ROAD DAMAGE PREDICTION USING MACHINE LEARNING

1Mrs. B. GOUTHAMI, 2M. SUMANTH, 3B. ASHRITH, 4B. NITHISH REDDY

1Assistant Professor, Department of CSE, Sri Indu College of Engineering and Technology-Hyderabad

234Under Graduate, Department of AI&DS, Sri Indu College of Engineering and Technology-Hyderabad

ABSTRACT: Road infrastructure maintenance is crucial for ensuring safe and efficient transportation, yet traditional manual inspection methods are often time-consuming, labor-intensive, and prone to human error, resulting in delayed maintenance and increased safety risks for road users. To overcome these limitations, this project introduces an AI-powered solution that leverages deep learning techniques—specifically the YOLOv8 (You Only Look Once version 8) object detection model—for automated road damage detection. The system is designed to accurately identify and classify various types of road surface damages, including longitudinal cracks, transverse cracks, alligator cracks, and potholes. Trained on the Crowdsensing-based Road Damage Detection Challenge 2022 (RDD2022) dataset, the model achieves high precision in detecting road anomalies under diverse environmental conditions. A user-friendly web interface built using Streamlit enables real-time analysis of images, videos, and live webcam feeds, allowing for immediate and interactive detection of road damages. This approach empowers infrastructure authorities with timely and reliable insights for planning and prioritizing maintenance tasks. By integrating advanced computer vision into traditional monitoring systems, the proposed solution enhances road safety, optimizes resource utilization, reduces operational costs, and minimizes traffic disruptions, thereby contributing to the development of smarter and more resilient transportation networks.

INTRODUCTION

Modern road infrastructure is an essential component of economic development, public safety, and urban planning. However, maintaining these roads in good condition poses a significant challenge due to factors such as increasing vehicle loads, environmental conditions, and aging infrastructure. Traditional manual methods of inspecting roads for damage—such as potholes, cracks, and surface deformities—are not only time-consuming and labor-intensive but also subject to human error and delays, resulting in inefficient maintenance operations and increased risks for road users.

To overcome these limitations, this project proposes an AI-driven solution that uses deep learning, particularly the YOLOv8 object detection algorithm, to automate road damage detection. The model is trained on the RDD2022 dataset, which includes annotated images representing various types of road damage, such as longitudinal cracks, transverse cracks, alligator cracks, and potholes. The project also integrates a Streamlit-based web interface, allowing users to upload images, videos, or access a webcam for real-time road condition analysis.

MOTIVATION

The motivation behind this project stems from the growing need for intelligent transportation systems and the limitations of manual road inspection methods. With increasing urbanization, road networks are expanding rapidly, making traditional monitoring approaches impractical. Automated road damage detection using AI offers a faster, more scalable, and more accurate alternative, which can assist municipalities and road management authorities in early detection and timely repair of road damage, thereby enhancing road safety and optimizing maintenance budgets.

LITERATURE SURVEY

Numerous studies and projects have explored the application of computer vision and machine learning for road surface analysis. Early approaches focused on classical image processing techniques, relying on edge detection and texture analysis to locate cracks or potholes. However, these methods often failed under varying lighting conditions or complex road textures.

Recent advancements in deep learning, particularly convolutional neural networks (CNNs), have significantly improved the accuracy of road damage detection. Projects utilizing models like Faster R-CNN, SSD, and YOLOv4 have shown promising results in detecting road anomalies. The release of large-scale annotated datasets like RDD2020, RDD2022, and GAPS has further enabled the training of robust models capable of performing well across diverse geographical regions and damage types.

In particular, YOLO (You Only Look Once) models have gained popularity due to their ability to perform object detection in real-time. YOLOv8, the latest version in the YOLO family, offers enhanced accuracy, speed, and flexibility. This project builds upon the success of previous models and datasets, utilizing

YOLOv8 and the RDD2022 dataset to create an end-to-end system for road damage detection and classification.

