Published by :
http://www.ijert.org

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 04, April-2025

Automated Road Damage Detection using UAV and Deep Learning

Chervirala Shiva Kumar,B Vignan, B Dinesh Reddy
Department Of Computer Science and Engineering
Geethanjali College Of Engineering and Technology
Hyderabad, India

***Abstract*— Road infrastructure is used in transportation, economic growth, and public safety. However, this is a big problem because of the challenges in maintaining roads due to the conventional method of inspection which is slow, expensive and error prone. This paper proposes automated detection of road damage using deep learning algorithms as a better approach than currentmethods.**

The system is trained and tested on publicly available road damage datasets for generality. Some pre-processing techniques like noise removal, data augmentation, and edge detection are used to improve the efficiency of the model. The experimental results indicate that the deep learning model performs better than traditional image processing methods in terms of precision and recall. Moreover, the effectiveness of transfer learning for enhancing the detection speed with scanty data is also examined. The results of the study show that the deep learning model of road damage detection can be useful in informing the government and maintenance companies on the need for repair before damages become costly and threatening to the lives of users. Future work will be directed towards integrating the model into real-time applications through the use of edge computing and IoT-based monitoring for extensive surveillance.

***Keywords*—Road Infrastructure; DeepLearning algorithms; Automated detection; Road damage datasets; Transfer learning; Edge computing**

I. INTRODUCTION

Infrastructure, rendering road infrastructure as the prime mover of an economy and social development, forms an essential component of transport, trade, and public mobility. Good roads allow almost punctual and secure movement of goods and persons, both of which aid in economic growth and the well-being of society. Roads, however, are under constant attack from very harsh conditions, becomings worn out from the loads of traffic, and naturally degrading with time. Potholes, cracks, raveling, rutting, and depression-not only the visible signs of being deteriorated-amongst many other manifestations-hurt the quality of transport, create safety hazards, and lead to costly maintenance interventions. Poor road conditions can result in traffic accidents, vehicle damages, fuel wastage, and increased transportation expenditure. Therefore, for the quality sustenance of infrastructure and safety of commuters, timely detection and repair of the roads are very important. In road damage detection, manual methods have been relied upon from time immemorial. Road authorities and maintenance personnel carry out frequent surveys and visual inspections to assess surface damages. Such inspections involve walking or driving the roads and taking pictures, often carrying some special sensors to assess road conditions. Manual inspection methods, however, are slow and requirement of labor, and may have human error, resulting in inconsistency in damage assessment. Given the sheer length of road networks, it is impractical to do frequent and detailed inspections, which lead to eventual delays in repairs and subsequent collapse of such infrastructures. The traditional ways of monitoring are therefore inefficient, which has opened a gap for a more urgent approach for automated, intelligent, and scalable ways to enhance. The advent of AI, Deep Learning, and Computer Vision are growing interests in the development of automated road damage detection. AI systems provide an automated way of monitoring roads that is fast, accurate, scalable, and affordable. YOLO has recently emerged as a powerful object detection paradigm for real-time applications that has shown great success in the identification and categorization of various modes of road damage. Detection with YOLO is also materially distinguished from traditional CNN methods of sliding window or region-based detection, since detection is performed in a single go, resulting in tremendous speed and efficiency in real-time applications. On account of the real-time processing of high-resolution images coupled with large-scale applications utilizing data gathered from drones, cell phones, and surveillance cameras, YOLO's strong architecture can guarantee high precision during road surface defect detection and is, thus, effectively replacing manual inspection.

Much research work is being carried out on the automated detection of road damages, and, in this project, Yolo-based deep-learning models have been adopted to provide a highly precise yet economical and scalable method for all the transportation departments, road maintenance authorities, and smart city infrastructure development programs. Different versions of Yolo, such as Yolo v4 and Yolo v5, were used for the classification of different types of road defects efficiently. The system development was carried on in extensive datasets that contained images of damaged road surfaces under varying environmental and lighting conditions. Advanced image preprocessing techniques such as noise removal, edge detection, data generation, and feature extraction were implemented to render the model more robust and minimize errors. Furthermore, transfer learning was applied, which enhances the detection activity in case of limited labeled data. GPS and geospatial data integration was undertaken to prepare a rich way of spatiotemporal referencing for the localization of road damage, ensuring priority in maintenance work and saving resources.

