

Full Length Article

Edge-Ready Road Damage Detection Using Enhanced Yolov10n With Hyperparameter Tuning

Mohammed Ahmed Ali¹, Md Abdul Sami², Abu Bakar Siddiq³, Mr. Harendra Reddy⁴

^{1,2,3}B.E.Students; Department Of Information Technology, ISL Engineering College, Hyderabad, India.

⁴Assistant professor; Department Of Information Technology, ISL Engineering College, Hyderabad, India.

Mail Id: mohammedahmedali42002@gmail.com, samimohdabdul14@gmail.com,

mohdsiddiq72306@gmail.com

Accepted 24-04-2026

Author(s) Retains the Copyrights of This Article

Abstract

Road infrastructure plays a critical role in transportation systems, economic development, and public safety. However, road defects such as potholes, cracks, and surface wear can lead to traffic accidents, vehicle damage, and increased maintenance costs if not detected at an early stage. Traditional road inspection methods rely heavily on manual monitoring, which is time-consuming, labor-intensive, and often inefficient for large-scale road networks. Recent advancements in deep learning and computer vision have enabled automated road damage detection systems capable of identifying defects with high accuracy in real time. This paper presents a lightweight road damage detection framework using YOLOv10n with hyperparameter tuning for deployment on edge computing devices. The proposed system is designed to detect multiple types of road damage, including potholes, longitudinal cracks, transverse cracks, and surface deterioration, while maintaining low computational complexity. The methodology includes data collection, annotation, preprocessing, augmentation, model training, hyperparameter optimization, testing, and deployment on embedded edge devices such as NVIDIA Jetson Nano and NVIDIA AGX Orin. Experimental results demonstrate strong detection performance with Precision of 98.6%, Recall of 97.3%, F1-Score of 97.8%, and mAP@0.5 of 98.8%. Real-time deployment evaluation shows inference speeds of 7.5 FPS on Jetson Nano and 67 FPS on AGX Orin, confirming the effectiveness of the proposed lightweight architecture for smart transportation and intelligent road maintenance applications. The system offers an efficient, scalable, and practical solution for smart city infrastructure management.

Keywords

Road Damage Detection, YOLOv10n, Hyperparameter Tuning, Deep Learning, Edge Computing, Smart Cities, Computer Vision, Pothole Detection, Intelligent Transportation Systems, Real-Time Detection.

Introduction

Road transportation networks are among the most important infrastructures in modern society because they support economic growth, industrial development, public mobility, and emergency services. The quality and condition of roads directly affect transportation safety, vehicle performance, fuel efficiency, and overall traffic management. However, road surfaces are continuously exposed to environmental conditions such as rainfall, temperature changes, heavy traffic loads, and natural wear and tear. Over time, these factors contribute to road defects including potholes, cracks, surface erosion, and pavement deterioration. Road damage creates serious safety concerns for drivers, passengers, and pedestrians. Potholes can cause vehicle accidents, tire damage, suspension failure, and traffic congestion. Cracks and surface wear reduce road durability and increase maintenance costs if not repaired at an early stage. According to transportation studies, delayed road maintenance leads to significantly higher repair

expenses and increased accident risks. Therefore, early detection and continuous monitoring of road damage are essential for maintaining safe and efficient transportation systems.

Traditional road inspection methods primarily rely on manual surveys conducted by maintenance personnel. These inspections are time-consuming, expensive, labor-intensive, and often inconsistent due to human error. Manual inspection also becomes impractical for large-scale road networks, highways, and urban infrastructure systems. In recent years, automated road damage detection systems based on computer vision and deep learning have emerged as promising alternatives.

Deep learning models, especially convolutional neural networks (CNNs), have shown remarkable success in image classification, segmentation, and object detection tasks. Among these approaches, the YOLO (You Only Look Once) family of object detection models has gained widespread popularity due to its balance between detection accuracy and real-time inference speed. YOLO models perform

260

object localization and classification simultaneously within a single forward pass, making them highly suitable for real-time applications.