Data Collection and Preprocessing Module: This module is responsible for gathering a large and diverse dataset of road images containing various types of damage. The Crowdsensing-based Road Damage Detection (RDD2022) dataset is used, which includes annotations for different classes such as longitudinal cracks, transverse cracks, alligator cracks, and potholes. Preprocessing steps include resizing images, converting annotation formats to YOLO-compatible structure, and augmenting data to improve model generalization. The goal is to prepare clean, well- structured data that improves training efficiency and accuracy.

Model Training and Optimization Module: This module focuses on training the YOLOv8 object detection model using the preprocessed dataset. The architecture is optimized for detecting small and fine-grained damages in real-time. Hyperparameters such as learning rate, batch size, and number of epochs are tuned for maximum performance. Model weights are continuously evaluated and saved at checkpoints to avoid overfitting. Once training is complete, the best model is selected based on evaluation metrics such as precision, recall, and mAP (mean Average Precision).

Real-Time Detection and Classification Module: This module implements the real-time inference system using the trained YOLOv8 model. Users can input images, videos, or activate webcam-based detection to analyze roads on the spot. The model detects and classifies road damages in the visual data and displays them with bounding boxes and damage labels. The fast inference capability of YOLOv8 ensures smooth, real-time performance on standard hardware.

Streamlit-Based User Interface Module: This module provides an intuitive and interactive web interface built with Streamlit. Users can upload media files or start webcam capture directly through the browser. The interface processes the input, runs the YOLOv8 detection model, and displays the results with visual overlays.

It enhances usability for non-technical users such as road maintenance staff and government authorities by allowing easy access to detection results without needing command-line operations.

Result Visualization and Reporting Module: This module focuses on presenting results in a meaningful way for analysis and decision- making. Along with the visual display of detections, the module includes confidence scores, damage type summaries, and potential location tagging (if enabled). The system can generate downloadable reports for future use or integration with maintenance planning systems. It may also include alerts for severe or multiple damages, aiding timely infrastructure response.

SYSTEM ANALYSIS

EXISTING SYSTEM

The monitoring and maintenance of road infrastructure is a critical task carried out by governmental agencies and municipal bodies across the globe. Traditionally, the detection of road surface damage—such as cracks, potholes, and surface wear—is conducted through manual inspection methods. In these methods, trained inspectors are deployed on-site to visually examine the road surface, document damages, and prioritize repairs. The approach, though reliable in certain cases, is highly inefficient and resource-intensive, especially when dealing with large-scale infrastructure in urban and rural areas.

In some scenarios, basic semi-automated systems are used, where images are captured using drones or surveillance vehicles and then analyzed by humans or with limited software assistance. However, even these semi-automated methods rely heavily on manual interpretation and are not scalable or adaptable for real-time damage detection.

PROPOSED SYSTEM

The proposed system is an intelligent road damage detection application that leverages the advanced capabilities of YOLOv8 (You Only Look Once, Version 8), one of the most accurate and fastest object detection models available. This deep learning model is trained on a large-scale, annotated dataset — the Crowdsensing-based Road Damage Detection Challenge 2022 — which includes various types of real-world road anomalies such as longitudinal cracks, transverse cracks, alligator cracks, and potholes. The model is integrated into a user-friendly interface built using Streamlit, a lightweight and interactive Python-based framework, allowing users to either upload road images or stream live video for real-time analysis. The system processes these inputs instantly, displaying detected damages with labeled bounding boxes, making it easier for users to assess road conditions efficiently without manual inspection. The use of computer vision ensures high precision and speed, making the system ideal for large-scale implementation by road authorities and urban planners.

