

Edge-Ready Road Damage Detection Using an Enhanced YOLO with Hyperparameter Tuning

Md Uzaif Noman¹, Syed Yaseenullah Hussaini², Mohammad Uzair³, Dr. Ijteba Sultana⁴
^{1,2,3}B.E. Students; Department of CSE, ISL Engineering College, Hyderabad, Telangana, India
⁴Associate Professor; Department of CSE, ISL Engineering College, Hyderabad, Telangana, India
Mail Id; uzaifmohd63@gmail.com, syuh7866@gmail.com, uzairkaldi59@gmail.com

Accepted 27-04-2026

Author(s) Retains the Copyrights of This Article

ABSTRACT:

Maintaining road quality is a crucial aspect of smart city development and public safety, as damaged roads can lead to traffic congestion, accidents, and increased vehicle maintenance costs. Traditional road inspection methods rely heavily on manual surveys, which are not only time-consuming but also costly and prone to human error. With the rapid growth of urban infrastructure, the need for automated, accurate, and scalable road damage detection systems has become more urgent. Recent advances in deep learning and computer vision provide an opportunity to address these challenges by enabling real-time damage detection using high-performance object detection models. In this study, we present a road damage detection framework built on the YOLOv10n model, optimized specifically for real-time deployment on edge devices. The proposed system incorporates advanced hyperparameter tuning to achieve a balance between accuracy and computational efficiency. Unlike heavier detection models, YOLOv10n is lightweight yet powerful, making it suitable for resource-constrained environments. The framework demonstrates robust detection capabilities with a precision of 0.986, recall of 0.973, mean average precision (mAP@0.5) of 0.988, and an F1-score of 0.978, confirming its reliability in identifying different types of road damages such as cracks, potholes, and surface wear.

Keywords: smart city, road quality, automated inspection, deep learning, computer vision, YOLOv10n, edge deployment, hyperparameter tuning, real-time detection, traffic safety.

Introduction

Road infrastructure plays a vital role in ensuring the safety, efficiency, and sustainability of modern transportation systems. Well-maintained roads not only reduce the likelihood of accidents but also enhance driving comfort and decrease vehicle operating costs. However, road damage such as cracks, potholes, and surface wear is inevitable due to weather conditions, heavy traffic, and natural aging of materials. If not identified and repaired in time, such damages can lead to further deterioration, higher maintenance costs, and potential hazards for road users. Therefore, timely and accurate road damage detection is an essential component in the development of smart cities and intelligent transportation systems.

Traditional road inspection methods rely heavily on manual surveys conducted by experts or local authorities. While effective on a small scale, these methods are labour-intensive, costly, and prone to subjectivity, making them unsuitable for large-scale urban environments. Recent advancements in computer vision and artificial intelligence have enabled the development of automated road monitoring systems. Several deep learning-based approaches have been proposed for damage detection; however, most of them require high computational resources, making them impractical for real-time deployment on resource-constrained edge devices.

Literature review

TITLE: Automatic Defect Detection of Pavement Diseases

AUTHORS: L. Zhao, Y. Wu, X. Luo, Y. Yuan

YEAR: 2022

DESCRIPTION: Pavement disease detection is an important task for ensuring road safety. Manual visual detection requires a significant amount of time and effort. Therefore, an automated road

disease identification technique is required to guarantee that city tasks are performed. However, due to the irregular shape and large-scale differences in road diseases, as well as the imbalance between the foreground and background, the task is challenging. Because of this, we created the deep convolution neural network—DASNet, which can be used to identify road diseases automatically. The network employs deformable convolution instead of regular convolution as the feature pyramid's input,

