

Real-Time Pothole Detection Using Deep Learning and Computer Vision Techniques

SRIMAT TIRUMALA KUNCHAPUDI SRI

PG Scholar, Department of MCA, DNR College, Bhimavaram, Andhra Pradesh

A. Naga Raju

(Assistant Professor), Master of Computer Applications, DNR College, Bhimavaram, Andhra Pradesh

ABSTRACT

Road infrastructure plays a crucial role in transportation efficiency and public safety. However, potholes remain one of the most common and hazardous issues affecting road quality, leading to vehicle damage, traffic congestion, and accidents. Traditional methods of pothole detection rely heavily on manual inspection, which is time-consuming, inefficient, and prone to human error. With advancements in deep learning and computer vision, automated pothole detection systems have emerged as a promising solution to address these challenges. This project presents a real-time pothole detection system using deep learning-based object detection techniques. The system leverages the YOLO (You Only Look Once) model, a state-of-the-art algorithm known for its speed and accuracy in detecting objects within images and videos. The model is trained on a dataset containing pothole images and is capable of identifying potholes in both static images and real-time video streams. The integration of ByteTrack tracking algorithm enhances the system's ability to track detected potholes across frames, ensuring continuity and improved detection reliability. The proposed system is implemented using Streamlit, providing an interactive web-based interface that allows users to upload images or videos or use sample data for testing. The system processes input data, detects potholes, and displays annotated results with bounding masks highlighting the detected regions. Additionally, it provides real-time performance metrics such as frame rate (FPS), ensuring transparency in system performance. The use of segmentation-based detection further improves the accuracy of pothole identification by precisely outlining the affected areas rather than just providing bounding boxes. This enables better assessment of pothole size and severity, which can be useful for road maintenance planning. The system also allows customization of parameters such as confidence thresholds and frame size, making it adaptable to different use cases and environments. Experimental results demonstrate that the proposed system achieves high accuracy and efficiency in detecting potholes under various lighting and road conditions. The integration of deep learning models with real-time processing capabilities makes this system suitable for deployment in smart transportation systems, autonomous vehicles, and road monitoring applications. In conclusion, this project highlights the potential of deep learning and computer vision in automating infrastructure monitoring tasks. The proposed pothole detection system not only reduces manual effort but also enhances road safety by enabling timely identification and repair of road damages. Future work may include integrating GPS tagging for detected potholes and deploying the system on edge devices for large-scale real-world implementation.

Keywords: Pothole Detection, Deep Learning, YOLO, Computer Vision, Road Safety, Image Processing, Object Detection, Streamlit, Smart Transportation, Autonomous Monitoring

I. INTRODUCTION

Road safety is a fundamental concern in modern transportation systems, directly impacting economic growth and human well-being. Among various road defects, potholes are one of the most prevalent and dangerous issues, especially in developing countries where road maintenance is often inconsistent. Potholes can cause severe damage to vehicles, increase the risk of accidents, and lead to traffic disruptions. Therefore, efficient detection and monitoring of potholes are essential for maintaining road quality and ensuring safety. Traditional pothole detection methods involve manual surveys conducted by road inspection teams. While effective to some extent, these methods are labor-intensive, time-consuming, and lack scalability. Moreover, they cannot provide real-time monitoring, which is crucial for immediate maintenance actions. With the increasing availability of cameras, sensors, and computational resources, automated solutions have become a viable alternative. Recent advancements in deep learning and computer vision have revolutionized the way objects are detected and analyzed in images and videos. Algorithms such as YOLO (You Only Look Once) have demonstrated remarkable performance in real-time object detection tasks. These models are capable of processing large amounts of visual data quickly and accurately, making them suitable for applications like pothole detection. The proposed system utilizes a deep learning-based approach to detect potholes in real-time using image and video inputs. By employing the YOLO model, the system can identify potholes with high precision and speed. The integration of tracking algorithms further enhances detection by maintaining consistency across video frames. This ensures that potholes are not missed during continuous monitoring. Another important aspect of the system is its user-friendly interface, developed using Streamlit. This allows users to easily interact with the system, upload data, and visualize results without requiring technical expertise. The system supports both image and video inputs, making it versatile for different applications such as surveillance cameras, vehicle-mounted cameras, and drone-based monitoring. The use of segmentation techniques provides an additional advantage by accurately outlining pothole regions. This enables better analysis of pothole characteristics, such as size and shape, which can be useful for prioritizing repairs. Furthermore, the system provides real-time feedback on performance metrics, ensuring transparency and reliability. In summary, this project aims to develop an efficient, accurate, and scalable pothole detection system using deep learning and computer vision techniques. By automating the detection process, the system reduces manual effort, improves detection accuracy, and enables timely maintenance actions, ultimately contributing to safer and better road infrastructure.

