Enhancing Road Safety: Real-Time Surface Anomaly Detection Using YOLOv8

Saeeda Varawalla[1] Sonali Bhutad[2], Manjusha Tatiya[4], Smita Bhagwat[3], Chitra Bhole[5] Reshma Pise [6]

1,2 : Department of Computer Engineering, Shah and Anchor Kutchhi Engineering College, Mumbai, India
3: Department of Artificial Intelligence and Data Science, Indira College of Engineering and Management , Pune
4: Department of Electronics and Telecommunication Engineering, Vishwakarma Institute of Technology, Pune
5: Department of Artificial Intelligence and Data Science, K.J.Somaiya Institute of Technology,Mumbai
6: Faculty of Science and Technology, Vishwakarma University , Pune
saeeda.16771@sakec.ac.in, sonali.bhutad@sakec.ac.in , manjusha.tatiya@indiraicem.ac.in, smita.bhagwat@vit.edu, cbhole@somaiya.edu, reshma.pise@vupune.ac.in

KEYWORDS **ABSTRACT**

Current navigation systems lack the capability to provide real-time information about road surface conditions, thereby posing significant safety risks. This study develops a robust system utilizing the YOLOv8 model for detecting road surface anomalies, including potholes, wet surfaces, sewer covers, drain holes, and unpaved roads. YOLOv8 demonstrates distinct advantages over its predecessors, such as enhanced accuracy, reduced inference times (18ms per image, 56 FPS), and an anchor-free detection mechanism, which simplifies training and improves detection precision. The study incorporates this model into a Flask web application, enabling real-time detection and visualization on a map. This integration, combined with persistent storage of detected anomalies along with their GPS coordinates, represents a novel approach to promoting road safety and transparency. Evaluations revealed a mean Average Precision (mAP) of 0.879 at Intersection over Union (IoU) 0.5 and 0.604 at IoU 0.95, validating YOLOv8's superior performance and efficiency.

Index Terms - Real-Time Detection, YOLOv8, Road Surface Anomalies, Road Safety, Deep L, Pothole Detection, Web Application Integration, Sustainable Transportation

I. INTRODUCTION

Road safety remains a critical global concern, with navigation systems often failing to detect real-time road anomalies such as potholes, cracks, and uneven surfaces. On November 24, 2024, a tragic incident in Uttar Pradesh, India, highlighted the urgent need for improved detection capabilities in navigation systems. Three men lost their lives after following GPS directions that led them onto an incomplete bridge, resulting in their vehicle plunging into the Ramganga River. The navigation system had failed to account for the fact that the bridge had been under construction and had collapsed in 2022 due to flooding, and the lack of proper warning signs or barricades further exacerbated the risk. This incident is not an isolated case. In September 2022, a man in North Carolina, United States, died after his GPS directed him to a collapsed bridge that had been out of service since 2013, with no warning or updates provided by the navigation system. Similarly, in 2019, nearly 100 drivers in Colorado became stranded after following GPS directions onto an impassable dirt road, underscoring the limitations of existing navigation technologies in detecting real-time road hazards.

These incidents underscore a fundamental limitation of contemporary navigation systems: their reliance on static or outdated data, which often fails to reflect real-time changes in infrastructure and environmental conditions. While modern navigation applications incorporate traffic updates and, in some cases, user-reported hazards, they generally lack autonomous detection mechanisms that can identify and respond to evolving road conditions. The consequences of this shortcoming range from minor inconveniences to fatal accidents, making it imperative to develop more sophisticated navigation systems with enhanced detection capabilities.

To address this challenge, this study proposes an advanced navigation framework that integrates real-time data acquisition with deep learning-based object detection. Utilizing YOLOv8, the proposed system aims to identify road surface anomalies, construction zones and other hazards. This system can significantly enhance road safety by providing real-time hazard detection, thereby minimizing risks associated with outdated navigation data. The significance of such advancements extends beyond immediate safety benefits, contributing to broader goals of sustainable and intelligent transportation systems. By improving real-time detection and response mechanisms, this approach fosters greater trust in navigation technologies while promoting safer and more efficient mobility solutions.

