Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

A Comprehensive Review on Road Condition Recognition Systems for Intelligent Driving: Deep Learning, Sensor Fusion, and Embedded Implementation Perspectives

[1]Md Musfiqur Rahman, [2]Md Abdullah Al Miraj, [3]Md Mahmudul Hassan, [4]Haque Medha, [5]Hossen Md Fahad

[1]Sadidmrahman@gmail.com

China University of Mining and Technology

Mechanical Engineering

[2]aamsobuj19@gmail.com

China University of Mining and Technology

Mechanical Engineering

[3]hasssantanvir9@gmail.com

China University of Mining and Technology

Mechanical Engineering

[4]marjiamedha1524@gmail.com

Jiangsu Nornal University

Mechanical Engineering

[5]fahad_hossen_cn@outlook.com

China University of Mining and Technology

Safety Engineer

ARTICLE INFO

Received: 01 Nov 2025

Revised: 05 Dec 2025

Accepted: 18 Dec 2025

ABSTRACT

This review summarizes the recent progresses of road-condition recognition systems for intelligent driving, with an emphasis on deep learning, multi-sensor fusion, and embedded deployment on edge devices. First, this work categorizes the methods proposed between the years 2020 and 2025 into the tasks of vision-based road damage detection and global road surface condition classification, pointing out the prevailing architectures, datasets, and evaluation metrics. Then, the performance of YOLOv8 and other state-of-the-art object detectors and segmentation models are explored for real-time identification of potholes, cracks, and surface anomalies by considering accuracy, latency, and robustness across diverse environmental conditions. The subsequent discussion points out the multi-sensor fusion approaches that merge camera data with mmWave radar and temperature and humidity and, finally, the visual texture information to detect hazardous surface states, such as wet roads, snow, and black ice. Following this, different aspects of their embedded implementation are discussed by comparing model complexity, frame rate, response time, and power consumption across various platforms, including NVIDIA Jetson, Raspberry Pi, and other resource-constrained

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

hardware. Limitations in current works are identified concerning the issue of dataset availability and diversity, low-light and adverse weather performance, domain generalization, and sensor calibration and synchronization. Finally, the paper concludes by drawing guidelines toward future research on accurate, cost-effective, and scalable road-condition monitoring frameworks, including lightweight architectures, self-supervised and domain adaptive learning, standardized benchmarks, and deployment-ready designs for large-scale intelligent transportation systems.

Keywords: Road condition recognition; deep learning; sensor fusion; embedded edge deployment; intelligent transportation systems.

Introduction

With advanced driver assistance and route planning, not to mention autonomous driving, intelligent transportation systems are increasingly dependent on road condition recognition. The road surface defects such as potholes, cracks, and rutting degrade ride comfort and contribute to vehicle damage and traffic accidents, particularly in developing regions where maintenance is not so frequent (Carrillo et al., 2020). At the same time, transient hazardous states such as wet surfaces, snow, and black ice reduce tire-road friction, while timely detection is an important contributor to safe vehicle control and speed adaptation. Growing lengths of road networks and increased traffic volume make manual inspection and traditional surveying methods impracticable; hence, a demand arises for automated, scalable, and cost-effective monitoring solutions.

Recent breakthroughs in deep learning and edge computing have empowered a whole new generation of road condition recognition systems to run real-time analyses using inexpensive sensors. The vision-based methods, which rely on either convolutional neural networks or transformer-based architectures, have achieved very good performance in the detection of road damage or in the classification of road surface conditions directly from RGB images or video streams (Huang et al., 2023). Specifically, one-stage object detectors like recent variants of YOLO have allowed for very fast pothole and crack detection, while lightweight models and pruning/quantization techniques have made deployments possible on embedded platforms such as NVIDIA Jetson and Raspberry Pi. Complementing vision, multi-sensor fusion techniques combine millimeter-wave radar, temperature, and humidity sensors, along with auxiliary weather information for robust identification of slippery or low-adhesion surfaces even under low light or visual degradation.

These advances notwithstanding, significant challenges remain to be surmounted before road condition recognition can become a trustworthy component of the intelligent driving system. Most public datasets are small in size, less geographically diverse, and badly annotated, which limits generalization across regions, seasons, and pavement types (Ren et al., 2023).These models also show reduced robustness under low-light and rain-fog-strong glare conditions. Multi-sensor systems have problems coping with calibration and synchronization issues, which degrade the fusion performance on resource-constrained hardware. These gaps provide motivation for a systematic review that synthesizes recent methods but also analyses real-time performance and deployment implications.

