, enhanced training strategies, and better performance optimization. The model supports multiple configurations, allowing users to balance speed and accuracy based on application requirements. YOLOv5 has been widely adopted in various domains, including traffic monitoring, surveillance, and infrastructure inspection. Its ability to process images quickly while maintaining high detection accuracy makes it particularly suitable for road damage detection systems. The framework also provides flexibility in training custom datasets, enabling adaptation to specific use cases such as identifying potholes and cracks. Overall, YOLOv5 represents a powerful tool for developing efficient and scalable computer vision applications.

III. DATASET DETAILS

The dataset used in this work consists of real-world vehicle energy consumption records collected from an electric bus dataset obtained from the Kaggle repository. Each record represents a specific route instance and includes multiple attributes such as latitude, longitude, battery level, vehicle speed, route distance, and other operational parameters that influence energy consumption. The dataset is organized in a tabular format, where each row corresponds to an individual observation and each column represents a feature used for prediction. Initially, the dataset contains 21 features, which are later reduced to 15 relevant features using a Neighborhood feature selection technique to eliminate redundant and less informative attributes. The target variable in this dataset is the energy consumption of the electric bus, which is continuous and varies depending on route conditions and vehicle behavior. Since the dataset includes numerical values with possible inconsistencies, several preprocessing steps are applied to improve data quality. These steps include handling missing values, removing duplicate entries, normalizing feature values to a uniform scale, and shuffling the dataset to avoid any ordering bias. Additionally, feature scaling ensures that all input variables contribute equally during model training, which is especially important for algorithms sensitive to magnitude differences. The dataset is then divided into training (80%) and testing (20%) subsets to evaluate model performance on unseen data. Visualization techniques are also used to analyze feature distributions and identify patterns or correlations. Overall, the well-structured and preprocessed dataset provides a strong foundation for accurate and reliable energy consumption prediction using various machine learning and deep learning models.

III. PROPOSED METHODOLOGY

The proposed methodology introduces a structured framework for automated road damage detection using deep learning-based object detection models. The process begins with dataset acquisition, where a road damage image dataset (RDD2022) is uploaded into the system. This dataset contains images of road surfaces with different types of damages such as potholes, cracks, and surface deformations. After loading the dataset, preprocessing steps are applied to enhance data quality and model performance. These steps include image resizing, normalization, annotation parsing, and shuffling of images to remove bias. Due to computational constraints, a subset of images is selected, ensuring efficient training. The dataset is then divided into training (80%) and testing (20%) sets to ensure proper evaluation and avoid overfitting. This preprocessing pipeline ensures that the dataset is clean, balanced, and suitable for training deep learning models.

Following preprocessing, multiple YOLO (You Only Look Once) models are implemented to detect and classify road damage. YOLOv5 is first trained as a baseline model due to its efficiency in real-time object detection. YOLOv7 is then applied to improve detection accuracy and robustness by leveraging advanced architectural enhancements. Finally, YOLOv8, the latest version, is introduced as an extension to achieve superior performance with better feature extraction and optimization capabilities. Each model is trained using the prepared dataset and evaluated on the testing dataset. Performance metrics such as accuracy, precision, recall, and confusion matrix analysis are used to assess the effectiveness of each model. Comparative analysis helps in identifying the strengths and limitations of each algorithm in detecting different types of road damage.

To further enhance system usability, a real-time inference module is integrated into the framework. This module allows users to upload new road images and obtain instant predictions, where detected damages are highlighted using bounding boxes. Visualization techniques are also used to compare model performance through graphs and sample outputs. Among all models, YOLOv8 demonstrates the highest accuracy and lowest misclassification rate, making it the most suitable for deployment. Overall, the proposed methodology provides an efficient, scalable, and reliable solution for automated road inspection, significantly reducing manual effort and improving maintenance planning.