To contextualize the current landscape of navigation applications, table 1 presents a comparative survey of existing systems with respect to their detection capabilities:

Table 1: App Survey of existing techniques

Navigation System	**Real-Time Traffic Updates**	**Road Hazard Detection**	**Infrastructure Change Detection**	**User Reporting**
Google Maps	Yes	Limited	No	Yes
Waze	Yes	Yes	No	Yes
Apple Maps	Yes	Limited	No	Yes
HERE WeGo	Yes	Limited	No	No
MapQuest	Yes	No	No	No
TomTom	Yes	Yes	No	Yes

As illustrated in Table 1, while most widely used navigation applications provide real-time traffic updates and some level of hazard reporting, none currently possess autonomous real-time detection of infrastructure changes. The proposed system intends to overcome this gap by using an integrated method to navigation, marking a significant step toward the development of safer and more reliable transportation networks.

II. Literature Survey

Several studies have proposed diverse approaches for detecting potholes and road anomalies, emphasizing cost-effectiveness, accuracy, and real-time applicability. Visual input-based methods [7,8,16,21] utilize image processing and deep learning techniques to detect pothole shapes and determine their count. While these methods are cost-effective and rely on 2D visual data, they struggle with measuring pothole depth and volume and are susceptible to lighting and shadow variations. More advanced 3D reconstruction-based techniques [4,11] estimate pothole volume and shape but face implementation challenges due to higher costs and difficulty in detecting anomalies obscured by water or dirt. Vibration-based methods [13,17,23] analyse data from vehicle sensors for pothole presence and depth, providing real-time, low-storage solutions. However, these methods lack the ability to accurately identify pothole shapes and are sensitive to sensor and vehicle variations.[9] compared Faster R-CNN and YOLOv3, finding YOLOv3 better suited for real-time pothole detection due to its higher detection speed, though the dataset lacked variability across weather conditions. Recent advancements in YOLO-based models have further pushed the boundaries of road anomaly detection. For instance, [18] introduced a lightweight YOLOv7-Tiny model enhanced with SimAM attention to improve detection accuracy, though its robustness in extreme lighting conditions remains a limitation. Similarly, [5] achieved computational efficiency by integrating lightweight attention mechanisms but faced challenges in detecting smaller anomalies. [20] enhanced detection on low-resolution images by incorporating generative adversarial networks, albeit at the cost of higher computational demands. [14] proposed a mobile sensing approach for road quality assessment, leveraging smartphone sensors for cost-effective monitoring; however, its accuracy is affected by variations in smartphone hardware, sensor noise, and inconsistent data collection environments. [19] employed mobile crowdsensing with drive recorders to assess road conditions, but the dataset quality is affected by uneven data collection, as it relies on user participation and specific driving routes. [15] integrated CycleGAN with YOLOv5 to detect pavement damages effectively, but its performance in varying weather conditions requires further testing. [6] used CycleGAN-based data augmentation to enhance road surface detection accuracy under varying conditions but the model's reliance on synthetic data may not fully generalize to real-world road anomalies with complex textures and lighting variations. [12] reported that the YOLOX-Nano model achieved higher accuracy in pothole detection while keeping computational costs low. Furthermore, utilizing mobile devices equipped with cameras and GPS sensors has been suggested as an affordable and efficient approach to pothole detection using YOLO algorithms, allowing for community-driven reporting and mapping of road conditions [10].

However, existing pothole detection algorithms are limited by the lack of diversity in the datasets used for training and testing. In this study, through the integration of novel datasets, we established YOLOv8 as an effective approach for real-time road anomaly detection, surpassing its predecessor YOLOv7 and showcasing substantial advancements in computer vision technology.

The key contributions of this study are outlined below:

Curated Dataset: The dataset of road surface images from studies [2,3,25], including potholes, wet surfaces, drain holes, sewer covers and unpaved roads is compiled. This dataset enhances YOLOv8 model training, improving its accuracy in identifying various road surface anomalies.