In addition to accurate and fast detection, the system offers a wide range of practical benefits. It significantly reduces reliance on time-consuming and error-prone manual inspections, saving operational costs and labour. By enabling real-time processing, it ensures timely insights that support preventive maintenance, reducing the risk of accidents and infrastructure degradation. The Streamlit interface ensures ease of access, even for non- technical users, while the backend can be extended to store detection logs for long-term analysis, maintenance planning, and damage trend forecasting. The system is also scalable, allowing deployment across entire cities or transportation networks using existing surveillance or drone infrastructure. Moreover, its ability to be retrained with new data allows for continuous improvement, making it a future-ready solution for smart transportation management. This approach not only enhances public safety but also aligns with modern smart city goals by integrating AI-driven decision-making into traditional civil infrastructure workflows.

IMPLEMENTATION AND RESULTS

ARCHITECTURE

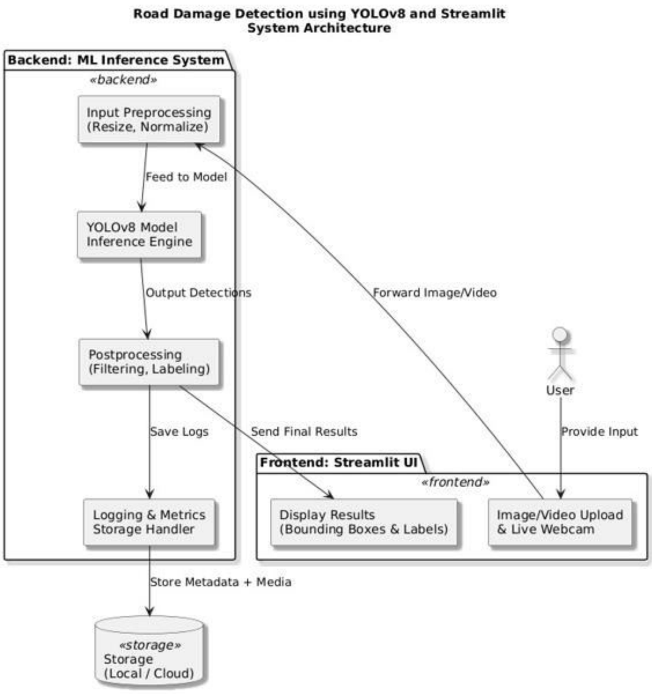

Right, both start with ABC. If we are talking about ABC in the Python context, it's clear that the programming language ABC is meant. ABC is a general-purpose programming language and programming environment, which had been developed in the Netherlands, Amsterdam, at the CWI

(Centrum Wiskunde &Informatica). The greatest achievement of ABC was to influence the design of Python. Python was conceptualized in the late 1980s. Guido van

Rossum worked that time in a project at the CWI, called Amoeba, a distributed operating system. In an interview with Bill Venners1, Guido van Rossum said: "In the early 1980s, I worked as an implementer on a team building a language called ABC at Centrum Wiskunde en Informatica (CWI). I don't know how well people know ABC's influence on Python. I try to mention ABC's influence because I'm indebted to everything I learned during that project and to the people who worked on it. Later on in the same Interview, Guido van Rossum continued: "I remembered all my experience and some of my frustration with

ABC. I decided to try to design a simple scripting language that possessed some of ABC's better properties, but without its problems. So, I started typing. I created a simple virtual machine, a simple parser, and a simple runtime. I made my own version of the various ABC parts that I liked. I created a basic syntax, used indentation for statement grouping instead of curly braces or begin-end blocks, and developed a small number of powerful data types: a hash table (or dictionary, as we call it), a list, strings, and numbers."

YOLOv8 TECHNOLOGY

YOLOv8 (You Only Look Once, version 8) is the latest and most optimized real-time object detection algorithm developed by Ultralytics. It belongs to the YOLO family of models known for balancing high speed with high accuracy. YOLOv8 builds upon previous versions with enhanced architecture, better bounding box regression, and strong support for both image and video data, making it ideal for tasks like road damage detection in real-world environments.

