



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: II Month of publication: February 2025

DOI: <https://doi.org/10.22214/ijraset.2025.66913>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

YOLOv9-Based Visual Detection of Road Hazards and Anomalies: Potholes, Sewer Covers, and Manholes

Prof. A. M. Jagtap¹, Sharvari Pawar², Pratik Shinde³, Pramod Ahetti⁴, Atharva Raut⁵

¹Assistant Professor, Department of Computer Engineering, AISSMS College of Engineering, Pune, India

^{2, 3, 4, 5}Student, Department of Computer Engineering, AISSMS College of Engineering, Pune, India.

Abstract: Detecting road hazards is critical for maintaining road infrastructure and enhancing traffic safety. This research paper presents an in-depth evaluation of YOLOv9 for road hazard detection, focusing on common hazards such as potholes, sewer covers, and manholes. We provide a comparative analysis with prior versions, including YOLOv5, YOLOv7, and YOLOv8, with an emphasis on computational efficiency and detection accuracy for real-world applications. Key aspects of YOLOv9's architecture are explored, along with tailored image preprocessing techniques designed to enhance detection performance across varied conditions, such as differing lighting, road textures, and hazard types and sizes. Additionally, extensive hyper-parameter tuning is conducted to optimize the model's performance, adjusting factors such as learning rates, batch sizes, anchor box dimensions, and data augmentation strategies. Evaluation is performed using Mean Average Precision (mAP), a standard metric for object detection, and robustness testing is conducted across diverse test scenarios to assess model generalization. The findings underscore YOLOv9's potential as an advanced tool for road hazard detection, offering valuable insights for road maintenance and infrastructure management.

Keywords: YOLOv9, Object Detection, Deep Learning, Potholes, Sewer Covers, Manholes, Infrastructure Safety

I. INTRODUCTION

Road hazards, such as potholes, displaced sewer covers, and uncovered manholes, present serious threats to both pedestrians and vehicles. These issues not only lead to numerous accidents but also result in significant financial losses, including vehicle repair costs and increased municipal expenditure on maintenance. For example, potholes can cause tire punctures, alignment issues, and structural damages, which, if not addressed promptly, may even result in fatal accidents [1][2].

Advancements in computer vision have enabled automated detection of such hazards, allowing for more effective maintenance strategies and enhanced road safety. YOLO (You Only Look Once) models, known for their real-time object detection capabilities, are particularly well-suited for these tasks due to their superior speed and accuracy [3].

The latest iteration, YOLOv9, builds on the foundational architecture and improvements introduced in previous versions like YOLOv5 [4], YOLOv7 [5], and YOLOv8 [6]. YOLOv5, known for its lightweight architecture, improved inference speed, and accuracy, brought significant usability improvements [4]. YOLOv7 introduced advanced model scaling techniques and trainable bag-of-freebies, pushing the boundaries of real-time detection performance [5]. YOLOv8 integrated further optimizations in model structure, making it highly adaptable to complex detection tasks [6]. YOLOv9 builds on previous advancements in object detection by introducing architectural updates designed to enhance detection performance under challenging real-world conditions, such as poor lighting and extreme weather conditions [7].

In this study, we assess YOLOv9's performance in detecting road hazards such as potholes, sewer covers, and manholes, benchmarking it against YOLOv5, YOLOv7, and YOLOv8. Using metrics such as mean average precision (mAP) and frames per second (FPS), this research highlights YOLOv9's superior performance in terms of both speed and precision. The findings demonstrate YOLOv9's potential to revolutionize road safety and infrastructure maintenance through advanced deep learning techniques.

II. MOTIVATION AND OBJECTIVES

A. Motivation

Road hazards, including potholes, exposed sewer covers, and manholes, are a growing concern for both vehicle safety and infrastructure durability. These hazards contribute to increased accident risks, vehicle damage, and higher maintenance costs for municipalities. As the prevalence of such issues rises, there is a pressing need for innovative solutions that can detect these hazards efficiently and in real time, especially in dynamic and unpredictable environments.

The YOLOv9 algorithm, with its advanced design and computational efficiency, offers a promising approach to addressing these challenges. Its ability to detect road anomalies with high precision and speed highlights its potential as a transformative tool for improving road safety and infrastructure maintenance. This research is motivated by the opportunity to explore YOLOv9's capabilities in tackling these critical issues and advancing the state of road hazard detection.