Innovative Flask-Based Web Application: We engineered a Flask-based web application integrating the YOLOv8 model. Users can capture live photos or upload images for real-time detection of road surface anomalies. Detected anomalies are logged with latitude, longitude and anomaly type in a database and visualized on a map interface for enhanced monitoring of road conditions.

Alignment with Sustainable Development Goals: This project supports Sustainable Development Goal 9 (Industry, Innovation, and Infrastructure) by enhancing resilient and sustainable transportation infrastructure through innovative technological solutions. By facilitating precise detection and proactive management of road surface anomalies, our approach contributes to enhanced road safety and infrastructure resilience, supporting sustainable development initiatives.

III. METHODOLOGY

Dataset

The dataset in this study combines images from multiple sources, including prior studies [2,3] and the Intel Unnati "Pothole Detection" dataset from Roboflow Universe [25]. These datasets were selected to ensure diversity, covering various road conditions such as unpaved roads, wet surfaces, and potholes under different environmental settings. The Intel Unnati dataset provided annotated images of potholes, sewer covers, and drain holes for object detection training.

To maintain consistency, all images were resized to 512×512 pixels and re-annotated using Roboflow. The dataset includes five road surface anomaly categories: potholes, sewer covers, drain holes, unpaved roads, and wet surfaces. It comprises 4,968 images, divided for optimal training: 75% (3,719 images) for training, 16% (778 images) for validation, and 9% (471 images) for testing. This structured split ensures effective model training, validation, and evaluation.

Convolutional Neural Network (CNN)

The CNN architecture employed for the road surface anomaly detection task is a straightforward yet effective design for extracting spatial features from images [1]. Given our objective of real-time road surface detection, we opted against CNNs, which are primarily geared towards image classification. Instead, we chose algorithms better suited for this purpose, such as YOLOv7 and YOLOv8. These algorithms are designed for object detection, providing superior performance in identifying and localizing road surface anomalies in real-time.

You Look Only Once (YOLO)v7

According to [24], YOLOv7 (You Only Look Once version 7) is an advanced object detection algorithm designed for real-time applications. It operates by dividing the input image into a grid and predicting bounding boxes and class probabilities directly from this grid.

Integration Into Web Application

The YOLOv8 model [22], integrated into a Flask-based web application, demonstrates real-time road surface anomaly detection. This application was developed on a local server to illustrate its functionality and potential when incorporated into a larger system. The application was designed with a lightweight architecture, leveraging Flask to enable efficient handling of image uploads, real-time processing, and visualization of detected anomalies. The integration allows the YOLOv8 model to process both uploaded images and live camera feeds, providing immediate feedback to the user. Anomalies detected are stored in an SQLite database, capturing essential information such as anomaly type, latitude, longitude, and timestamp for historical monitoring.

The web application's user interface (UI) was designed to prioritize simplicity and usability, featuring the following components:

1. *Image Uploading:* Users can easily upload images of road surfaces for anomaly detection.
2. *Real-Time Processing:* The application processes live camera feeds or uploaded images, providing real-time detection of road anomalies.
3. *Map Visualization:* Detected anomalies are displayed on a map with precise location markers (latitude and longitude).
4. *Database Integration:* Anomalies detected by the model are stored in an SQLite database, enabling users to review historical data and monitor road surface conditions over time.

Technical Details:

1. *Server Specifications:* The web application was deployed on a local server with the following specifications:
    - **CPU:** Intel Core i7-9700K 3.6 GHz
    - **RAM:** 16 GB DDR4
    - **GPU:** NVIDIA Tesla T4 (for YOLOv8 inference)
    - **Flask Version:** 2.0.1
    - **SQLite Version:** 3.31.1

The Flask framework, along with SQLite for lightweight data storage and map visualizations for real-time feedback, provides an efficient prototype for demonstrating the capabilities.