Research Question:

(1) How effective are recent deep learning and multi-sensor fusion approaches, including YOLOv8-based models, in achieving accurate and robust road condition recognition for intelligent driving under diverse environmental conditions?

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

(2) To what extent can current road condition recognition models be efficiently deployed on embedded edge platforms, considering constraints on computation, latency, and power consumption, while remaining scalable and cost-effective for large-scale intelligent transportation systems?

Fundamentals of Road Condition Recognition (RCR)

Definitions and problem formulation

It is important to note that RCR handles the automatic estimation of pavement state and surface hazards using sensor data which detects defects, such as potholes, cracks, rutting, or surface states relevant to vehicle safety and comfort. Formally, it could be cast as a supervised learning problem, mapping an input stream x_t from one or more sensors-which may include RGB images, radar returns, or inertial signals, or even tire forces-into a discrete label set, for example, {dry, wet, icy, pothole, debris} or continuous severity scores for each road segment (Liu et al., 2023). Depending on the application, RCR has been implemented as image classification, object detection, semantic/instance segmentation, or sequence modelling over spatiotemporal road segments.

Types of road conditions

Typical RCR taxonomies distinguish:

- Global surface states: dry, wet, snow, slush, icy/black ice and contaminated (mud, waterfilled potholes).
- Localised defects and hazards: cracks including longitudinal, transverse and alligator, potholes, rutting, patching, loose debris, manholes, and edge breaks.
- Such datasets as RDD2019/2022 and Cracks-and-Potholes-in-Road-Images annotate multiple types of cracks and potholes, while road-surface recognition datasets include such categories as dry/wet/icy to enable friction-aware driving and winter maintenance.

Challenges and constraints in real-world RCR

Real-world RCR needs to handle high intra-class variability due to different pavements, aging, and various materials. The nuisance factors include changes in illumination, shadows, occlusions, and camera motion. Degradation of visual cues due to adverse weather conditions such as rain, fog, or snow, nighttime conditions, and contamination resulting from water or dirt worsens the situation (Liu et al., 2023). Vibrations and variations in mounting height further hamper inertial and depth measurements. Moreover, in practical deployments, computing resources, bandwidth, and energy on vehicles or smartphones are limited. Very few large, diverse datasets with consistent annotations are available for training robust models.

Traditional approaches to road condition detection

Vision-based classical methods

Early methods using computer vision in RCR have focused on handcrafted features and low-level image processing, such as edge detection, thresholding, morphological operations, and texture descriptors like Gabor filters, LBP, HOG, by segmenting cracks or potholes from pavement backgrounds (Dasari et al., 2024). These methods often assume uniform lighting and pavement texture to a certain extent and conduct rule-based post-processing to filter out noise and estimate crack length or pothole area.

Statistical and rule-based techniques

Apart from pure image processing, other traditional approaches rely on statistical models like SVMs, random forests, or k-means clustering applied to features derived through feature engineering from

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

images, laser profiles, or accelerometer signals. Other threshold-based rules have also been developed for the purpose of vibration signatures, elevation profiles, or gray-level statistics estimation of potholes and roughness in smartphone- or vehicle-mounted inertial sensing systems.

Limitations of classical approaches

These methods are very sensitive to illumination, shadows, surface markings, and camera pose, all of which contribute to high false-positive and false-negative rates whenever real-world conditions differ from the assumptions made during design. Additionally, generalization across different types of pavements, locations, and weather conditions is poor, often requiring labour-intensive manual tuning of parameters, and hence these are unsuitable for large-scale heterogeneous deployments (Zhong et al., 2025).

Deep learning for road condition recognition

CNN-based methods

Convolutional Neural Networks are dominating recent research on RCR, with VGG, ResNet, U-Net, and YOLO variants for the classification, detection, and segmentation of road damage. Commonly, object detectors like YOLOv5/YOLOv8, Faster R-CNN, and SSD are trained on road damage datasets such as RDD2019/2022, Cracks-and-Potholes, and Kaggle road damage sets to localize and categorize each pothole and crack in real time (Gadekar et al., 2025).