Although advanced object detection models provide high accuracy, many of them require substantial computational resources and high-end GPUs. Such requirements limit their deployment on resource-constrained edge devices commonly used in smart city applications. Lightweight models such as YOLOv10n are specifically designed to reduce parameter count and computational complexity while maintaining strong detection performance.

This project proposes a lightweight road damage detection system using YOLOv10n with hyperparameter tuning for edge computing environments. The proposed framework focuses on achieving high detection accuracy while maintaining low memory usage and real-time performance. The optimized model is deployed on NVIDIA Jetson Nano and NVIDIA AGX Orin platforms to evaluate practical deployment capability. The proposed system aims to support intelligent transportation systems, smart city infrastructure management, and proactive road maintenance operations.

Literature Review

Road damage detection has become an important research area within intelligent transportation systems and smart infrastructure monitoring. Several traditional and deep learning-based approaches have been proposed to automate pavement inspection and defect analysis.

Early road damage detection systems mainly relied on image processing techniques such as edge detection, thresholding, texture analysis, and morphological operations. These methods attempted to identify cracks and potholes using handcrafted features extracted from pavement images. Although these techniques were computationally efficient, they often failed under varying lighting conditions, shadows, and complex road textures.

With the advancement of deep learning, convolutional neural networks (CNNs) became widely used for automated road damage detection. CNN-based classification models demonstrated improved feature extraction capability compared with traditional image processing methods. Researchers applied architectures such as AlexNet, VGGNet, ResNet, and DenseNet for pavement distress classification tasks. While these models improved detection accuracy, they were not suitable for real-time object localization.

Object detection frameworks such as Faster R-CNN, SSD, and YOLO significantly improved road defect detection performance. Faster R-CNN achieved high detection accuracy through region proposal mechanisms, but its inference speed was relatively

slow for real-time applications. SSD introduced faster detection performance, but it struggled with small and irregular road defects.

The YOLO family of models revolutionized real-time object detection by performing localization and classification simultaneously. YOLOv3 and YOLOv4 demonstrated strong performance for road damage detection tasks. However, these models often required powerful GPUs and substantial computational resources.

Recent studies introduced YOLOv7 and YOLOv8-based road damage detection systems capable of achieving higher accuracy and improved feature extraction. Researchers also explored UAV-based pavement monitoring systems and smartphone-based road inspection platforms. Although these methods provided strong detection capability, many remained computationally expensive for deployment on edge devices.

YOLOv10n introduces a lightweight architecture optimized for edge computing applications. It reduces parameter count while maintaining high inference speed and strong detection accuracy. Hyperparameter tuning further improves model generalization and detection performance by optimizing learning rate, batch size, image resolution, and optimizer settings.

The proposed work builds upon these advancements by developing an optimized YOLOv10n-based framework for efficient real-time road damage detection on edge devices.

System Architecture

The proposed road damage detection system consists of several major components including image acquisition, preprocessing, model training, inference, and deployment.

The first stage involves capturing road images using cameras mounted on vehicles, drones, or mobile devices. These images are collected under different environmental conditions including sunlight, shadows, rain, and varying road textures.

The collected images are then passed through preprocessing and augmentation pipelines. Preprocessing improves image quality by resizing, normalization, and noise reduction. Data augmentation increases dataset diversity and improves model robustness against real-world conditions.

The processed images are used to train the YOLOv10n model. Hyperparameter tuning is performed to optimize model performance and computational efficiency. The trained model is validated using testing datasets and evaluated using standard object detection metrics.

Finally, the optimized model is deployed on edge devices such as NVIDIA Jetson Nano and AGX Orin for real-time road damage detection.

Methodology

The methodology of the proposed system includes multiple stages such as dataset preparation, annotation, preprocessing, training, tuning, testing, and deployment.

The first stage involves collecting road images from public datasets and real-world environments. The dataset contains different categories of road defects including potholes, longitudinal cracks, transverse cracks, and surface wear.