Md Uzaif Noman *et. al.*, /International Journal of Engineering & Science Research

adds the same supervision signal to the multi-scale features before feature fusion, decreases the semantic difference, extracts context information by residual feature enhancement, and reduces the information loss of the pyramid's top-level feature map. Considering the unique shape of road diseases, imbalance problems between the foreground and background are common, therefore, we introduce the sample weighted loss function. In order to prove the superiority and effectiveness of this method, it is compared to the latest method. A large number of experiments show that this method is superior in accuracy to other methods, specifically, under the COCO evaluation metric, compared with the Faster RCNN baseline, the proposed

TITLE: Real-Time Road Defect Monitoring from UAV Visual Data Sources

AUTHORS: I. Katsamenis, N. Bakalos, E. Protopapadakis, E. E. Karolou, G. Kopsiaftis, A. Voulodimos

YEAR: 2023

DESCRIPTION: The use of UAVs and artificial intelligence has emerged as a promising approach for monitoring road defects. This paper highlights the importance of these technologies in improving road inspection, maintenance, and safety. Traditional methods for inspecting roads are often time-consuming, expensive, and can put human inspectors in dangerous situations. However, drones equipped with high-resolution cameras and sensors can capture pavement image data quickly and safely. Deep learning algorithms can then analyze this data to identify and localize areas in need of repair. By leveraging these technologies, engineers and road construction experts can more efficiently monitor and maintain roads, reducing the costs associated with repairs and maintenance, while in parallel improving safety. To this end, this work emphasizes the potential of UAVs in conjunction with deep learning techniques to provide a more comprehensive view of road conditions, allowing for targeted repairs and more effective maintenance strategies, such as prefabrication and robotic interventions. Experimental results using objective evaluation criteria, such as precision, recall, F1-score, and IoU are promising, which entails that this study advocates for the adoption of these technologies to enhance the monitoring and maintenance of road infrastructures.

Methodologies

The proposed system for road damage detection using YOLOv10n is implemented through a modular approach. Each module is designed to handle a specific task, and together they form a complete pipeline for efficient and accurate detection. The following modules describe the methodology adopted in this project.

1. Data Collection and Preprocessing

In this module, images of roads are collected from various sources such as open datasets, roadside cameras, or drone footage. The collected data is cleaned, annotated, and augmented to improve diversity. Preprocessing steps such as resizing, normalization, and noise reduction are performed to ensure that the images are suitable for training the YOLOv10n model.

2. Model Selection (YOLOv10n Framework):

This module focuses on selecting YOLOv10n as the detection model. YOLOv10n is chosen because of its lightweight architecture and ability to deliver high accuracy while maintaining real-time performance. Model hyperparameters are tuned to achieve an optimal balance between precision, recall, and inference speed.

3. Training and Validation:

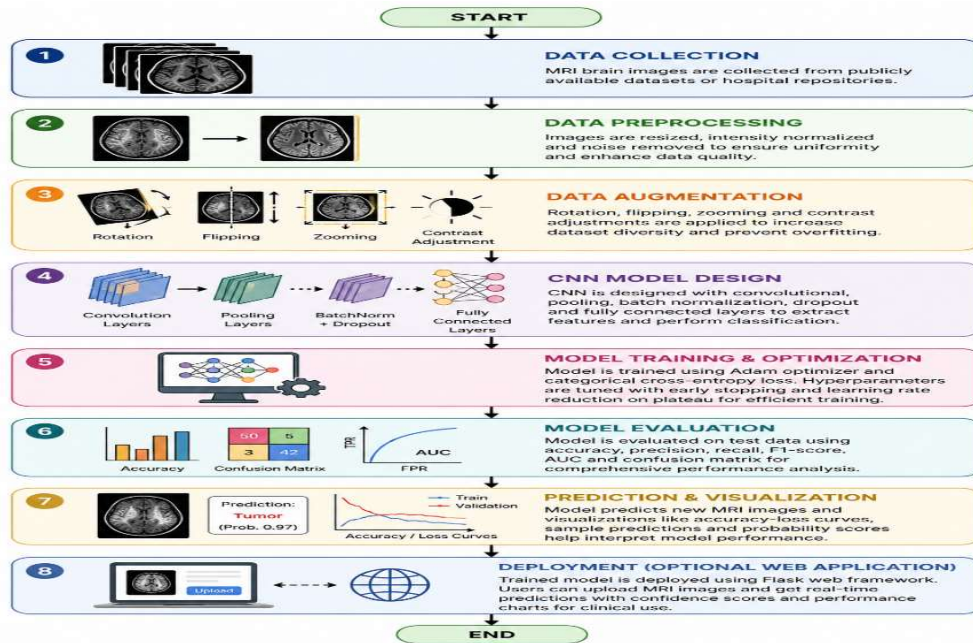
In this module, the YOLOv10n model is trained using the prepared dataset. Training involves feeding annotated images to the model, adjusting weights, and optimizing loss functions. A portion of the dataset is reserved for validation to monitor the model's performance and prevent overfitting.