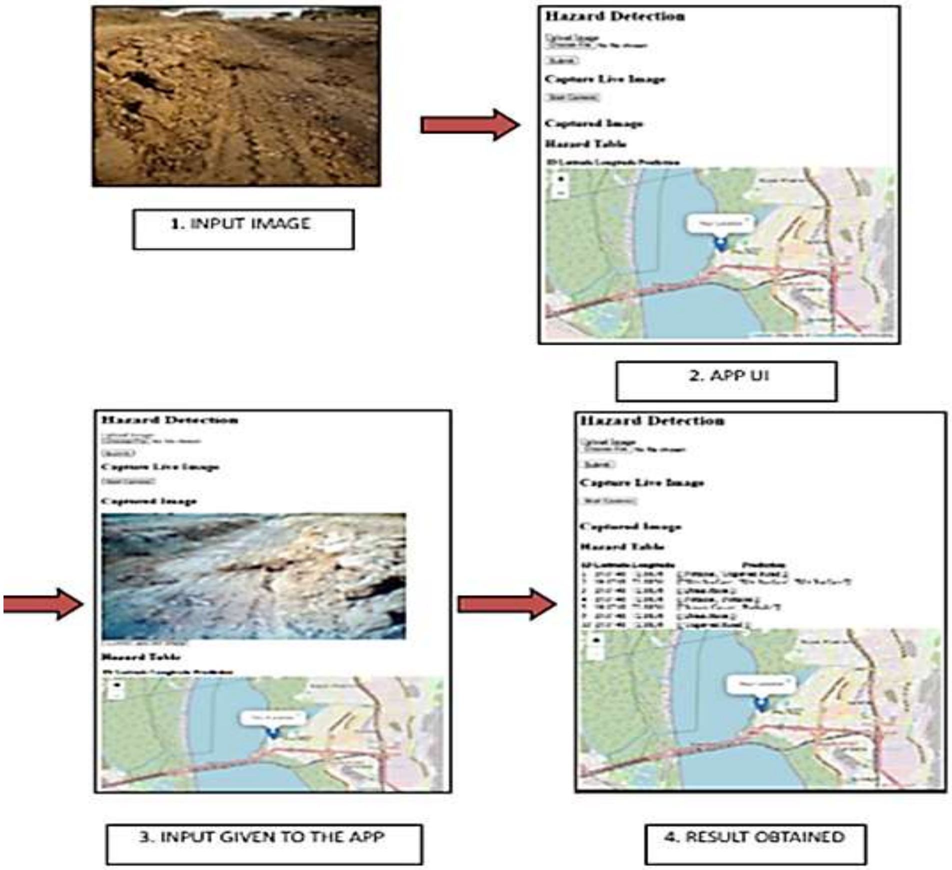

Fig. 1: Web Application Interface and Working

Figure 1 illustrates the user interface and functionality of the web application, showcasing screenshots that depict real-time anomaly detection, map visualization of detected anomalies, and the interface for reviewing historical data stored in the database. Table 2 gives the view of the table obtained in the app after anomaly detection.

Table 2: Results obtained from app

ID	LATITUDE	LONGITUDE	PREDICTION
1	19.0748	72.8856	[['Pothole', 'Unpaved Road']]
2	19.0748	72.8856	[['Wet Surface', 'Wet Surface', 'Wet Surface']]
3	19.0748	72.8856	[['Drain Hole']]
4	19.0748	72.8856	[['Pothole', 'Pothole']]
5	19.0748	72.8856	[['Sewer Cover', 'Pothole']]
6	19.0748	72.8856	[['Drain Hole']]
7	19.0748	72.8856	[['Unpaved Road']]

IV. Results And Discussion

The Convolutional Neural Network (CNN) model was tested on the dataset using Google Colab, running on an NVIDIA T4 GPU with 2,560 CUDA cores and 16 GB GDDR6 memory. The model achieved an accuracy of 58.09%, correctly classifying over half of the cases. Precision, which measures the proportion of correctly identified positive detections, was 48.18%, while recall, indicating the model's ability to detect all actual positives, stood at 29.21%. These results highlight the CNN's limitations in identifying road surface anomalies, emphasizing the need for more advanced models like YOLOv7 and YOLOv8 for improved feature extraction and real-time performance. Figure 2 presents the CNN's classification results.