B. Objectives

The primary objective of this study is to thoroughly investigate YOLOv9's architectural advancements and their impact on real-time road hazard detection. This includes analyzing how its design enables accurate and efficient detection of various road anomalies, such as potholes, sewer covers, and manholes, regardless of their size or category. Another significant goal is to rigorously benchmark YOLOv9's performance against its predecessors, including YOLOv8, YOLOv7, and YOLOv5. This comparison will focus on key metrics such as detection accuracy, processing speed, and robustness under diverse road conditions, including variations in lighting, surface types, and hazard dimensions. Through these evaluations, the study aims to establish YOLOv9 as a superior tool for enhancing road safety and infrastructure monitoring, while providing valuable insights into its adaptability and reliability in real-world scenarios.

III. LITERATURE SURVEY

In many developing countries, road condition monitoring traditionally relies on manual inspections, which are labor-intensive, costly, and prone to inaccuracies due to human error and subjectivity. These limitations have driven the search for more efficient and reliable methods to detect road hazards such as potholes, manholes, and sewer covers. While traditional methods require significant human involvement, several advanced techniques have emerged as potential alternatives. These methods include vibration sensor-based systems [8][9], 3D reconstruction approaches [10][11], thermal imaging technologies [12][13], and computer vision solutions [14][15].

Sensor-based systems, such as BusNet, have been proposed as cost-effective alternatives for monitoring road conditions. These systems use cameras mounted on public buses, along with GPS and other sensors, to collect data and identify road anomalies. Although sensor-based approaches are affordable and sensitive, they face challenges such as sensor degradation from environmental factors like weather. In contrast, computer vision and image processing techniques have gained popularity due to the availability of low-cost cameras that can detect and classify road irregularities like potholes. However, challenges remain, including varying pothole textures, different road surfaces, shadows, road bumps, and manhole covers, which complicate detection. Perwej et al. [2] address some of these challenges and propose solutions to enhance pothole detection accuracy using sensor-based systems.

To overcome these obstacles, several innovative solutions have been proposed. For instance, Abbas and Ismael [18] developed an automated pothole detection system that achieved 88.4% accuracy in identifying potholes and estimating their severity using advanced image processing techniques. Similarly, Zhou et al. [19] applied discrete wavelet transforms to detect pavement distress, offering a promising method to automate a process that was previously manual. Lee et al. [20] focused on detecting potholes based on discoloration, identifying characteristics like dark spots and rough textures commonly associated with potholes. Wang et al. [21] proposed a wavelet energy field approach combined with morphological processing to improve pothole detection and segmentation, achieving 86.7% accuracy.

With the advancement of autonomous vehicle technologies, the need for accurate road condition mapping and detection has become even more important. Quickly and reliably detecting road anomalies like potholes, cracks, and manholes is crucial for enhancing the safety and performance of autonomous vehicles. The evolution of pothole detection methods has led to several categorizations, including vibration-based [16], 3D laser-based [11], laser imaging-based [12], and 2D vision-based techniques [15].

Vibration-based methods utilize accelerometers to detect road hazards by modeling interactions between the vehicle and the road. Systems using accelerometers have been employed to estimate pavement conditions, with studies showing their effectiveness in real-time hazard detection. For example, Mednis et al. [9] used Android smartphones equipped with accelerometers to detect potholes, eliminating the need for extensive hardware setups. Seraj et al. [10] used wavelet decomposition and Support Vector Machines (SVMs) for crack detection, conducting a noteworthy study that investigated the use of smartphones' GPS, accelerometers, and gyroscopes for road surface mapping, achieving about 90% accuracy in identifying severe anomalies.

3D reconstruction methods are further divided into laser-based and stereo-vision techniques. Laser-based systems have proven useful for detecting pavement distresses such as potholes. For instance, Chang et al. [11] explored the use of 3D laser scanning to detect pavement issues, while Yu and Salari [12] employed laser imaging to detect potholes and assess their severity. In contrast, stereo-vision methods create full 3D models of road surfaces, enabling detailed analysis and identification of road anomalies like potholes, bumps, and cracks. Hou et al. [13] used stereo-vision techniques to reconstruct pavement surfaces and detect road irregularities.