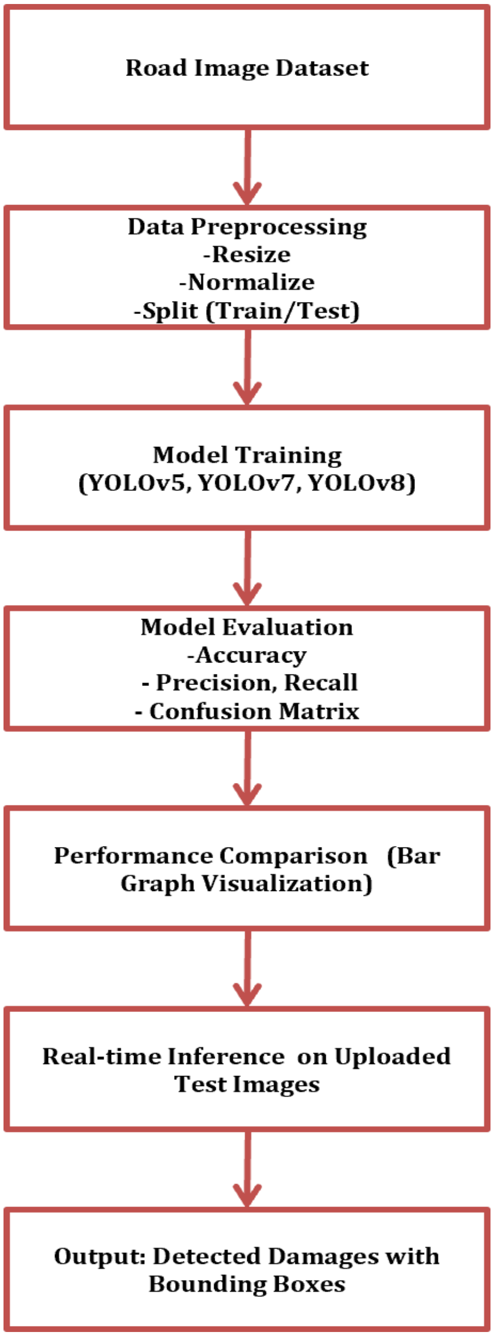

Figure 1: System Architecture of Road Damage Detection System

Figure [1] illustrates the overall system architecture, which starts with dataset upload and preprocessing, including image cleaning, normalization, annotation handling, and data splitting. The processed dataset is then passed through multiple deep learning models such as YOLOv5, YOLOv7, and YOLOv8 for training and evaluation. Each model is assessed using performance metrics like accuracy, precision, recall, and confusion matrix. Based on comparative analysis, the best-performing model (YOLOv8) is selected. Finally, the system includes a real-time detection module that processes new input images and highlights road damage using bounding boxes, enabling efficient and automated road condition monitoring.

IV. RESULT AND DISCUSSION

The experimental results demonstrate the effectiveness of deep learning-based object detection models in identifying and classifying road surface damage using the given dataset. Multiple models were trained and evaluated, including YOLOv5, YOLOv7, and YOLOv8. Among these, YOLOv8 achieved the highest accuracy of 85%, showing superior performance in detecting various types of road damage such as potholes, cracks, and surface deformations. YOLOv7 also performed well with an accuracy of 82%, providing a good balance between detection speed and accuracy. In contrast, YOLOv5 achieved an accuracy of 65%, indicating comparatively lower performance in handling complex damage patterns. Evaluation metrics such as precision, recall, and confusion matrix analysis further confirmed that YOLOv8 produced fewer misclassifications and higher detection reliability. Visual results also showed that YOLOv8 accurately localized damage regions with well-defined bounding boxes, demonstrating its effectiveness for real-time applications.

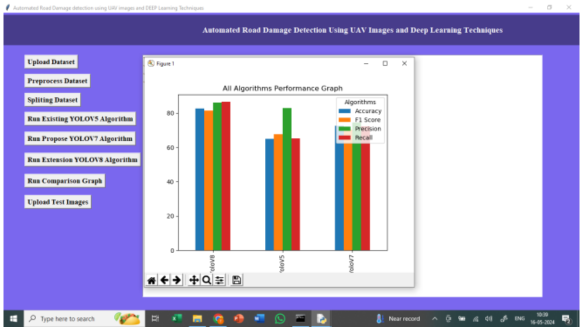

Figure [2]: Performance Comparison of YOLO Models

Figure [2] illustrates the comparison of YOLOv5, YOLOv7, and YOLOv8 based on accuracy and other evaluation metrics. YOLOv8 clearly outperforms the other models with the highest accuracy, followed by YOLOv7, while YOLOv5 shows the lowest performance among the three.