Fig. 2: Classification by CNN

We compared the inference efficiency of YOLOv7 and YOLOv8 on NVIDIA Tesla T4 GPUs (Kaggle). YOLOv7 had an inference time of 24.9 ms per image, while YOLOv8 performed significantly faster at 11.4 ms, benefiting from architectural improvements that enhance real-time applicability.

Despite strong performance, both models had failure cases. YOLOv8 sometimes misclassified road anomalies under conditions like partial occlusions, small potholes, and low-contrast lighting. It also prioritized larger features, occasionally missing minor cracks. Both models struggled with higher IoU thresholds, where precise localization became challenging. YOLOv8 showed a mean average precision (mAP) of 0.604 at IoU@0.95, outperforming YOLOv7's 0.520, but still highlighting limitations in detecting smaller or less prominent anomalies.

A comparison using normalized confusion matrices (Fig. 3 & 4) showed YOLOv8 outperformed YOLOv7 in most categories. YOLOv8 achieved 82% accuracy for Sewer Covers (vs. 68% for YOLOv7), 95% for Unpaved Roads (vs. 87%), and 72% for Wet Surfaces (vs. 66%). Both models performed equally in Drain Hole classification at 87%. These results confirm YOLOv8's superior detection accuracy, particularly for complex road features. Misclassification rates into the Background category were notably higher for YOLOv7, with significant confusion observed into categories such as Pothole, Sewer Cover, and Wet Surface. This suggests that YOLOv8 exhibits superior feature discrimination capabilities, particularly for background elements.

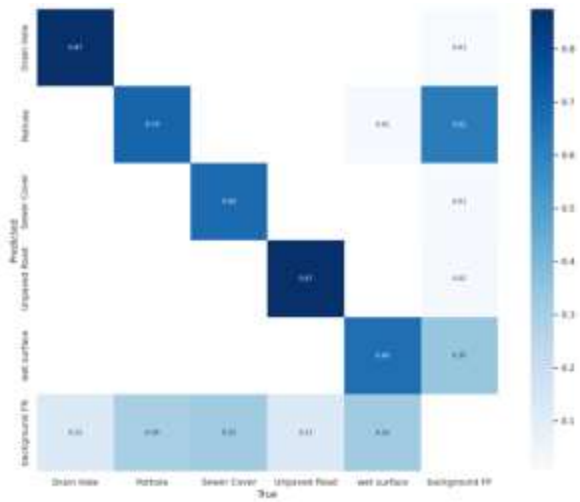

Fig. 3: YOLOv7Confusion Matrix

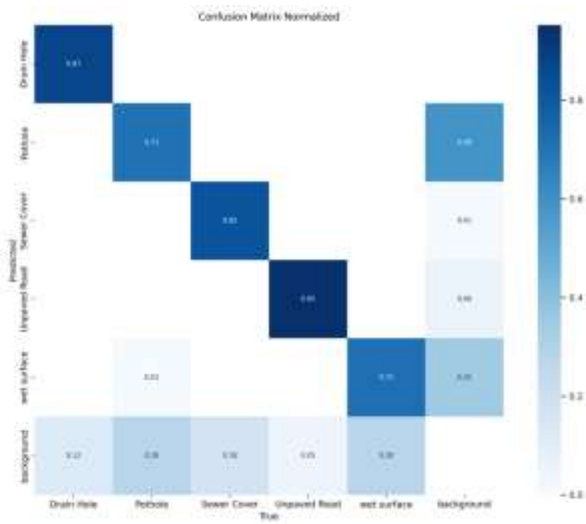

Fig. 4 YOLOv8 Confusion Matrix

Fig. 5: Predictions by YOLOv8

While YOLOv8 excels in detecting road surface anomalies, it has limitations. The model struggles with higher IoU thresholds, where precise bounding box localization is essential, leading to reduced mAP at IoU 0.95, particularly for smaller anomalies in complex environments.

Additionally, despite its faster inference speed, YOLOv8 still presents trade-offs in memory and processing demands, which may impact large-scale real-time applications. Further optimization is needed to enhance localization accuracy and maintain efficiency in more demanding scenarios.