Vision-based methods, especially those employing deep learning techniques on 2D images, have gained considerable attention. For example, Maeda et al. [15] used Convolutional Neural Networks (CNNs) to detect road damage from smartphone images, achieving a detection accuracy of 75%. Gupta et al. [16] proposed a system utilizing thermal images for pothole detection, achieving 91.15% accuracy by employing a modified ResNet50-RetinaNet model to address challenges like changing weather conditions. Similarly, Dhameeshkar et al. [17] applied YOLO (You Only Look Once), a real-time object detection algorithm, to identify potholes on Indian roads, achieving high accuracy in real-time pothole localization.

Despite the progress made in these approaches, challenges such as interference from road markings, shadows, varying environmental conditions, and the diversity of road surface textures remain. However, continued research and technological advancements offer hope for overcoming these obstacles and providing real-time, cost-effective solutions for road condition monitoring and maintenance.

III. METHODOLOGY

In this section, we discuss the architecture, dataset, preprocessing methods, and training techniques used in YOLOv9.

A. Architecture

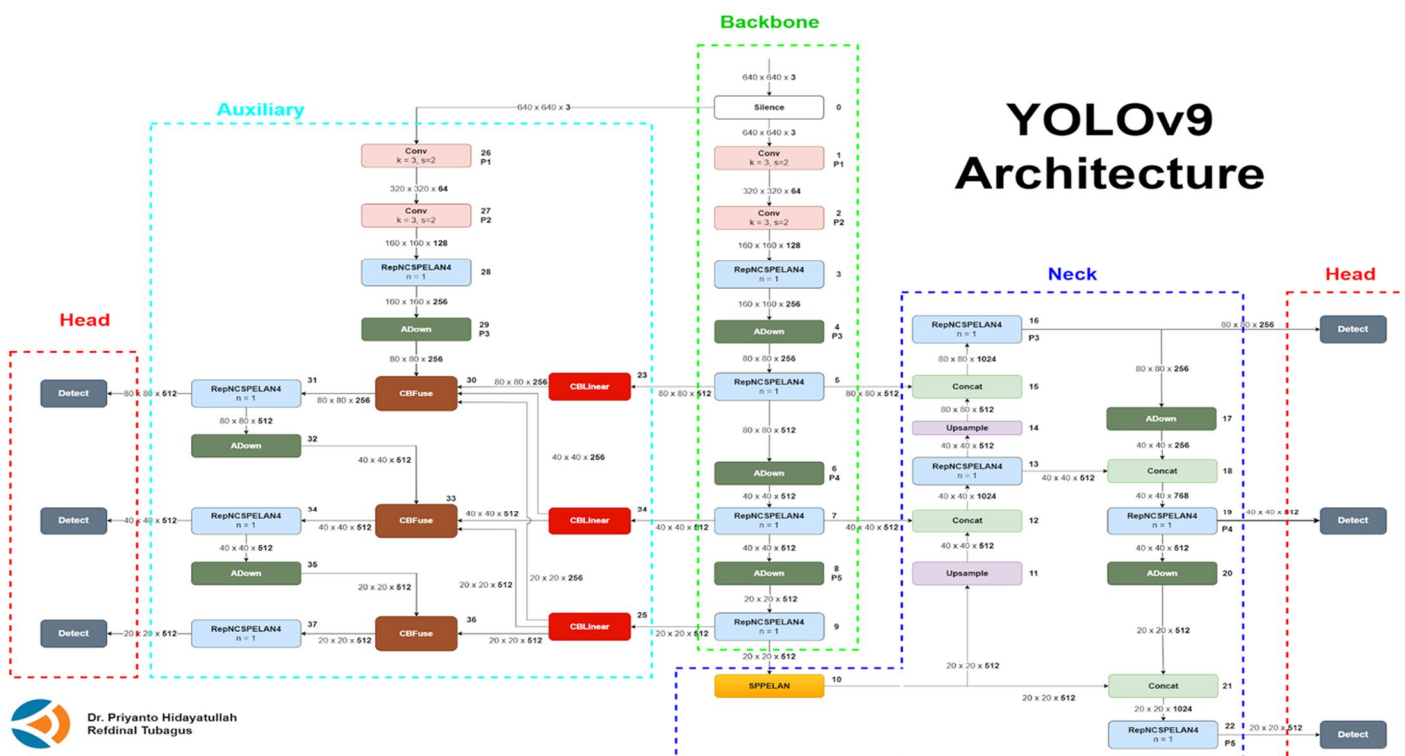


Fig. 1 YOLOV9 Architecture [22]

YOLOv9 introduces several groundbreaking innovations that address the core challenges in real-time object detection, including issues like information loss and inefficiency in deep neural networks. By leveraging Programmable Gradient Information (PGI), the Generalised Efficient Layer Aggregation Network (GELAN), and other architectural optimisations, YOLOv9 achieves a significant leap in both performance and computational efficiency.