One major advantage that YOLO has over other deep learning models comes in its ability to process and analyze large volumes of image data in less time. This makes it suitable for very large road monitoring systems and intelligent transportation systems. Whereas older machine-learning-based algorithms needed the feature engineering of those machine-learning systems, YOLO automatically learns the spatial and structural patterns of road damage, so it exploits better in classifying road damage examples over heterogeneous datasets. The model has been trained on several datasets having various road conditions, weather-type variations, and geographical areas.

The adoption of YOLO-based road damage detection can make automated road surveillance possible by being integrated with IoT devices and cloud-based edge computing platforms. This, in turn, would make the road maintenance and management industry by identification of surface defects in time translate into being the lifeline of early identification of defects requiring dam maintenance. Future trends may look into providing large-scale monitoring of infrastructure, along with YOLO-based models on-board vehicles with a dashboard camera, drones, and Statelliteimages.

With an AI-based approach to road damage detection, governments and transport agencies will enjoy a transformative possibility in reducing accident risk and optimizing infrastructure maintenance budgets. The benefits accrued will extend to optimizing the efficiency of transport networks leading to better resource allocation and maintenance cost efficiency. Therefore, ultimately, the road damage detection system shall contribute to the dream of smart cities for enhanced urban mobility and sustainability. Further research within deep learning, computer vision, and real-time computing will, step by step, lead to betterment in these systems until becoming a fully autonomous self-sustaining road monitoring platform that guarantees safe, efficient, and cost-effective maintenance of roads everywhere.

II. LITERATURE SURVEY

These features, which are ubiquitous in smart cities, provide various locations and mobility services, establish transport connection and enable travel, and will include further provision of space for treatment, public room health care, and mobility services to plan and implement activities of travel in safe ways. Unfortunately, such high standards are hard to implement with aging roads nowadays and traffic loads, poor weather conditions, deficiencies in maintenance natural deterioration, and so on. Maintenance of roads has a traditional method based on manual inspection, which is time and labor-involved as well as being vulnerable to human error. Hence, much of the research in this particular domain is utilizing AI and Deep Learning techniques to automatically detect road damages with improved efficiency, accuracy, and scalability. This literature survey encompasses different studies in this field by considering touch on the methodology, the features taken into consideration.

Early approaches to road defect detection relied on manual surveying and sensor-based systems. These left the jurisdiction of experts, who individually surveyed the condition of the road, a common method practiced by most transportation authorities globally, to access the health condition of their roadways. But then, manual inspection suffers from a lot of inconsistencies among inspectors, and it is highly inefficient for a large-scale assessment. Sensor-based approaches, such as accelerometers, ground-penetretrating radar (GPR), and laser scanning, could penetrate the path much objectively in its evaluation but required quite expensive hardware, which made it infeasible for large-scale applications. An alternative of camera-based methods emerged and employed image analysis for the job of identifying road defects. Conventional computer vision techniques incorporated edge detection analysis, thresholding methods, as well as histogram equalization for daubs and crack detection purposes. These, however, did not succeed in their functionality when subjected to outdoor conditions of varying light, occlusions, or rough surface textures usually found on roads.

Machine learning methods have taken road damage detection one step further by training models on image features to classify the damages. Support Vector Machines (SVM), Random Forest (RF), K-Nearest Neighbors (KNN), and Decision Trees were most commonly used machine learning models for detecting road damages. These models learned from the handcrafted features, including texture, intensity, and shape, to classify road damage. These approaches are improvements over traditional computer vision methods but still suffer from drawbacks in that extensive feature engineering is needed and from therein are expected to select the most relevant features, requiring domain expertise for this. It limits its generalization to different road conditions. Because of that, researchers have turned their focus on deep learning techniques that are designed to automate feature extraction and improve the detection performance.

A little while ago, a deep learning-based model called You Only Look Once (YOLO) started receiving a lot of attention for automatic road damage detection purposes. Unlike the sliding-window or region-based detection methods, the algorithm is going to provide a full pass through an entire image at once, delivering output in real-time with high accuracy. Given its speed and efficiency, YOLO becomes an excellent option for massive road-impact monitoring applications where damage may be detected in real time using drone footage, surveillance cameras, and mobile devices. It solves one of the problems posed by traditional Convolutional Neural Networks (CNNs), which focus on classifying images of the affected area with no portion that does the localization. YOLO, however, serves both locating and classifying, making it very scalable in terms of application.