Figure 1: YOLOv8 Pothole/Crack Detection Example

Transformer and attention-based models

More recent works integrate attention mechanisms and transformer-based backbones to better capture long-range context and subtle texture variations. Vision transformers and hybrid CNN–transformer networks have been applied to road surface classification and crack segmentation, improving robustness to complex backgrounds and variable scales by adaptively focusing on informative regions.

Multi-task learning approaches

Multi-task learning frameworks jointly optimize related tasks such as defect detection, severity estimation, and surface-type classification, sharing features across tasks to improve sample efficiency

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/ **Research Article**

(Dasari et al., 2024). For example, some systems combine segmentation with depth or damage severity regression, enabling both localization and quantification of pavement distress from the same model.

Comparative analysis of deep learning models

Comparisons typically report mean Average Precision (mAP), precision–recall metrics, and inference speed (FPS) across detectors on standard benchmarks like RDD2022 and Cracks-and-Potholes-in-Road-Images. Lightweight YOLO-based models on embedded GPUs often achieve a trade-off where minor accuracy loss relative to heavier backbones is offset by substantial gains in FPS and energy efficiency, especially on Jetson-class hardware (Gadekar et al., 2025).

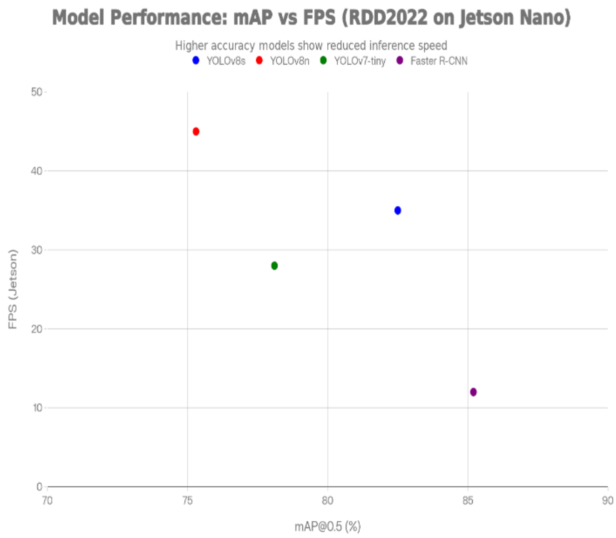

Figure 2: Model Performance Comparison (mAP vs FPS)

Datasets and benchmarking techniques

Representative RCR datasets

Dataset / Source	**Target conditions / classes**	**Modality**	**Scale (approx.)**	**Notes**
RDD2019 / RDD2022	Multiple crack types, potholes	RGB images	5k–10k+ images	Widely used for road damage detection with YOLO/Faster R-CNN

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

Dataset / Source	Target conditions / classes	Modality	Scale (approx.)	Notes
Cracks-and-Potholes-in-Road-Images	Cracks, potholes, road	RGB images	2,235 images, 4,720 objects	Designed for object detection with three classes
Road Damage Dataset (Kaggle, variants)	Potholes, cracks, manholes	RGB images	Several thousand images	Collected in varying urban and rural regions
Road Surface Recognition Dataset	Dry, wet, snowy/icy, others	RGB / weather	Thousands of samples	Focus on global surface condition classification
Multi-modal driving datasets (e.g., ZOD, RSXD)	Rich autonomous driving scenes	Camera, LiDAR, radar	Hundreds of sequences	Useful for sensor fusion, but may lack explicit damage labels

Sensor fusion approaches

Vision + LiDAR fusion

Vision–LiDAR fusion exploits camera texture and LiDAR geometry to identify surface anomalies and estimate depth or severity of potholes and rutting. Typical pipelines project LiDAR point clouds onto image planes, using point-wise features or bird's-eye-view grids fused with image features to detect elevation changes consistent with defects.

Vision + radar fusion

Camera–radar fusion uses mmWave or microwave radar to sense surface reflectivity and roughness, which correlates with friction and road condition, especially under low visibility. Fusion architectures either augment image-based networks with radar feature maps or incorporate radar-based surface classification outputs as additional inputs for road condition recognition (Liu et al., 2023).

Tire-force and vehicle dynamics-based sensors

Some systems detect the road condition by analysing tire-road interaction signals, such as wheel slip, longitudinal/lateral acceleration, yaw rate, and suspension responses. Machine learning models trained with these signals classify low-friction states, such as wet or icy, and estimate surface roughness, often in combination with environmental measurements of temperature and humidity (Liu et al., 2023).