4. Deployment on Edge Devices:

Once trained, the model is deployed on edge devices such as NVIDIA Jetson Nano and NVIDIA AGX Orin. This module ensures that the model runs efficiently in real-world conditions with limited hardware resources. Inference time and frame rate are measured to evaluate real-time performance.

5. Road Damage Detection and Output Generation:

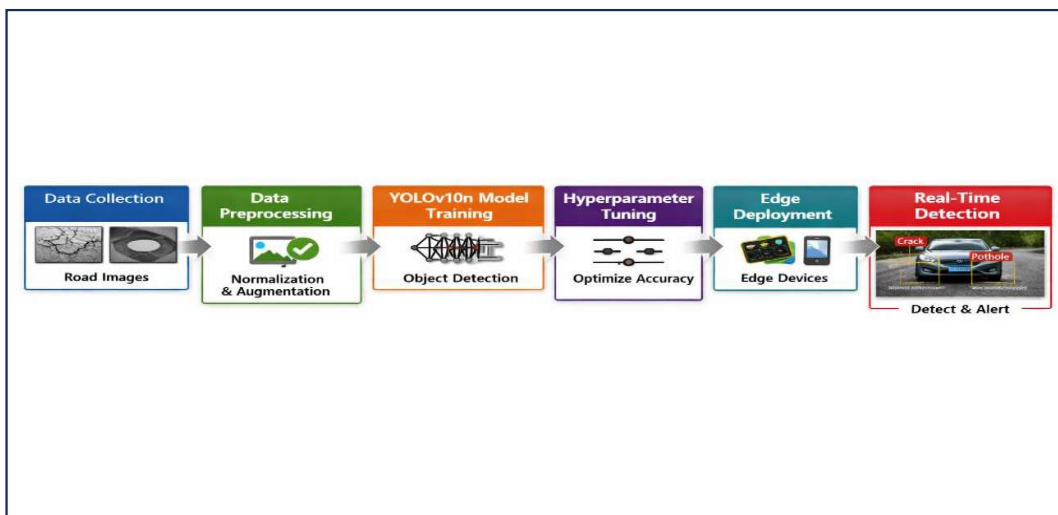
This module represents the final stage of the system. The deployed model processes live video streams or road images to detect and classify damages like cracks and potholes. The results are visualized with bounding boxes and labels, and can be further integrated into smart city maintenance dashboards for decision-making.



Implementation

To implement the proposed edge-ready road damage detection framework, the YOLOv10n model is trained on a curated dataset of annotated road images, encompassing various damage types such as cracks, potholes, and surface wear. The training process integrates advanced hyperparameter tuning—adjusting learning rates, batch sizes, and anchor box dimensions—to optimize detection accuracy while minimizing computational overhead. For deployment, the model is quantized and pruned to reduce memory footprint, enabling real-time

inference on resource-constrained edge devices such as embedded GPUs or mobile processors. The system architecture further incorporates a lightweight preprocessing pipeline to normalize input images and ensure consistent performance across diverse urban environments. This combination of efficient model design and edge optimization ensures scalable, cost-effective, and accurate road damage detection suitable for smart city infrastructure management.