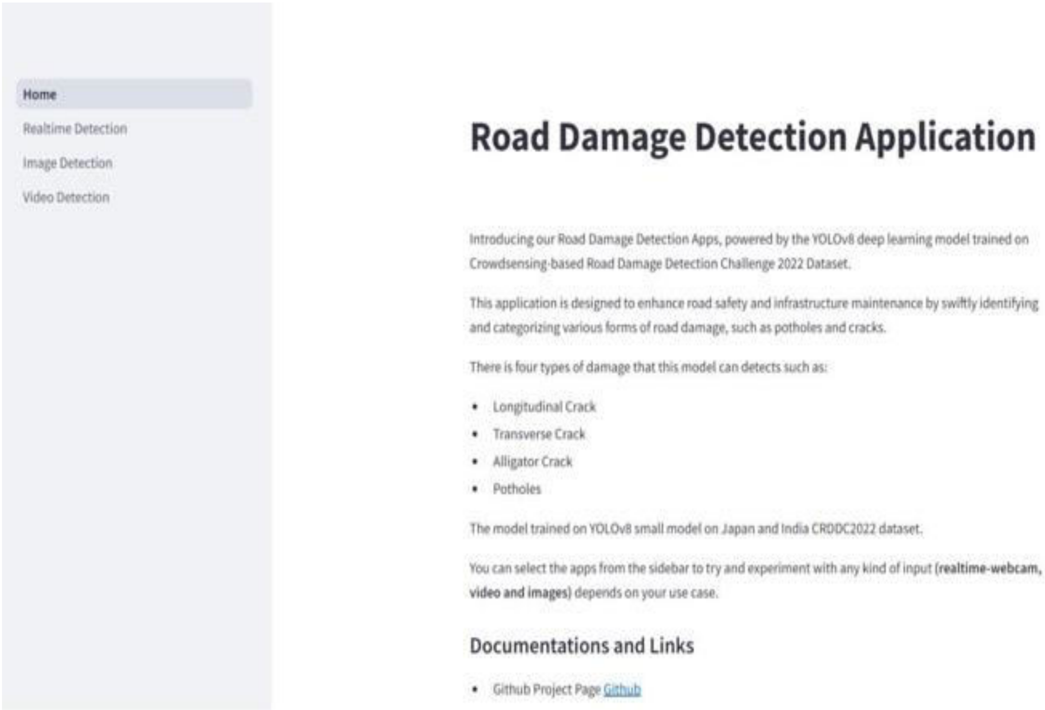
Home
Realtime Detection
Image Detection
Video Detection
Road Damage Detection Application
Introducing our Road Damage Detection Apps, powered by the YOLOv8 deep learning model trained on Crowdsensing-based Road Damage Detection Challenge 2022 Dataset.
This application is designed to enhance road safety and infrastructure maintenance by swiftly identifying and categorizing various forms of road damage, such as potholes and cracks.
There is four types of damage that this model can detects such as:
Longitudinal Crack
Transverse Crack
Alligator Crack
Potholes
The model trained on YOLOv8 small model on Japan and India CRDDC2022 dataset.
You can select the apps from the sidebar to try and experiment with any kind of input (realtime-webcam, video and images) depends on your use case.
Documentations and Links
Github Project Page Github