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

Multi-modal deep learning architectures

In multi-modal deep learning frameworks, image and LiDAR, radar, inertial, and weather data are combined using separate encoder branches and fusion layers. Architectures vary from the simple concatenation of modality-specific features to advanced attention-based fusion blocks that weight modalities according to reliability and context, for example, trusting radar more when there is fog.

Fusion strategies: early, late, and hybrid fusion

Early fusion means early network layer merging of raw or low-level features from different sensors; hence it maximizes joint representation learning but demands tight calibration and synchronization (Liu et al., 2023). In contrast, late fusion sums high-level decisions or class probabilities from modality-specific networks; thus, it eases integration at the cost of some performance. Hybrid fusion combines both such strategies (for example, mid-level feature fusion with decision refinement) for improving robustness.

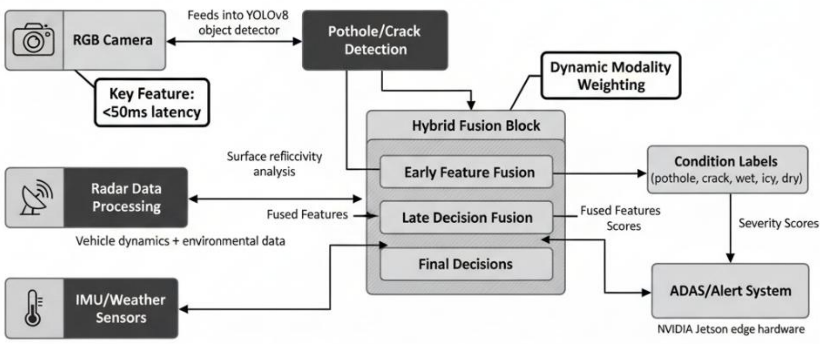

Figure 3: System Architecture Diagram

Embedded and real-time implementation

Requirements for onboard automotive systems

This requires onboard systems to operate under strict constraints of tens of milliseconds of latency, reliability, and functional safety within the limited computational budgets and thermal envelopes. The automotive-grade implementations have to be robust to vibrations, to a wide temperature range, and also should be able to interface with vehicle networks such as CAN and Ethernet along with existing ADAS stacks.

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

Hardware platforms (NVIDIA Jetson, FPGA, ARM, DSP, etc.)

Typical embedded platforms include NVIDIA Jetson Nano/Xavier/Orin for GPU-accelerated deep learning, ARM-based SoC boards like Raspberry Pi, and specific FPGA/DSP platforms for low-latency signal processing (Zhong et al., 2025). The literature indicates that optimized CNN/YOLO models achieve real-time performance at 20-60 FPS on Jetson devices, while Raspberry Pi often needs smaller models, accelerators (e.g., Edge TPU), or reduced frame rates.

Model compression and acceleration

Quantization-for example, INT8-reduces the model size and increases the inference speed for edge deployment. Pruning and knowledge distillation are also widely used techniques. With proper care and combination with architecture search for a lightweight backbone, they can reduce memory footprints and power usage significantly with modest accuracy degradation.

Energy efficiency and latency constraints

Automotive roadside units are required to operate continuously with constrained power budgets in deployments that are either battery-powered or solar-backed. Thereby, evaluation considers not only the mAP and F1 but also the per-frame energy consumption, throughput Aka FPS, and end-to-end response time, often measured under realistic driving workloads (Arya et al., 2021).

Deployment challenges in real vehicles

Real-vehicle deployment involves issues related to sensor mounting variability, calibration drift, network delays, and adherence to safety and cybersecurity regulations. Therefore, for the purpose of safe operation in case sensor failures occur, or extreme condition changes and model predictions with high uncertainty, over-the-air updates, fault detection, and fallback strategies will be necessary.

Evaluation metrics and experimental protocols

Accuracy, precision, recall, F1

Accuracy, Precision, Recall, F1-score are standard metrics applied to the tasks of global surface-state recognition and binary damage detection. On object detection, mAP at different IoU thresholds, per-class precision-recall curves, and confusion matrices give a more fine-grained view of performance regarding cracks, potholes, and other damage types (Gadekar et al., 2025).