After data collection, annotation tools are used to create bounding box labels for each road defect. Proper annotation is essential for supervised object detection training.

Image preprocessing techniques are applied to improve dataset quality. These include resizing images to a fixed resolution, normalization, contrast enhancement, and noise reduction. Data augmentation methods such as rotation, flipping, scaling, brightness adjustment, and random cropping are applied to increase training diversity.

The YOLOv10n model is then trained using optimized hyperparameters. Important training parameters include learning rate, optimizer selection, batch size, momentum, and number of epochs. Hyperparameter tuning is performed to identify optimal configurations that maximize detection accuracy while minimizing computational overhead.

The trained model is validated using separate testing datasets. Performance metrics such as Precision, Recall, F1-Score, mAP@0.5, inference speed, and memory consumption are measured.

Finally, the optimized model is deployed on edge computing devices for real-time inference.

Implementation

The implementation of the proposed system combines deep learning algorithms, computer vision techniques, and embedded deployment strategies.

The first step involves collecting and annotating road images using annotation software. Bounding boxes are created around potholes, cracks, and damaged road surfaces.

The images are resized and normalized before being fed into the YOLOv10n network. The model architecture consists of backbone, neck, and detection head components optimized for lightweight inference.

Hyperparameter tuning is performed to improve convergence and detection performance. Parameters such as learning rate scheduling, optimizer selection, confidence threshold, and IoU threshold are adjusted experimentally.

The implementation algorithm follows these steps:

1. Collect road damage images from multiple environments.

2. Annotate defects using bounding box labeling tools.
3. Resize and normalize input images.
4. Apply augmentation techniques to increase dataset diversity.
5. Train YOLOv10n using optimized hyperparameters.
6. Validate detection performance using testing datasets.
7. Deploy the trained model on edge devices.
8. Perform real-time road damage detection.
9. Store and visualize detection results.

The system is implemented using Python, OpenCV, PyTorch, and the Ultralytics YOLO framework.

Hyperparameter Tuning

Hyperparameter tuning plays an important role in improving model performance and generalization capability. Improper parameter selection can lead to overfitting, underfitting, or unstable training.

Several important hyperparameters are optimized in this project:

Learning rate

Batch size

Momentum

Weight decay

Confidence threshold

IoU threshold

Image resolution

Number of epochs

Learning rate scheduling is used to improve convergence stability during training. Batch size is selected based on GPU memory constraints and training efficiency.

Confidence and IoU thresholds are carefully adjusted to balance false positives and missed detections. Data augmentation parameters are also optimized to improve robustness under varying environmental conditions.

The optimized hyperparameters significantly improve detection accuracy and reduce training instability.

Testing and Validation

Testing and validation are essential for evaluating system reliability and practical deployment performance.

Unit testing verifies individual components such as image preprocessing, dataset loading, annotation parsing, and model initialization.

Functional testing evaluates the ability of the model to correctly detect and classify road defects under different environmental conditions.

Performance testing measures:

Inference speed

GPU memory usage

CPU utilization

Detection latency

Stress testing evaluates system stability during continuous real-time operation.

Testing is performed on both NVIDIA Jetson Nano and NVIDIA AGX Orin platforms to analyze deployment capability under different hardware conditions.

Results and Discussion

Experimental results demonstrate the effectiveness of the proposed YOLOv10n-based road damage detection framework.

The optimized model achieved:

Precision = 98.6%

Recall = 97.3%

F1-Score = 97.8%

mAP@0.5 = 98.8%

These results indicate strong localization accuracy and reliable classification performance for multiple road defect categories.

The deployment evaluation shows that the model achieves approximately 7.5 FPS on NVIDIA Jetson Nano and 67 FPS on NVIDIA AGX Orin. These inference speeds confirm that the system is suitable for real-time deployment.

The lightweight architecture significantly reduces memory usage and computational complexity compared with larger object detection models. Despite reduced parameter count, the model maintains high detection performance due to optimized feature extraction and hyperparameter tuning.