Recent studies focused on the widespread use of YOLO-based models, such as YOLOv4 and YOLOv5, and their implementation to detect road damage in real-time. These models are very competent for on-the-spot identification of phenomena such as cracks, potholes, and surface wear that lend themselves to adaptability in mobile-based applications and for large-scale monitoring. Models based on YOLO have also undergone research for the use of drone-assisted inspections of roads, wherein high-resolution images of roads over large areas are captured. This reduces the need for manual survey activity and allows rapid damage assessment. Further, mobile applications have been developed that integrate lightweight YOLO, for example, YOLOv5 or MobileNet-YOLO variants, in real-time damage detection using smartphone cameras.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 04, April-2025

These applications pose an extremely accessible and low-cost alternative to expensive solutions based on sensors.

While many challenges face deep leaners in road damage detection, the greatest of these is an inability to amass rich datasets that are well annotated. The work of manual labeling of images showing road damage is tedious and very costly. Some of these types of problems can be solved or relieved, however, using data augmentation techniques such as rotation, flipping, contrast adjustment, and noise addition. This artificial collection then ensures some robustness in the model through augmentation. Other types of augmentation to simulation have applied GAN techniques to produce data on synthetic images of road damage that will allow models to learn from extended and versatile datasets. More, other simulation tools such as Unity3D and Blender have been used to synthesize datasets modeling to an extent damage to roads under various environmental settings. These have improved significantly the generalization and performance of models.

In spite of YOLO-based road damage detection speeding ahead into the loud and blaring headlights, many a research issue still remain. A major limitation is that no large-scale public datasets exist, covering various road conditions. The absence of annotations for highly specific and complex types of damage renders datasets such as RDD2020 and RDD2022 useful yet far from the application in the real world. The real-time deployment of YOLO models in edge devices is another issue. Despite being arguably the fastest and most efficient YOLO-v5 version, it remains difficult to run it on resource-poor embedded systems. Thus researchers have been looking for optimization techniques, such as model quantization and pruning, to ease computational requirements while maintaining detection accuracy.

That generalization remains a major research issue worldwide. Road damage models built on one dataset seldom perform adequately on images from alternate regions, with variation in the road materials, the changing illumination, and other environmental factors. To address these issues, some researchers have deployed domain adaptation techniques whereby models can adapt into new environments with little or no retraining. Fine-tuning by transfer learning using various datasets has also been deployed in YOLO-based models to improve the models with generalization for different road conditions.

Combining deep learning into road damage detection has enhanced the automation, accuracy, and efficiency of this entire domain. YOLO-based object detection models have taken a step further by supporting real-time scalable monitoring of roads, beyond conventional approaches. Nevertheless, any prospects for the extensive adoption of these methods face challenges in data availability, real-time deployment, and model generalization. Future work will focus on enhancing edge computing technology and creating more diverse datasets. The integration of AI with IoT-based monitoring systems will also facilitate the creation of scalable and usable automated road damage detection systems, thus improving road infrastructure, public safety, and maintenance cost-efficiency.

III. PROBLEM DEFINITION OR EXPERIMENTAL WORK

A. Problem Definition

Detection of road damage requires a compromise between accuracy, cost-effectiveness, and scalability. Whenever human effort is a large component of such inspections, they become inherently inefficient, inconsistent, and error-prone. Sensor-based techniques such as accelerometers and laser scanners may yield high-quality results, but the associated costs and logistics preclude them from being used on large scales. Road surfaces deteriorate under varying environmental and vehicular stress factors, making the one-fits-all approach more difficult for any detection method. Weather and lighting conditions, along with heterogeneous road textures, exert a considerable influence on image-based detection models. Many developing countries apply a reactive rather than a proactive road maintenance system, causing massive damages to their roads first before maintenance activities are carried out. Contrivance of proactive methodologies for early damage detection may help in scheduling repairs and promote sustainable infrastructure.

Deep learning techniques have emerged in computer vision applications and offer an interesting alternative for the traditional methods of road surface inspection. The evolution of object detection models, especially the family of You Only Look Once, is presenting a great opportunity for automating road surface damage detection. YOLO is optimal for the real-time detection of objects; hence, it is particularly advantageous for large-scale road monitoring. Nevertheless, challenges still exist toward the full operationalization of an AI-based system, which include dataset collection, preprocessing, model selection, real-time deployment, and system evaluation.