Latency and real-time performance

Latency here includes per-frame inference time, the end-to-end processing delay from sensor capture to output, and throughput in frames per second. In general, real-time performance is defined with respect to requirements at the application (e.g., ≥30 FPS for onboard detection) and will be directly evaluated on target hardware rather than desktop GPUs (Zhong et al., 2025

Robustness to weather, lighting, and occlusion

It is also considered in terms of robustness to weather conditions, such as clear, rain, and snow, during daytime, dusk, or night, by stratified subsets, and/or domain shift experiments. Some augment training by adding synthesized weather conditions or using domain adaptation techniques that measure degradation under shifted conditions in mAP/F1.

Simulation vs real-world testing

These allow controlled experiments, generation of rare events, and safety validation for extreme cases; simulation platforms and synthetic datasets typically have domain gaps relative to real roads (Arya et

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

al., 2024). Best practice therefore combines simulation-based development with extensive on-road testing, including cross-region and cross-season evaluations.

Case studies and real-world applications

Autonomous driving systems

Autonomous vehicles adjust motion planning, speed profiles, and braking strategies based on road condition estimates, mainly in low-friction scenarios. The integration of RCR with perception and planning enhances safety margins due to the adaptation of following distances and lane-change behaviour to the state of the surface and detected damage (Arya et al., 2021).

Advanced Driver Assistance Systems (ADAS)

These include warning drivers of impending potholes, suggesting speed reductions on slippery segments, and enabling adaptive cruise control and emergency braking. RCR applications based on smartphones or retrofit devices for crowdsourcing of damage locations and surface states complement OEM-mounted sensors (Liu et al., 2023).

Smart cities and intelligent transportation systems

These RCR outputs are used by municipalities and highway agencies for prioritizing maintenance, optimizing inspection schedules, and supporting pavement management decisions. The data can also feed dynamic speed limit, route guidance, and maintenance planning dashboards when integrated into ITS infrastructure at a city scale (Arya et al., 2021).

Current limitations and open research challenges

Generalization across environmental conditions

Models tend to perform worse when deployed in regions, seasons, or pavement types not well represented in the training data, which indicates limited domain generalization. The solution lies in more diverse datasets and handling methods of domain shift, including domain adaptation and robust feature learning.

Real-time multimodal fusion constraints

Real-time fusion of high-bandwidth modalities is computationally burdensome and sensitive to calibration and synchronization errors. The approach would require developing efficient architectures and sensor scheduling strategies that keep latency and power consumption within onboard limits with bounded loss of accuracy.

Lack of diverse and annotated datasets

Geographical coverage, variability in weather conditions, and richness of the annotation such as severity labels and friction estimates remain limited. More comprehensive benchmarking thus remains limited. Aligned multimodal datasets with camera-LiDAR-radar-vehicle dynamics signals, specifically labelled for RCR tasks, are still rare (Liu et al., 2023).

Safety and reliability considerations

RCR systems need to cope with uncertainty, sensor failures, and adversarial conditions, where misclassifications may have unsafe decisions as consequences. Formal verification, uncertainty estimation, and safety-aware design remain largely unexplored compared to model optimization for accuracy.

Future directions

Self-supervised and few-shot learning

The self-supervised and contrastive learning methods will try to exploit the large volumes of unlabelled driving data so as to learn robust representations with reduced reliance on expensive manual annotation. Few-shot and meta-learning could enable fast adaptation to new regions or pavement types with few labelled examples.

Edge AI and ultra-efficient embedded solutions

The future RCR systems will increasingly leverage ultra-lightweight architectures, which are co-designed with hardware accelerators, using techniques such as neural architecture search, aggressive quantization, and sparse computation (Huang et al., 2023). These edge AI solutions will support continuous, vehicle-wide deployment and large-scale crowdsourced monitoring without overwhelming communication and energy budgets.

V2X integration (Vehicle-to-Everything)

Integrating RCR outputs into V2X communication will provide the capability for vehicles and infrastructure to share surface condition information, enabling cooperative hazard warnings and coordinated speed management. It requires standardized message formats and appropriate security mechanisms to integrate RCR into existing V2X ecosystems (Huang et al., 2023).

Federated learning and privacy-preserving systems

Federated learning will enable training global RCR models across fleets of vehicles without centralizing raw sensor data, thereby improving performance while preserving privacy. In particular, combining federated optimization with differential privacy and secure aggregation will be crucial for large-scale RCR services deployed according to regulations.