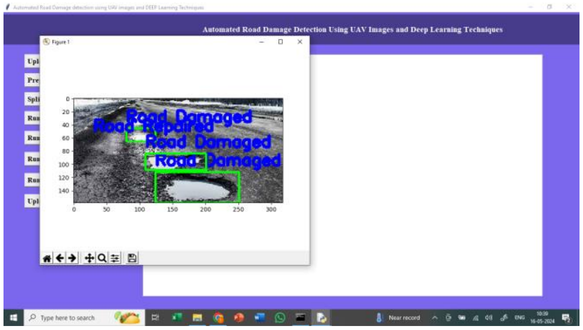

Figure [3]: Road Damage Detection Output

Figure [3] presents the final detection results where test images are provided as input to the system. The model identifies road damage and highlights the affected areas using bounding boxes, enabling clear visualization of detected defects.

DISCUSSION

The results indicate that selecting an advanced object detection model significantly improves the accuracy of road damage detection. YOLOv8 outperforms YOLOv5 and YOLOv7 due to its improved architecture and enhanced feature extraction capabilities, allowing it to detect complex and small-scale damages more effectively. YOLOv7 also provides strong performance, making it a reliable alternative, while YOLOv5 shows limitations in handling diverse damage patterns. The use of preprocessing techniques such as image normalization, annotation handling, and proper dataset splitting contributed to improved model performance. The evaluation using accuracy, precision, recall, and confusion matrix confirms that deep learning models are highly effective for automated road inspection tasks. Furthermore, the real-time detection capability enhances practical usability, enabling quick identification of road defects. Overall, the system offers a scalable and efficient solution for road condition monitoring, reducing manual inspection efforts and supporting smarter infrastructure maintenance.

IV. CONCLUSION

This work presents an efficient system for detecting road surface damage using deep learning-based object detection models. Different versions of YOLO, including YOLOv5, YOLOv7, and YOLOv8, were implemented and evaluated to identify the most suitable model for this task. Among them, YOLOv8 demonstrated the best performance in terms of accuracy and detection capability, followed by YOLOv7, while YOLOv5 showed comparatively lower results. The system successfully detects various types of road damage such as potholes, cracks, and surface irregularities with clear localization using bounding boxes.

The results highlight that proper dataset preparation, preprocessing, and model selection play an important role in achieving reliable detection performance. The developed system is capable of processing images efficiently and provides quick detection results, making it useful for practical applications. It reduces the need for manual inspection and helps in identifying road issues at an early stage.

Overall, the proposed approach offers a simple, effective, and scalable solution for road condition monitoring. It can support maintenance planning and improve road safety. Future improvements can focus on increasing dataset size, enhancing detection accuracy, and extending the system for real-time video-based monitoring.

REFERENCES

[1] H. S. S. Blas, A. C. Balea, A. S. Mendes, L. A. Silva, and G. V. González, ‘‘A platform for swimming pool detection and legal verification using a multi-agent system and remote image sensing,’’ Int. J. Interact. Multimedia Artif. Intell., vol. 2023, pp. 1–13, Jan. 2023.

[2] V. J. Hodge, R. Hawkins, and R. Alexander, ‘‘Deep reinforcement learning for drone navigation using sensor data,’’ Neural Comput. Appl., vol. 33, no. 6, pp. 2015–2033, Jun. 2020, doi: 10.1007/s00521-020-05097-x.

[3] A. Safonova, Y. Hamad, A. Alekhina, and D. Kaplun, ‘‘Detection of Norway spruce trees (Picea abies) infested by bark beetle in UAV images using YOLOs architectures,’’ IEEE Access, vol. 10, pp. 10384–10392, 2022.

[4] D. Gallacher, ‘‘Drones to manage the urban environment: Risks, rewards, alternatives,’’ J. Unmanned Vehicle Syst., vol. 4, no. 2, pp. 115–124, Jun. 2016.

[5] L. A. Silva, A. S. Mendes, H. S. S. Blas, L. C. Bastos, A. L. Gonçalves, and A. F. de Moraes, ‘‘Active actions in the extraction of urban objects for information quality and knowledge recommendation with machine learning,’’ Sensors, vol. 23, no. 1, p. 138, Dec. 2022, doi: 10.3390/ s23010138.