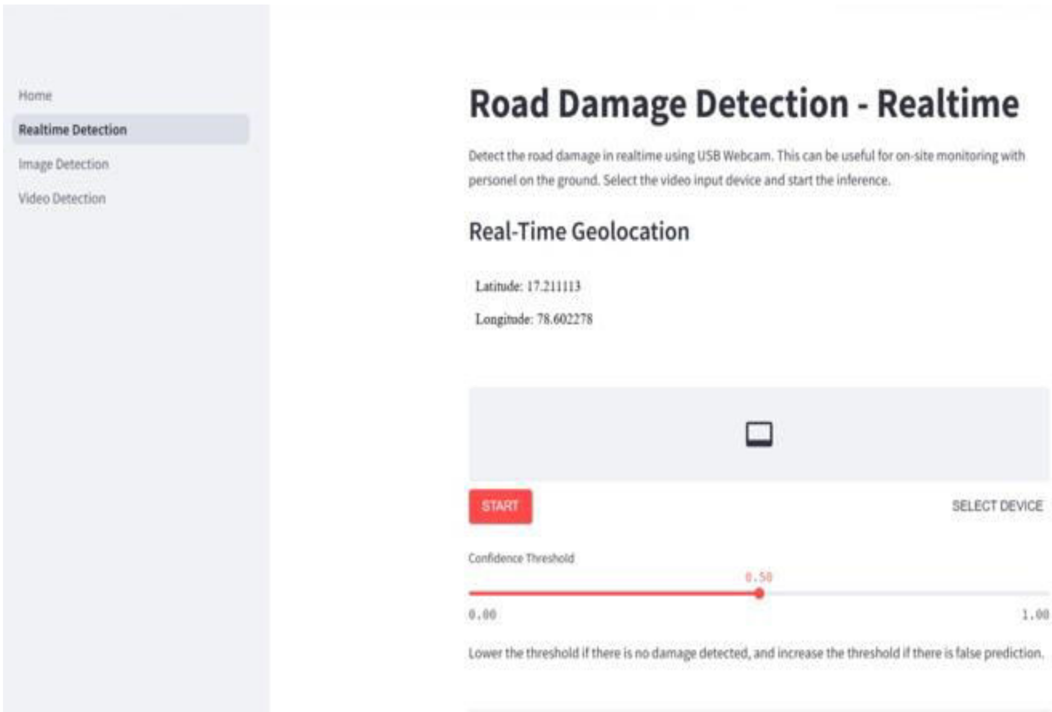
Home
Realtime Detection
Image Detection
Video Detection
Road Damage Detection - Realtime
Detect the road damage in realtime using USB Webcam. This can be useful for on-site monitoring with personel on the ground. Select the video input device and start the inference.
Real-Time Geolocation
Latitude: 17.211113
Longitude: 78.602278
START
SELECT DEVICE
Confidence Threshold
0.50
0.00
1.00
Lower the threshold if there is no damage detected, and increase the threshold if there is false prediction.

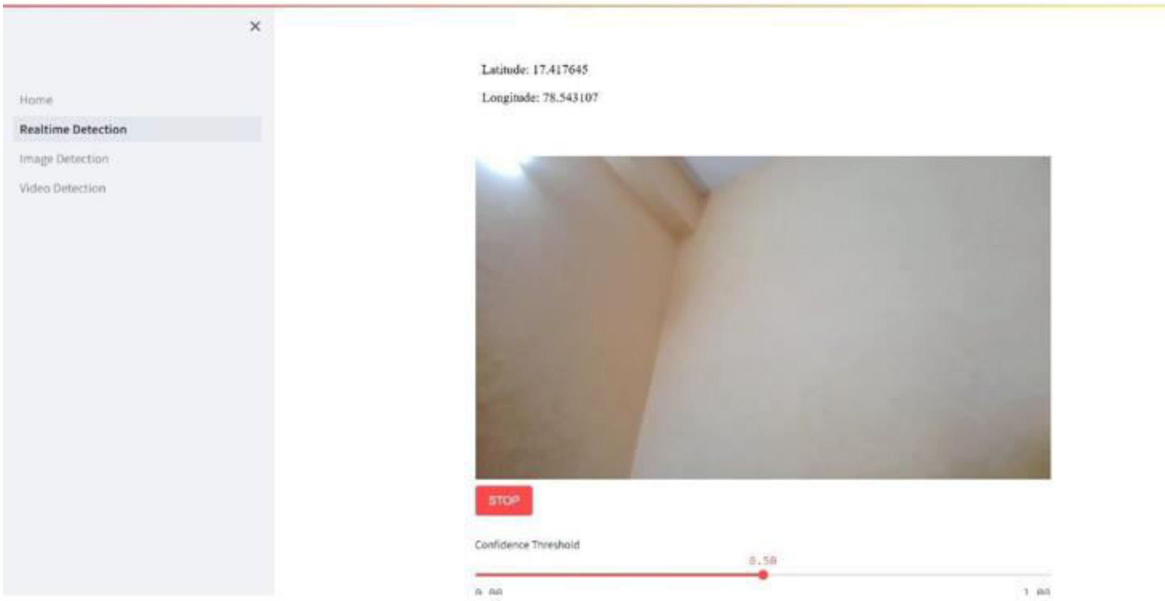
Home
Realtime Detection
Image Detection
Video Detection
Latitude: 17.417645
Longitude: 78.543107
STOP
Confidence Threshold
0.50
0.00
1.00