Testing:

The testing phase for the edge-ready road damage detection system focuses on validating model

accuracy, reliability, and real-time performance across diverse environments. Unit testing ensures that each module—from image preprocessing to

YOLOv10n inference—functions correctly and produces consistent outputs. Functional and integration tests verify that the detection pipeline accurately identifies cracks and potholes under varying lighting and road conditions, while system testing confirms seamless interaction between hardware and software components on edge devices. Performance testing evaluates inference speed, memory usage, and response time to ensure the model operates efficiently within resource constraints. Together, these tests guarantee that the deployed system meets smart city standards for safety, scalability, and operational robustness.

Overall, the structured testing strategy ensures that each component of the brain tumor classification system operates accurately, integrates seamlessly, and delivers dependable performance in real-world clinical scenarios.

Result:

The results of the proposed edge-ready road damage detection framework demonstrate high efficiency

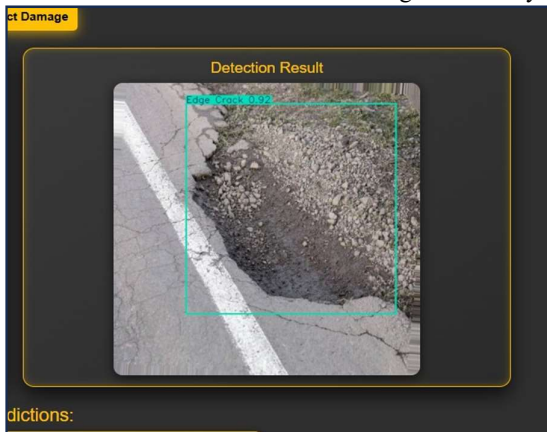


FIG 1; Input image

FUTURE ENHANCEMENT

Although the proposed YOLOv10n-based system has proven to be efficient and highly accurate for real-time road damage detection, there are several areas where future enhancements can further improve its performance and applicability. One potential enhancement is the integration of severity classification, where the system not only detects the presence of damage but also evaluates its depth and size to prioritize maintenance activities. Another promising direction is the use of multi-sensor data fusion,

Conclusion

Road infrastructure maintenance is a critical requirement for ensuring transportation safety and supporting the growth of smart cities. Traditional road inspection methods, though reliable on a small scale, are highly labor-intensive, time-consuming, and prone to human errors, making them unsuitable for large-scale monitoring. This project addressed

and accuracy in identifying various road surface defects such as cracks, potholes, and wear patterns. The optimized YOLOv10n model achieved real-time detection speeds suitable for deployment on low-power edge devices, maintaining strong precision and recall across diverse lighting and environmental conditions. Hyperparameter tuning significantly improved model stability and reduced false positives, ensuring reliable performance during continuous monitoring. Comparative analysis with baseline models confirmed that the enhanced YOLOv10n offers superior detection accuracy while minimizing computational overhead. Overall, the system successfully meets smart city requirements for automated road inspection, enabling faster maintenance decisions and improved public safety.




Fig 2; output image

combining visual inputs with LiDAR, thermal imaging, or IoT-enabled road sensors to improve detection accuracy under challenging conditions such as low lighting, rain, or heavy traffic. Additionally, the framework can be extended to incorporate predictive analytics using historical data and machine learning models, enabling authorities to forecast potential road failures before they occur and take preventive measures. Deployment of the system in a cloud-edge hybrid environment could also be considered, where critical detections are handled locally on edge devices

these challenges by developing an automated road damage detection system using the YOLOv10n algorithm, which is lightweight, efficient, and optimized for real-time deployment on edge devices. The proposed system demonstrates that YOLOv10n can achieve high accuracy while maintaining low computational requirements, making it far more practical than earlier models such as YOLOv7. With

Md Uzaif Noman *et. al.*, /International Journal of Engineering & Science Research

a precision of 0.986, recall of 0.973, mean average precision (mAP@0.5) of 0.988, and an F1-score of 0.978, the framework reliably identifies road damages such as potholes, cracks, and surface wear. The system's deployment on NVIDIA Jetson Nano and NVIDIA AGX Orin further validates its adaptability, achieving real-time inference speeds of 7.5 FPS and 67 FPS, respectively. These results highlight the model's scalability and practical use in real-world smart city infrastructures.