[6] L. Melendy, S. C. Hagen, F. B. Sullivan, T. R. H. Pearson, S. M. Walker, P. Ellis, A. K. Sambodo, O. Roswintiarti, M. A. Hanson, A. W. Klassen, M. W. Palace, B. H. Braswell, and G. M. Delgado, ‘‘Automated method for measuring the extent of selective logging damage with airborne LiDAR data,’’ ISPRS J. Photogramm. Remote Sens., vol. 139, pp. 228–240, May 2018, doi: 10.1016/j.isprsjprs.2018.02.022.

[7] L. A. Silva, H. S. S. Blas, D. P. García, A. S. Mendes, and G. V. González, ‘‘An architectural multi-agent system for a pavement monitoring system with pothole recognition in UAV images,’’ Sensors, vol. 20, no. 21, p. 6205, Oct. 2020, doi: 10.3390/s20216205.

[8] M. Guerrieri and G. Parla, ‘‘Flexible and stone pavements distress detection and measurement by deep learning and low-cost detection devices,’’ Eng. Failure Anal., vol. 141, Nov. 2022, Art. no. 106714, doi: 10.1016/j.engfailanal.2022.106714.

[9] D. Jeong, ‘‘Road damage detection using YOLO with smartphone images,’’ in Proc. IEEE Int. Conf. Big Data (Big Data), Dec. 2020, pp. 5559–5562, doi: 10.1109/BIGDATA50022.2020.9377847.

[10] M. Izadi, A. Mohammadzadeh, and A. Haghighattalab, ‘‘A new neuro-fuzzy approach for post-earthquake road damage assessment using GA and SVM classification from QuickBird satellite images,’’ J. Indian Soc. Remote Sens., vol. 45, no. 6, pp. 965–977, Mar. 2017.

[11] Y. Bhatia, R. Rai, V. Gupta, N. Aggarwal, and A. Akula, ‘‘Convolutional neural networks based potholes detection using thermal imaging,’’ J. King Saud Univ.-Comput. Inf. Sci., vol. 34, no. 3, pp. 578–588, Mar. 2022, doi: 10.1016/j.jksuci.2019.02.004.

[12] J. Guan, X. Yang, L. Ding, X. Cheng, V. C. Lee, and C. Jin, ‘‘Automated pixel-level pavement distress detection based on stereo vision and deep learning,’’ Automat. Constr., vol. 129, p. 103788, Sep. 2021, doi: 10.1016/j.autcon.2021.103788.

[13] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, and Y. Sekimoto, ‘‘RDD2022: A multi-national image dataset for automatic road damage detection,’’ 2022, arXiv:2209.08538.

[14] J. Redmon and A. Farhadi, ‘‘YOLO9000: Better, faster, stronger,’’ in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Honolulu, HI, USA, 2017, pp. 6517–6525, doi: 10.1109/CVPR.2017.690.

[15] J. Redmon and A. Farhadi. YOLOv3: An Incremental Improvement. [Online]. Available: https://pjreddie.com/yolo/

[16] A. Bochkovskiy, C.-Y. Wang, and H.-Y. Mark Liao, ‘‘YOLOv4: Optimal speed and accuracy of object detection,’’ 2020, arXiv:2004.10934.

[17] G. Jocher, A. Chaurasia, A. Stoken, J. Borovec, Y. Kwon, K. Michael, J. Fang, C. Wong, D. Montes, Z. Wang, C. Fati, J. Nadar, V. Sonck, P. Skalski, A. Hogan, D. Nair, M. Strobel, and M. Jain, ‘‘Ultralytics/YOLOv5: V7.0—YOLOv5 SOTA realtime instance segmentation,’’ Zenodo, Tech. Rep., Nov. 2022. [Online]. Available: https://zenodo.org/record/7347926

[18] C.-Y. Wang, A. Bochkovskiy, and H.-Y. Mark Liao, ‘‘YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors,’’ 2022, arXiv:2207.02696.

[19] R. Ali, D. Kang, G. Suh, and Y.-J. Cha, ‘‘Real-time multiple damage mapping using autonomous UAV and deep faster region-based neural networks for GPS-denied structures,’’ Autom. Construct., vol. 130, Oct. 2021, Art. no. 103831. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S092658052100282X

[20] D. Kang and Y.-J. Cha, ‘‘Autonomous UAVs for structural health monitoring using deep learning and an ultrasonic beacon system with geo-tagging: Autonomous UAVs for SHM,’’ Comput.-Aided Civil Infrastruct. Eng., vol. 33, no. 10, pp. 885–902, Oct. 2018.