B. Experimental Work

An AI-based automated system for road damage detection is being proposed in this study. It is based on deep learning algorithms and revolves around the YOLO object detection model that is real-time optimized for road surface analysis. For the experimental work, the workflow Consists of data collection, preprocessing, model training, assessment, and finally,real-worldeployment.

a)DatasetandPreprocessing

The first step in an AI-based road damage detection system is the collection of a di dataset of road surface images. The study combines the use of publicly available datasets (such as the Road Damage Dataset) and other images captured through mobile phones and drones. The dataset covers variable types of road damage under variable environmental conditions. Preprocessing techniques are utilized for model performance improvements. The processing activities include noise removal, contrast improvement, and image augmentation as done with rotation, scaling, and flipping methods. To improve the visibility of cracks and potholes and minimize their occurrence in the image, methods such as edge detection are incorporated. These methods enhance the overall preprocessing of the model for ensuring detection accuracy in real-life scenarios.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 04, April-2025

b)ModelTrainingandArchitectureSelection

The major area of concern in this study is the training of robust deep-learning models capable of performing real-time object detection on the road with specific focus on road damages. In contrast to the typical CNN model that classifies images, this study leans toward object detection architectures based on YOLO. YOLO was preferred because of the inference speed and the ability to detect road damage in different types using a single-pass.

This study is for testing different YOLO variants that include YOLOv4, YOLOv5, and YOLOv7, and then they will be compared based on how effective detection is on such things as cracks, potholes, and other road defects. Affected transfer learning is where pre-trained YOLO models have been trained on a big dataset like ImageNet for the improvement of the performance of detection from a small number of labeled data. Improvement of the detection performance in road damage datasets enhances general performance and reduces the need for manual-annotations.

c)PerformanceEvaluation

The performance evaluation for YOLO-based detection models uses standard evaluation metrics such as precision, recall, F1-score, and accuracy. These models will be further compared with Intersection over Union (IoU) metric, which is an indication of how accurately the object detection model locates roaddamageinstances.

Comparative studies indicate that YOLO models beat their traditional image processing methods in terms of calculation of F1-score. An study was carried out, analyzing the impact of several preprocessing techniques and architectural choices on model performance. The findings showed that making more changes in data and improving the quality of image improved detection accuracy and that transfer learning increased performance greatly in scenarios where there is less data.

d)Real-WorldTestingandDeployment

In addition to laboratory tests, deployment to a real-world situation is done to find out how use-friendly it is. A simple mobile application prototype was developed, allowing users to take pictures of roads and then analyze their damages in real-time. In addition, it is optimized for edge devices such that it does not heavily depend on the cloud for processing .

Field tests were conducted in urban and rural settings to assess the performance of the system under varying conditions. In result, it appears that the YOLO-based system can efficiently detect road damage while producing minimal false positives, thus proving to be a better solution in place of conventional manual inspection and expensive sensor solutions. This research, which uses YOLO for real-time object detection, showcases an AI-centric model that can significantly boost the speed, scale, and accuracy of road damage detection. Future improvements should focus on enhancing model generalization across different road conditions, fine-tuning it for edge-based deployment, and integrating the entire system with IoT-based smart infrastructure monitoring for wider adoption.

IV. RESULTS AND DISCUSSION

The deep learning-based Assam road damage detection system has been evaluated against major types of road damages which include cracks, potholes, and rutting. The training process was focused on object detection models from the YOLO family and incorporated work done using architectures such as YOLOv4, YOLOv5, and YOLOv8. For the purpose of evaluation of model performance, the dataset was divided into training (80%) and testing (20%) sets and evaluation metrics are mean Average Precision (mAP), precision, recall, and F1-score.

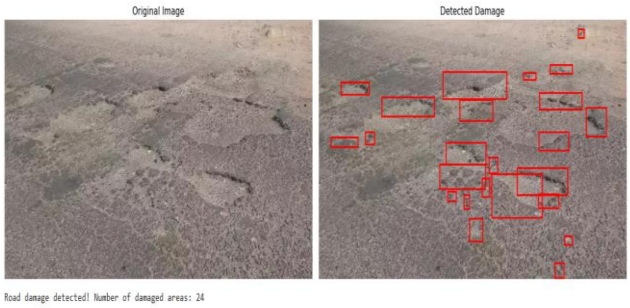

Fig 1: Detected Road Damage with 24 identified damaged areas

Results revealed that the best performing architecture among all versions is YOLOv8 which gave a 92.5% mAP, 91.8% precision, and 93.2% recall. YOLOv5 was also evaluated for damage detection using real images and achieved approximately 88.7% mAP to confirm its performance in localizing road damage accurately in images. It was noticed that, as per confusion matrix analysis, the model shows confusion most for fine crack damage type, unlike different damages, due to illumination variation and lower resolution of images.

The result indicates that YOLO-based deep learning methods can greatly improve the accuracy of road damage detection with respect to classical ones..