Conclusion

It has evolved from classical image and statistical approaches to sophisticated deep learning architectures and multi-sensor fusion techniques for real-time detection of potholes, cracks, and hazardous surface states, such as wet roads and black ice. YOLOv8 and other models prove strong performance on benchmarks like RDD2022 and Cracks-and-Potholes datasets, while embedded deployments on Jetson and Raspberry Pi platforms achieve viable FPS and latency within resource constraints. Fusion strategies-early, late, and hybrid-promote further robustness by incorporating vision with radar, LiDAR, and vehicle dynamics data to address the limitations imposed by a single-modality approach under adverse weather conditions and low light.

Key Insights

This review highlights that with the improvement in the accuracy metrics of mAP and F1, challenges still persist in generalization, dataset diversity, and safety-critical reliability for ITSs. Embedded implementations strike a balance between speed and efficiency through quantization and pruning but multimodal synchronization and energy constraints remain bottlenecks towards scalable ADAS and autonomous driving applications.

Outlook

Future developments should take into consideration self-supervised learning, V2X integration, and federated frameworks to provide cost-effective privacy-aware solutions deployable across different global road networks. These directions promise safety and intelligent transportation by bridging existing robustness and scalability gaps.

References

1. Arya, D., Maeda, H., Ghosh, S. K., Toshniwal, D., Mraz, A., Kashiyama, T., & Sekimoto, Y. (2021). Deep learning-based road damage detection and classification for multiple countries. *Automation in Construction*, 132, 103935. https://doi.org/10.1016/j.autcon.2021.103935
2. Arya, D., Maeda, H., Ghosh, S. K., Toshniwal, D., & Sekimoto, Y. (2024). RDD2022: A multi-national image dataset for automatic road damage detection. *Geoscience Data Journal*, 11(4), e260. https://doi.org/10.1002/gdj3.260
3. Gadekar, P., Tajne, N., Kadu, A., Shinde, T., & Sadaphal, S. (2025). Pothole and road hump detection using YOLO model. *International Journal of Scientific Research and Engineering Development*, 8(2), IJSRED-V8I2P171. https://ijsred.com/volume8/issue2/IJSRED-V8I2P171.pdf
4. Liu, Y., Liu, K., Mi, J., & Luo, M. (2023). Distributed multi-MMW radar fusion for target detection and tracking in highway traffic environment. *Mathematical Problems in Engineering*, 2023, 5537122. https://doi.org/10.1155/2023/5537122
5. Carrillo, J., Crowley, M., Pan, G., & Fu, L. (2020). Design of efficient deep learning models for determining road surface condition from roadside camera images and weather data. *arXiv preprint arXiv:2009.10282*. https://doi.org/10.48550/arXiv.2009.10282
6. Ren, M., Zhang, X., Chen, X., Zhou, B., & Feng, Z. (2023). YOLOv5s-M: A deep learning network model for road pavement damage detection from urban street-view imagery. *International Journal of Applied Earth Observation and Geoinformation*, 120, 103335. https://doi.org/10.1016/j.jag.2023.103335
7. Guo, G., & Zhang, Z. (2022). Road damage detection algorithm for improved YOLOv5. *Scientific Reports*, 12, 15523. https://doi.org/10.1038/s41598-022-19674-8
8. Huang, Y. T., Jahanshahi, M. R., Shen, F., & Mondal, T. G. (2023). Deep learning-based autonomous road condition assessment leveraging inexpensive RGB and depth sensors and heterogeneous data fusion: Pothole detection and quantification. *Journal of Transportation Engineering, Part B: Pavements*, 149(2), 04023010. https://doi.org/10.1061/JPEODX.PVENG-1194
9. Dasari, D., Vijaya Lakshmi, D. N. V. S., & Tatapudi, A. (2024). Real-time detection of road damages using YOLOv8: An innovative deep learning approach. *International Journal of Emerging Technologies and Innovative Research (JETIR)*, 11(12), JETIR2412755. https://www.jetir.org/papers/JETIR2412755.pdf
10. Zhong, J., Kong, D., Wei, Y., & Pan, B. (2025). YOLOv8 and point cloud fusion for enhanced road pothole detection and quantification. *Scientific Reports*, 15(1), 11260. https://doi.org/10.1038/s41598-025-94993-0