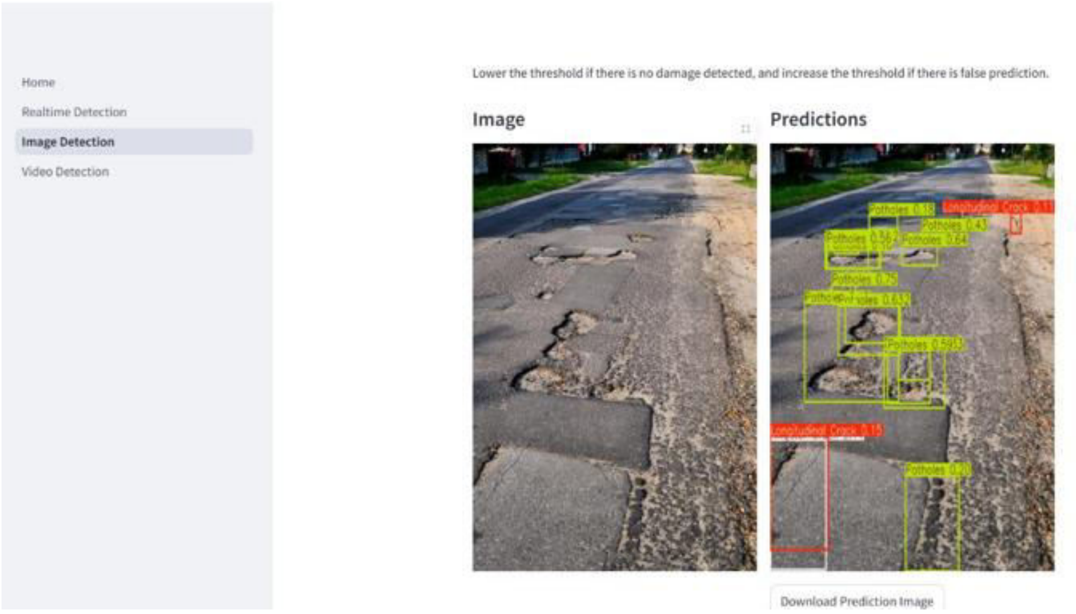
Home
Realtime Detection
Image Detection
Video Detection
Lower the threshold if there is no damage detected, and increase the threshold if there is false prediction.
Image
Predictions
Download Prediction Image