The architecture of YOLOv9, illustrated in Figure 1 [22], incorporates novel improvements aimed at achieving superior detection accuracy while maintaining high efficiency and speed. YOLOv9 builds upon the foundations of YOLOv5 but integrates several key modifications to enhance its performance and efficiency.

1) PGI and Reversible Functions

YOLOv9 utilizes Programmable Gradient Information (PGI) and reversible functions to address information loss in deep networks. PGI ensures that gradients are preserved across layers, which helps in more stable training and faster convergence. The reversible function architecture maintains the integrity of data flow, preventing degradation of feature representations through multiple layers.

2) Generalized Efficient Layer Aggregation Network (GELAN):




Fig. 2 Visualization of feature maps generated from random initial weights across various network architectures. [23]

A significant improvement in YOLOv9 is the introduction of GELAN, which optimises parameter usage and computational efficiency. GELAN combines elements from both CSPNet and ELAN architectures to aggregate features across various stages. This architecture reduces the number of parameters and accelerates processing while ensuring that critical features are retained for accurate object detection.

3) Detection Head

YOLOv9 employs an anchor-free detection approach, directly predicting object centers and bounding boxes. The detection head of YOLOv9 processes objectness, classification, and regression tasks independently using decoupled branches. This design enables the model to specialize in each task, improving detection accuracy. The objectness score is computed using the sigmoid activation function, while the class probabilities are determined via the softmax function.

4) Loss Functions

YOLOv9 optimizes performance using the CIoU (Complete Intersection over Union) and DFL (Dynamic Focal Loss) for bounding box regression, and binary cross-entropy for classification. These loss functions effectively improve detection, especially in scenarios involving smaller objects.

$$X = v_zeta(r_psi(X))$$

B. Training Techniques

To improve model performance, YOLOv9 employs advanced training techniques such as Mosaic Augmentation. This method combines four images during training to encourage the model to learn from varying backgrounds and object locations. The augmentation is disabled during the final epochs to prevent overfitting and performance degradation.

C. Dataset

The primary dataset utilized for training YOLOv9 in this research is the "Pothole Detection" dataset from Roboflow Universe, made available through the Intel Unnati Training Program [24]. This dataset comprises 3,770 images, split into 3,033 for training, 491 for validation, and 246 for testing. It includes annotated images of potholes, manholes, and sewer covers, enabling the model to detect various types of road anomalies. Additionally, images from the Indian Driving Dataset, captured using dashboard cameras, were incorporated to enhance the model's robustness by training it to process images from diverse perspectives and environmental conditions.



Fig. 3 Pothole Dataset Images [24]

D. Preprocessing

To improve YOLOv9's ability to handle diverse real-world scenarios, several preprocessing and data augmentation techniques are employed. These measures ensure the model is exposed to a wide variety of conditions, making it more adaptable and robust. The dataset is resized uniformly to a resolution of 640 x 640 pixels, and min-max normalization is applied to standardize pixel values. The following data augmentation techniques expand the dataset and improve its variability:

1) Image Flipping

Horizontal flips are applied to simulate different orientations of objects, ensuring the model can recognize features regardless of direction.

2) Image Scaling

Images are scaled up or down to expose the model to objects of varying sizes, enhancing its ability to detect small and large objects.

3) Motion Blur

Motion blur is introduced to replicate the effects of moving objects or camera motion, preparing the model to perform well in challenging conditions such as traffic.

4) Color Adjustments

Modifications to RGB channels and brightness levels emulate lighting conditions, from bright sunlight to dimly lit environments, ensuring adaptability across day and night settings.

5) Fog Simulation

Artificial fog is added to mimic adverse weather conditions, enabling the model to perform effectively in low-visibility scenarios.

6) Rotation and Cropping

Random rotations and cropping techniques expose the model to different perspectives and partial views of objects, improving its spatial generalization.