A. Feature Map Visualization for Model Decision-Making

To analyse the YOLOv8 model's decision-making process, feature maps from various layers were visualized to understand how it detects and localizes road surface anomalies. These maps illustrate the model's hierarchical learning, starting with basic feature extraction and advancing to high-level object detection.

In the initial convolutional layer (Layer 0, Fig. 5), the model captures low-level features like edges and textures, forming the foundation for identifying anomalies such as sewer covers. By Layer 19, it begins recognizing more complex structures, highlighting potential road hazards. At intermediate depths (Layer 58), the model refines localization by focusing on contours and shapes, improving differentiation from the surrounding environment. In deeper layers (Layer 82), strong activations indicate precise anomaly detection, confirming the model's effectiveness.

These visualizations (Fig. 6) validate the model's interpretability and reliability, showcasing its ability to progressively refine features for accurate road hazard detection. They also enhance transparency and trust in its real-world applicability.

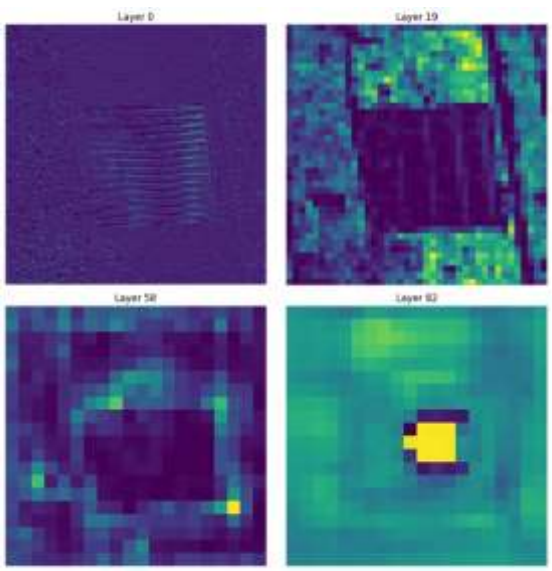

Fig. 6: Visualization of Feature Maps Across Different YOLOv8 Layers

V. FUTURE SCOPE

Future research directions for enhancing the YOLOv8 model in road surface anomaly detection encompass refining its architecture to improve precision and computational efficiency through the pruning of redundant layers. Efforts to adapt the model for varying environmental factors, such as extreme weather or low-light detection, would further enhance its practical applicability, contributing to safer and more efficient infrastructure management.

VI. Conclusion

The deployment of YOLOv8 in a Flask-based web application, capable of providing real-time anomaly detection and visualizing results on a map, offers a user-friendly solution that enhances transparency and facilitates effective road maintenance decision-making.

This paper contributes to the growing body of research exploring the use of deep learning for infrastructure monitoring, particularly in the context of road surface anomaly detection. The combination of YOLOv8's superior performance with real-time applications represents a substantial step forward in automating the inspection and maintenance of transportation infrastructure, reducing human error, and ensuring safer roads. The integration of YOLOv8 into practical systems could transform the way municipalities and organizations approach road safety, enabling more efficient detection and timely interventions.