The system performs effectively under varying lighting conditions, shadows, and road textures. Data augmentation techniques contribute to improved robustness and generalization capability. Overall, the proposed framework successfully balances detection accuracy, computational efficiency, and real-time capability.

Applications

The proposed road damage detection system has several important applications in intelligent transportation systems and smart city infrastructure. One major application is automated road inspection. Government agencies and municipal corporations can use the system to continuously monitor road conditions and identify damaged areas requiring maintenance.

The framework can also support predictive maintenance systems by enabling early defect detection before severe road deterioration occurs.

Another important application is autonomous vehicles and driver assistance systems. Accurate road condition awareness improves driving safety and navigation efficiency.

The system may also be integrated with cloud-based smart city platforms for centralized road monitoring and infrastructure management.

Because the framework is lightweight, it is suitable for deployment on:

Edge devices

Smart cameras

Drones

Mobile inspection vehicles

Embedded systems

Future Scope

Several future enhancements can further improve the proposed system.

One possible improvement is road damage severity assessment. Instead of only detecting defects, future systems may estimate damage severity levels and maintenance priority.

Multi-sensor fusion using LiDAR, thermal cameras, and GPS data may improve robustness and environmental awareness.

Cloud-edge integration can enable centralized monitoring of large-scale road networks while maintaining low-latency edge inference.

Future work may also include:

Real-time road condition mapping

3D pavement analysis

Automated repair recommendation systems

Federated learning for distributed training

Integration with autonomous vehicle platforms

These advancements can significantly improve intelligent infrastructure management in future smart cities.

Conclusion

The proposed YOLOv10n-based road damage detection system provides an efficient, lightweight, and accurate solution for real-time pavement monitoring. The framework combines deep learning, hyperparameter tuning, and edge computing deployment to achieve high detection performance with low computational requirements. Experimental results demonstrate strong accuracy with Precision of 98.6%, Recall of 97.3%, F1-Score of 97.8%, and mAP@0.5 of 98.8%. Deployment on NVIDIA Jetson Nano and AGX Orin confirms the practical feasibility of the system for real-world applications.

The lightweight design enables deployment on resource-constrained edge devices while maintaining reliable real-time performance. The proposed framework supports proactive road maintenance, intelligent transportation systems, and smart city development.

Future enhancements involving severity analysis, multi-sensor fusion, and cloud-edge integration can further improve system capability and scalability.

References

- 1) T. Qiu et al., "Edge Computing in Industrial Internet of Things: Architecture, Advances and

- Challenges,” IEEE Communications Surveys & Tutorials, vol. 22, no. 4, pp. 2462–2488, 2020.
- 2) H. Mahmudah et al., “Digital Twin: Challenge Road Damage Detection on Edge Device,” Proceedings of International Smart Systems Conference, 2024.
 - 3) L. Zhao et al., “Automatic Defect Detection of Pavement Diseases Using Deep Learning,” Remote Sensing, vol. 14, no. 8, 2022.
 - 4) C. Wang et al., “YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors,” Proceedings of CVPR, 2023.
 - 5) J. Redmon et al., “You Only Look Once: Unified, Real-Time Object Detection,” Proceedings of CVPR, 2016.
 - 6) A. Bochkovskiy et al., “YOLOv4: Optimal Speed and Accuracy of Object Detection,” arXiv preprint arXiv:2004.10934, 2020.
 - 7) Z. Ge et al., “YOLOX: Exceeding YOLO Series in 2021,” arXiv preprint arXiv:2107.08430, 2021.
 - 8) K. He et al., “Deep Residual Learning for Image Recognition,” Proceedings of CVPR, 2016.
 - 9) T. Lin et al., “Feature Pyramid Networks for Object Detection,” Proceedings of CVPR, 2017.
 - 10) A. Dosovitskiy et al., “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale,” International Conference on Learning Representations, 2021.