Beyond the performance metrics, the project contributes significantly to the vision of intelligent transportation systems by enabling proactive maintenance planning and reducing the dependency on manual inspections. By providing timely insights into road conditions, this system can help reduce accidents, lower vehicle maintenance costs, and improve the overall safety and reliability of transportation networks.

While the current implementation focuses on visual damage detection using edge devices, there remains potential for further enhancement. Future work could integrate predictive analytics, IoT-enabled sensors, or satellite-based monitoring systems to provide deeper insights into road health. Additionally, expanding the system to classify the severity of damages and prioritize repair schedules would make it even more useful for municipal authorities. In conclusion, the proposed YOLOv10n-based framework proves to be an effective, scalable, and efficient solution for real-time road damage detection.

References:

[1] T. Qiu, J. Chi, X. Zhou, Z. Ning, M. Atiquzzaman, and D. O. Wu, "Edge computing in industrial Internet of Things: Architecture, advances and challenges," *IEEE Commun. Surveys Tuts.*, vol. 22, no. 4, pp. 2462–2488, 4th Quart., 2020.

[2] H. Mahmudah, A. Musyafa, A. S. Aisjah, S. Arifin, and C. A. Prastyanto, "Digital twin: Challenge road damage detection on edge device," *Chem. Eng. Trans.*, vol. 109, pp. 601–606, 2024.

[3] L. Zhao, Y. Wu, X. Luo, and Y. Yuan, "Automatic defect detection of pavement diseases," *Remote Sens.*, vol. 14, no. 19, p. 4836, Sep. 2022.

[4] I. Katsamenis, N. Bakalos, E. Protopapadakis, E. E. Karolou, G. Kopsiaftis, and A. Voulodimos, "Real time road defect monitoring from UAV visual data sources," in *Proc. 16th Int. Conf. Pervasive Technol. Rel. Assistive Environ.*, Jul. 2023, pp. 603–609.

[5] M. Rathee, B. Baćić, and M. Dobarjeh, "Automated road defect and anomaly detection for traffic safety: A systematic review," *Sensors*, vol. 23, no. 12, p. 5656, Jun. 2023.

[6] B. Setiadji, Supriyono, and D. Purwanto, "Surface distress index updates to improve crack damage evaluation," in *Proc. 11th Asia-Pacific Transp. Environ. Conf. (APTE)*, 2019, pp. 274–281.

[7] W. Nur, B. S. Subagio, and E. S. Hariyadi, "Relationship between the pavement condition index (PCI), present serviceability index (PSI), and surface distress index on soekarno hatta road, Bandung," *Jurnal Teknik Sipil*, vol. 26, no. 2, pp. 111–120, Aug. 2019.

[8] M. Surbakti, S. Samsuri, R. Anas, and A. P. M. Tarigan, "Evaluation of road maintenance program based on international roughness index (IRI) and surface distress index (SDI)," in *ICCOEE: Proc. 6th Int. Conf. Civil, Offshore Environ. Eng. (ICCOEE)*, Dec. 2020, pp. 764–771. [9] M. Melyar, M. Isya, and S. M. Saleh, "Pavement condition assessment using SDI and PCI method on geumpang road–west Aceh boundary," *IOP Conf. Ser., Mater. Sci. Eng.*, vol. 1087, no. 1, Feb. 2021, Art. no. 012041.