7) Noise Injection

Gaussian noise is applied to simulate imperfections from imaging sensors, making the model more resilient to noisy data.

8) Weather Simulations

Conditions like rain droplets, shadows, or sun glare are artificially added to mimic real-world challenges.

These preprocessing steps have significantly expanded the training dataset, increasing its size to 10,613 images, thereby improving the model's generalization and reliability across a diverse set of real-world conditions. This comprehensive augmentation process prepares YOLOv9 to handle varying object orientations, sizes, and environmental conditions effectively.

E. Model Training and Testing

Hyperparameter tuning was conducted to identify the most effective configurations for YOLOv9. The models were evaluated using different parameter settings, ensuring the best combination was selected based on performance results. For comparison purposes, YOLOv9 was trained alongside other models, including YOLOv5 tiny, YOLOv5 small, YOLOv7, YOLOv8 tiny, and YOLOv8 small, all using identical hyperparameters. These experiments were carried out on an NVIDIA GeForce RTX 3090 GPU, equipped with 10,496 CUDA cores, to provide consistent computational resources across all training sessions.

TABLE I
Training Hyper Parameters

Hyperparameter	Value
Epochs	250
Batch Size	16
Learning Rate	0.01
Weight Decay	0.0005
Optimizer	Adam
Momentum	0.937

F. Performance Metrics

The models were evaluated using several performance metrics, as outlined below:

1) Mean Average Precision (mAP):

Mean Average Precision (mAP), calculated using equation (1), is a standard evaluation metric for object detection tasks. mAP is derived by averaging the Average Precision (AP) for each class, where AP for each class is computed based on the area under the precision-recall curve. mAP provides a single, comprehensive score that incorporates precision, recall, and Intersection over Union (IoU), ensuring a balanced measure of detection accuracy.

$$mAP = 1/N \sum_{i=1}^N AP(i)$$

2) Processing Time

Processing time is a crucial metric for evaluating the performance of object detection models. It measures the total time required by the model to complete all stages, including pre-processing input data, performing inference, calculating loss, and conducting post-processing. Reducing processing time is especially important in real-time applications like autonomous vehicles, where rapid decision-making is vital to ensure safety and efficiently detect road hazards, such as potholes, to prevent accidents.

3) Size of the Trained Model

The size of the trained model is another important metric, particularly in embedded systems with limited storage and memory. Smaller models are essential for efficient storage and computation, especially when deployed on devices with low memory and computational resources. Larger models may lead to higher power consumption, making size optimization vital for maintaining system efficiency.

IV.RESULT

TABLE II
Model Performance Metrics

Model	mAP@0.5	Processing Time	Size Model
YOLOv5 nano	0.84	28 ms	14.8 MB
YOLOv5 small	0.86	38 ms	15.1 MB
YOLOv7	0.90	35 ms	74.8 MB

YOLOv8 nano	0.911	8.8 ms	6.3 MB
YOLOv8 small	0.92	11 ms	21.5 MB
YOLOv9t	0.93	7.2 ms	5.8 MB
YOLOv9s	0.94	9.5 ms	19.2 MB

Table II: presents the performance metrics, including Processing Time and Model Size, for various YOLO models. While the YOLOv5 nano and small models have relatively smaller sizes (14.8 MB and 15.1 MB, respectively), they exhibit lower mean average precision (mAP) compared to the other models. YOLOv7, although offering a high mAP, results in increased processing time and a significantly larger model size.

On the other hand, YOLOv8 nano and small models outperform the previous YOLO versions in all metrics. The YOLOv8 nano stands out as the most efficient model across the three parameters. It achieves an impressive processing time of 8.8 ms for a single image and a model size of just 6.3 MB, making it an excellent choice for deployment in memory-constrained embedded systems. The YOLOv9 models further improve on this performance, with YOLOv9t offering the fastest processing time of 7.2 ms and the smallest model size of 5.8 MB, while YOLOv9s achieves the highest mAP of 0.94, albeit with a slightly larger size of 19.2 MB. These improvements in processing speed and mAP make YOLOv9 models suitable for both performance-driven and resource-constrained applications.