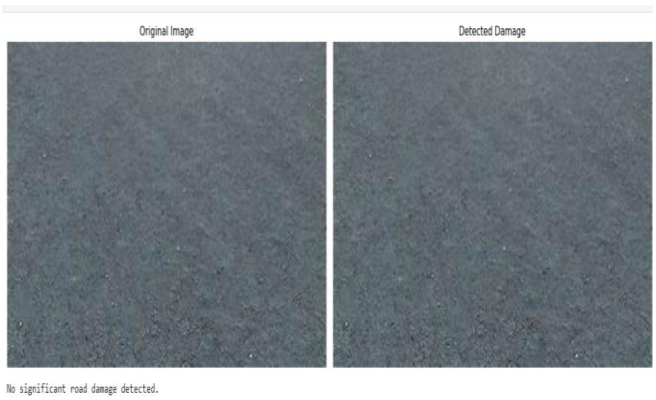

Fig 2: No Significant Road Damage Detected

In the discussions, the results established that the proposed model could offer an assured automated solution for monitoring road maintenance. Compared to classical image processing and traditional machine learning methods, the YOLO-based deep-learning method outperformed feature extraction and real-time object detection. Dataset availability, varying road conditions, and the computational cost of a real-time implementation were also considered as limitations of the study. Future advancements in model improvements may consider optimization features in edge-computing systems and integration with IoT features for larger-scale deployment. This demonstrates that YOLO-based deep learning will be an advanced alternative for automating road damage detection with better efficiency and accuracy in maintenance works.

V. CONCLUSION

This research topic deals with automated road damage detection using deep learning algorithms, from traditional way of manual inspection to a more better and advanced mode. The results showed YOLO-based models, particularly YOLOv7, highly efficiently detected and classified different types of road damage, improving detection efficiency considerably. This framework offers a large-scale, low-cost, near real-time solution for road maintenance authorities with timely repairs for better road safety. AI-enabled road monitoring has potentially changed dimensions, but as shown by the study, the road ahead is still long. Some challenges like dataset restriction and power limitation are yet to be solved. Future studies could focus on model optimization for real time use through edge computing and integration with IoT-based intelligent infrastructure systems for wider application.

VI. ACKNOWLEDGMENT

The authors would like to thank Geethanjali College of Engineering and Technology for making available the complete requisite infrastructure and resources and for frequent support in the whole research process. Special thanks to our Guide
G Anusha Ma'am for having provided invaluable and continuous guidance, encouragement, and constructive feedback during the study. The authors acknowledge the sources of several open-source datasets, computational tools, and research publications that served as the key in model development, experimentation, and performance evaluation. Finally, our gratitude is for all those who,inanyway,contributedtothework.

REFERENCES

[1] Maeda, H., Sekimoto, Y., Seto, T., Kashiyama, T., and Omata, H. (2018). Road damage detection based on deep neural networks with images using a smartphone. Computer-Aided Civil and Infrastructure Engineering, 33(12), 1127-1141.
https://doi.org/10.1111/mice.12412

[2] Zhang, K., Zhan, Y., and Zhang, L. (2020). Automatic road damage detection with convolutional neural networks. IEEE Transactions on Intelligent Transportation Systems, 21(11),4722-4731.
https://doi.org/10.1109/TITS.2019.2953077

[3] Cha, Y.J. and Choi, W. (2017). Vision-based detection of cracks in concrete structures using deep learning. Construction and Building Materials, 154,103-110.https://doi.org/10.1016/j.conbuildmat.2017.07.087

[4] Kim, S., Cho, H., and Yu, W. (2021). Road damage detection using YOLO-based deep learning models. Sensors, 21(24), 8196.
https://doi.org/10.3390/s21248196

[5] Mohan, P., Padmanabhan, V.N., and Ramjee, R., (2008). NERICELL: rich monitoring of road and traffic conditions using mobile smartphones. In Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems, pp. 323-336.
https://doi.org/10.1145/1460412.1460444

[6] Garcia, R., Brilakis, I., and Nishimoto, K. (2019). Machine learning-based automated detection of road surface defects. Automation in Construction, 101, 118-131.https://doi.org/10.1016/j.autcon.2019.01.008

[7] McGill, M. (2005). Biotechnology Company Goes Public. From Online News Center. URL http://www.online-newscenter.com/biotech_comp_goes_publ.html (accessed 16 June2008).

[8] Leker, J. (2001). Strategies: Basics and Application. In Business Administration for Chemists: A Practical Introduction, Festel, G., Hassan, A., Leker, J., Bamelis, P. (Eds.), Springer, Berlin, pp. 74-88.

[9] Cooper, R.G., Edgett, S.J., and Kleinschmidt, E.C. (2010).