References

[1] Albawi, S., Bayat, O., Al-Azawi, S., & Ucan, O. N. (2018). Social touch gesture recognition using convolutional neural network. Computational Intelligence and Neuroscience, 2018(1), 6973103.
[2] Bhutad, S., & Patil, K. (2022). Dataset of road surface images with seasons for machine learning applications. Data in brief, 42, 108023.
[3] Bhutad, S., Patil, K., & Khare, N. (2023). Differentiating stagnant water from wet surface for detecting potential mosquito breeding sites in real time. Indian Journal of Science and Technology, 16(5), 331-338.
[4] Chang, K. T., Chang, J. R., & Liu, J. K. (2005). Detection of pavement distresses using 3D laser scanning technology. In Computing in civil engineering (2005) (pp. 1-11).
[5] Chen, J., Yu, X., Li, Q., Wang, W., & He, B. G. (2024). LAG-YOLO: efficient road damage detector via lightweight attention ghost module. Journal of Intelligent Construction, 2(1), 1-10.
[6] Choi, W., Heo, J., & Ahn, C. (2021). Development of road surface detection algorithm using cyclegan-augmented dataset. Sensors, 21(22), 7769.
[7] Dharneeshkar, J., Aniruthan, S. A., Karthika, R., & Parameswaran, L. (2020, February). Deep Learning based Detection of potholes in Indian roads using YOLO. In 2020 international conference on inventive computation technologies (ICICT) (pp. 381-385). IEEE.
[8] Gupta, S., Sharma, P., Sharma, D., Gupta, V., & Sambyal, N. (2020). Detection and localization of potholes in thermal images using deep neural networks. Multimedia tools and applications, 79, 26265-26284.
[9] Gupta, P., & Dixit, M. (2022, December). Image-based road pothole detection using deep learning model. In 2022 14th International Conference on Computational Intelligence and Communication Networks (CICN) (pp. 59-64). IEEE.
[10] Hiremath, R., Malshikare, K., Mahajan, M., & Kulkarni, R. V. (2021). A Smart App for Pothole Detection Using Yolo Model. In ICT Analysis and Applications: Proceedings of ICT4SD 2020, Volume 2 (pp. 155-164). Springer Singapore.
[11] Hou, Z., Wang, K. C., & Gong, W. (2007, July). Experimentation of 3D pavement imaging through stereovision. In International Conference on Transportation Engineering 2007 (pp. 376-381).
[12] KC, S. (2022). Enhanced pothole detection system using YOLOX algorithm. Autonomous Intelligent Systems, 2(1), 22.
[13] Kim, T., & Ryu, S. K. (2014). Review and analysis of pothole detection methods. Journal of Emerging Trends in Computing and Information Sciences, 5(8), 603-608.
[14] Kumar, A., & Kumar, S. (2018). Road quality management using mobile sensing. In International Conference on Innovations in information Embedded and Communication Systems.

[15] Li, Z., Xiao, X., Xie, J., Fan, Y., Wang, W., Chen, G., ... & Wang, T. (2024). Cycle-YOLO: A Efficient and Robust Framework for Pavement Damage Detection. arXiv preprint arXiv:2405.17905.

[16] Maeda, H., Sekimoto, Y., Seto, T., Kashiyama, T., & Omata, H. (2018). Road damage detection using deep neural networks with images captured through a smartphone. arXiv preprint arXiv:1801.09454.

[17] Mednis, A., Strazdins, G., Zviedris, R., Kanonirs, G., & Selavo, L. (2011, June). Real time pothole detection using android smartphones with accelerometers. In 2011 International conference on distributed computing in sensor systems and workshops (DCOSS) (pp. 1-6). IEEE.

[18] Ning, Z., Yang, Y., & Zhou, L. (2023, November). Research on Road Anomaly Detection Algorithm Based on SimAM Attention Mechanism. In 2023 3rd International Conference on Electronic Information Engineering and Computer (EIECT) (pp. 230-234). IEEE.

[19] Piao, B., & Aihara, K. (2017, October). Detecting the road surface condition by using mobile crowdsensing with drive recorder. In 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC) (pp. 1-8). IEEE.

[20] Rout, N. K., Dutta, G., Sinha, V., Dey, A., Mukherjee, S., & Gupta, G. (2023). Improved pothole detection using YOLOv7 and ESRGAN. arXiv preprint arXiv:2401.08588.

[21] Staniek, M. (2013). Stereo vision techniques in the road pavement evaluation. In XXVIII International Baltic Road Conference (pp. 1-5).

[22] Terven, J., Córdova-Esparza, D. M., & Romero-González, J. A. (2023). A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas. Machine learning and knowledge extraction, 5(4), 1680-1716.

[23] Yu, B. X., & Yu, X. (2006). Vibration-based system for pavement condition evaluation. In Applications of advanced technology in transportation (pp. 183-189).

[24] Wang, C. Y., Bochkovskiy, A., & Liao, H. Y. M. (2023). YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 7464-7475).

[25] I. U. T. Program. (2023, August). Pothole detection dataset. Roboflow. Retrieved from https://universe.roboflow.com/intel-unnati-training-program/pothole-detection-bqu6s