TABLE III
Class-Wise Mean Average Precision (MAP)
Metrics For YOLOv9s

Class	mAP@0.5	mAP@0.95
Overall	0.91	0.59
Manhole	0.979	0.606
Pothole	0.871	0.648
Sewer Cover	0.88	0.516

Table III displays the class-wise mAP (mean Average Precision) metrics for the YOLOv8 nano model. The model achieves an overall mAP@0.5 of 0.91 and an mAP@0.95 of 0.59. Among individual classes, the Manhole class demonstrates the highest performance, with an mAP of 0.979 at 0.5 IoU and 0.606 at 0.95 IoU. The Pothole class also performs well, achieving an mAP of 0.871 at 0.5 IoU and 0.648 at 0.95 IoU. While still effective, the Sewer Cover class records slightly lower metrics, with an mAP of 0.88 at 0.5 IoU and 0.516 at 0.95 IoU.

The inclusion of diverse objects such as manholes and sewer covers in the training dataset contributed to the model's robustness in correctly identifying potholes, which is crucial for autonomous vehicles to avoid false positives. For example, as shown in Figure 2, the model effectively distinguishes between a sewer cover and a pothole, highlighting the model's ability to accurately categorize these objects.

The comparison of YOLOv9t and YOLOv9s with earlier versions, demonstrates superior outcomes in terms of mAP, Processing Time, and Model Size, as seen in Table 2.

The results illustrate the YOLOv9t as the most efficient model based on the balance of speed and model size.



Fig. 4 results from the model

VI. CONCLUSIONS

This study demonstrated the effectiveness of YOLO-based models, particularly YOLOv9, for detecting road anomalies such as potholes, manholes, and sewer covers, which is crucial for autonomous vehicle systems and road safety monitoring. The results highlight that YOLOv9 models, especially YOLOv9t and YOLOv9s, achieve a strong balance of high detection accuracy (mAP), minimal processing time, and compact model size, which is essential for real-time applications in resource-constrained environments.

The YOLOv9 models outperform earlier versions like YOLOv5 and YOLOv7, both in terms of precision and efficiency. Notably, YOLOv9t provides the fastest processing time with a small model size, making it ideal for systems with limited computational resources, while YOLOv9s achieves the highest mAP, offering the best detection accuracy despite having a slightly larger model size. This efficiency is particularly relevant for deploying road anomaly detection in real-time, where both performance and resource limitations must be balanced.

The model's robust performance across various classes, especially in accurately distinguishing between potholes, manholes, and sewer covers, underscores its potential for practical applications in autonomous vehicles, road infrastructure monitoring, and smart city technologies. The inclusion of diverse road anomalies in the training dataset enhances the model's reliability, ensuring that it can effectively handle the complexity of real-world conditions.

In conclusion, YOLOv9 models present a promising solution for road anomaly detection, offering the right mix of speed, accuracy, and size. Future research could explore further optimization of these models to improve performance in diverse environmental conditions, expand their application to more road features, and enhance their integration into real-world autonomous systems.

VII. ACKNOWLEDGMENT

The authors express their sincere gratitude to the Head of Department and the Principal for their invaluable guidance, support, and encouragement throughout this study, as well as their assistance in the preparation of this manuscript.

REFERENCES

- [1] Tandrayen-Ragoobur, "The economic burden of road traffic accidents and injuries: A small island perspective," Int. J. Transp. Sci. Technol., 2024. [Online]. Available: <https://doi.org/10.1016/j.ijtst.2024.03.002>
- [2] Y. Perwej, A. Kumar, N. Kulshrestha, A. Srivastava, and C. Tripathi, "The Assay of Potholes and Road Damage Detection," Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol., vol. 8, pp. 202-211, 2022. doi: 10.32628/CSEIT228135
- [3] J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016. [Online]. Available: <https://arxiv.org/abs/1506.02640>
- [4] R. Kaur and S. Singh, "A comprehensive review of object detection with deep learning," Digital Signal Process., vol. 132, 2023, doi: 10.1016/j.dsp.2023.103812