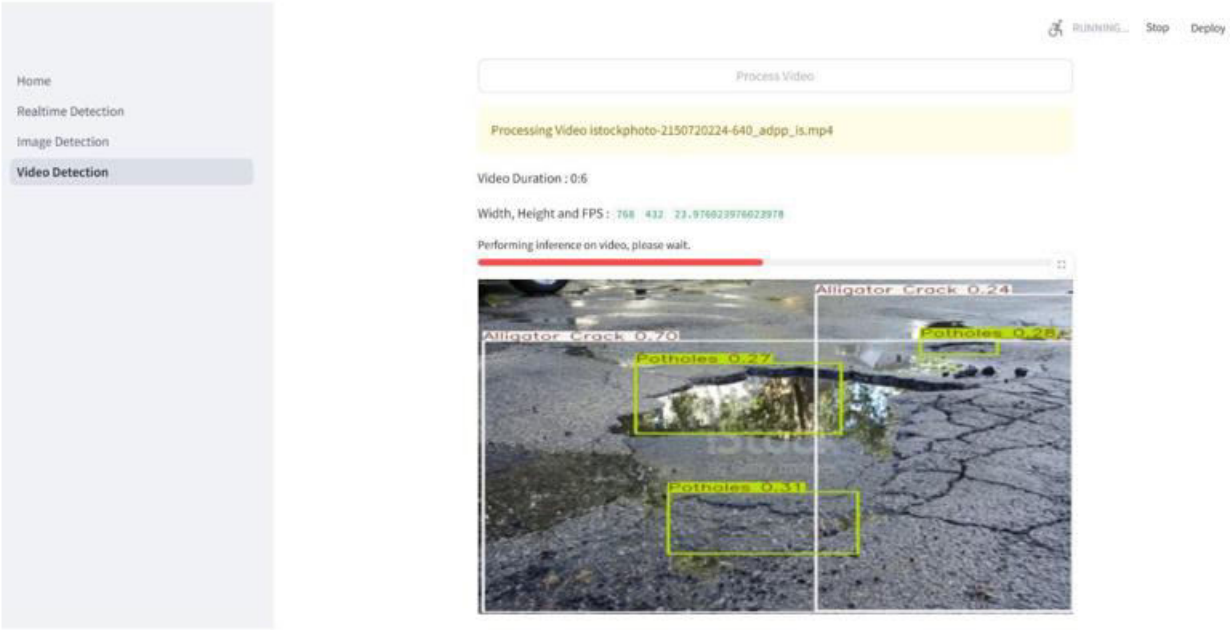

CONCLUSION

The Road Damage Detection System using YOLOv8 and Streamlit effectively addresses the challenge of identifying and categorizing road damages from images and video inputs. By integrating a powerful deep learning model (YOLOv8) with a user-friendly web interface (Streamlit), the system enables real-time detection and visualization of various types of road defects such as cracks, potholes, and surface damage.

The modular architecture ensures: Efficient preprocessing and inference, making the system scalable for real-world applications. Interactive and responsive UI, allowing users to easily upload images/videos or use a live webcam feed. Accurate localization and labeling, which helps municipal authorities and road maintenance teams prioritize repair work. The project also demonstrates the potential of combining state-of-the-art AI models with accessible deployment tools to build practical, impactful solutions for infrastructure management. In conclusion, this system not only improves the efficiency and accuracy of road condition assessments but also lays the foundation for smart city initiatives and automated monitoring systems. Future enhancements can include integrating GPS metadata, mobile app development, and cloud-based dashboards for large-scale deployment.

REFERENCES

- Bochkovskiy, A., Wang, C.Y., & Liao, H.Y.M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934. https://arxiv.org/abs/2004.10934
- Jocher, G., Chaurasia, A., Qiu, J., & Stoken, A. (2023). YOLOv8 - Ultralytics. GitHub Repository. https://github.com/ultralytics/ultralytics
- Streamlit Team. (2023). Streamlit: The fastest way to build and share data apps. https://streamlit.io/
- Sekilab. (2021). Road Damage Detector GitHub Repository. https://github.com/sekilab/RoadDamageDetector
- Kaggle. (2021). Global Road Damage Detection Dataset (GRDDC). https://www.kaggle.com/datasets/saurav9786/global-road-damage-detection
- OpenCV. (2023). Open Source Computer Vision Library. https://opencv.org/
- Kingma, D.P., & Ba, J. (2015). Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980.
- Microsoft. (2023). Visual Studio Code - Code Editing. Redefined. https://code.visualstudio.com/