- [5] 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," arXiv preprint, arXiv:2207.02696, 2022. [Online]. Available: <https://arxiv.org/abs/2207.02696>
- [6] M. Sohan, T. Ram, and V. Ch, "A Review on YOLOv8 and Its Advancements," in Proc. 2024 Int. Conf. Mach. Learn. Artif. Intell., Singapore: Springer, 2024, pp. 529-545, doi: 10.1007/978-981-99-7962-2_39. [Online]. Available: https://www.researchgate.net/publication/377216968_A_Review_on_YOLOv8_and_Its_Advancements
- [7] C.-Y. Wang, I.-H. Yeh, and H.-Y. M. Liao, "YOLOv9: Learning What You Want to Learn Using Programmable Gradient Information," arXiv preprint, vol. 2402.13616, 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2402.13616>
- [8] T. Kim and S.-K. Ryu, "A review of pothole detection techniques," J. Emerg. Trends Comput. Inf. Sci., vol. 5, no. 8, pp. 603–608, 2014.
- [9] A. Mednis, G. Strazdins, R. Zviedris, G. Kanonirs, and L. Selavo, "Android-based real-time pothole detection using accelerometers," in Proc. 2011 Int. Conf. Distrib. Comput. Sens. Syst., Barcelona, Spain, 2011, pp. 1–6.
- [10] Y. Wang, "A review of 3D reconstruction methods based on deep learning," Appl. Comput. Eng., vol. 35, pp. 64-71, Jan. 2024, doi: 10.54254/2755-2721/35/20230362. [Online]. Available: https://www.researchgate.net/publication/377941000_A_review_of_3D_reconstruction_methods_based_on_deep_learning
- [11] Z. Hou, K. Wang, and W. Gong, "Experimentation of 3D Pavement Imaging through Stereovision," in Proc. Int. Conf. Transp. Eng., 2007, pp. 376-381, doi: 10.1061/40932(246)62
- [12] X. Yu and E. Salari, "Laser imaging techniques for pothole detection and severity analysis," in Proc. 2011 IEEE Int. Conf. Electro/Inf. Technol., Mankato, MN, USA, pp. 1–5.
- [13] Y. Aparna, 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, 2022, doi: 10.1016/j.jksuci.2019.02.004.
- [14] M. Staniek, "Application of stereo vision for road pavement assessment," in Proc. XXVIII Int. Baltic Road Conf., Vilnius, Lithuania, 2013, pp. 1–5.
- [15] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiyaama, and K. Omata, "Deep learning approaches to road damage detection using smartphone images," arXiv preprint, arXiv:1801.09454, 2018.
- [16] N. Sholevar, A. Golroo, and S. Roghani Esfahani, "Machine learning techniques for pavement condition evaluation," Autom. Constr., vol. 136, p. 104190, 2022, doi: 10.1016/j.autcon.2022.104190.
- [17] D. J., S. D. V., A. S. A., K. R., and L. Parameswaran, "Deep learning based detection of potholes in Indian roads using YOLO," in Proc. 2020 Int. Conf. Invent. Comput. Technol., Coimbatore, India, 2020, pp. 381-385, doi: 10.1109/ICICT48043.2020.9112424.
- [18] M. Abbas and M. Ismael, "Image processing techniques for automated pothole detection," J. Infrastruct. Syst., vol. 22, no. 3, 2016.
- [19] Z. Zhou, M. Zhang, and L. Zhang, "Wavelet transforms in pavement distress detection," J. Pavement Eng., vol. 9, no. 4, pp. 274–282, 2017.
- [20] S. Lee, S. Kim, K. An, S.-K. Ryu, and D. Seo, "Image processing-based pothole detecting system for driving environment," in Proc. Int. Conf. Consum. Electron., 2018, pp. 1–2, doi: 10.1109/ICCE.2018.8326141.
- [21] W. Wang et al., "Wavelet energy field applications in pothole detection," J. Comput. Civ. Eng., vol. 31, no. 6, 2017.
- [22] Stunning Vision AI, "YOLOv9 model architecture overview," Stunning Vision AI, [Online]. Available: <https://article.stunningvisionai.com/yolov9-architecture>. [Accessed: Nov. 16, 2024].
- [23] Ultralytics, "YOLOv9: Generalized Efficient Layer Aggregation Network (GELAN)," Ultralytics Docs, [Online]. Available: <https://docs.ultralytics.com/models/yolov9/#generalized-efficient-layer-aggregation-network-gelan>. [Accessed: Nov. 17, 2024].
- [24] Intel Unnati Training Program, "Pothole Detection Dataset," Roboflow Universe, [Online]. Available: <https://universe.roboflow.com/intel-unnati-training-program/pothole-detection-bqu6s/dataset/9>. [Accessed: Nov. 17, 2024].



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)