II. LITERATURE SURVEY (WITH EXISTING METHODS)

The problem of pothole detection has been widely studied in recent years, with various approaches proposed to improve detection accuracy and efficiency. Early methods relied on traditional image processing techniques such as edge detection, thresholding, and texture analysis. These methods were limited in their ability to handle complex road conditions and varying lighting environments. One of the earliest approaches involved using vibration sensors and accelerometers mounted on vehicles to detect road anomalies. While effective in detecting irregularities, these methods lacked the ability to visually confirm potholes and often produced false positives due to speed variations and other factors. Similarly, laser scanning and 3D reconstruction techniques were used to identify road surface defects, but these methods required expensive equipment and were not suitable for large-scale deployment. With the advent of machine learning, researchers began exploring classification-based approaches for pothole detection. Support Vector Machines (SVM) and Random Forest algorithms were used to classify road images into pothole and non-pothole categories. These methods improved accuracy but still required manual feature extraction, which limited their performance. Deep learning techniques, particularly Convolutional Neural Networks (CNNs), have significantly improved pothole detection capabilities. CNN-based models can automatically learn relevant features from raw image data, eliminating the need for manual feature engineering. Models such as Faster R-CNN, SSD (Single Shot Detector), and YOLO have been widely used for object detection tasks, including pothole detection. Among these, YOLO stands out due to its real-time performance and high accuracy. Unlike other models that process images in multiple stages, YOLO performs detection in a single pass, making it extremely fast. Recent versions of YOLO also support segmentation, allowing precise localization of objects. In addition to detection, tracking algorithms such as ByteTrack have been introduced to improve consistency in video-based applications. These algorithms track detected objects across frames, reducing false detections and improving reliability. Recent studies have also explored the integration of deep learning models with web-based frameworks for user-friendly deployment. Tools like Streamlit have made it easier to develop interactive applications for real-time monitoring and visualization. Despite these advancements, challenges remain in handling diverse road conditions, varying lighting environments, and occlusions. However, ongoing research continues to improve model robustness and efficiency.

III. EXISTING SYSTEM

The existing systems for pothole detection primarily rely on manual inspection and traditional sensing methods. Road maintenance authorities typically deploy personnel to physically inspect roads and identify potholes. This approach is not only time-consuming but also inefficient, especially for large road networks. Additionally, manual inspection is prone to human error and lacks consistency. Sensor-based approaches have also been used, where accelerometers and vibration sensors are installed in vehicles to detect road irregularities. These systems analyze changes in vehicle motion to identify potential potholes. However, they often produce inaccurate results due to variations in vehicle speed, suspension systems, and road conditions.

Another approach involves using basic image processing techniques such as edge detection and thresholding. While these methods can detect surface irregularities, they struggle with complex backgrounds and varying lighting conditions. As a result, their accuracy is limited. Machine learning-based systems improved detection by using classifiers like SVM and Decision Trees. However, these systems require manual feature extraction, which is time-consuming and may not capture all relevant features. Overall, existing systems lack real-time capabilities, scalability, and accuracy, making them unsuitable for modern smart transportation systems.

IV. PROPOSED METHOD

The proposed system introduces a real-time pothole detection framework using deep learning and computer vision techniques. It leverages the YOLO model for fast and accurate object detection, enabling the identification of potholes in both images and video streams. The system integrates ByteTrack for object tracking, ensuring that detected potholes are consistently tracked across video frames. This reduces false detections and improves overall reliability. Additionally, segmentation-based detection is used to precisely outline pothole regions, providing more detailed information about their size and shape. A user-friendly interface is developed using Streamlit, allowing users to upload images or videos and view detection results in real time. The system also provides performance metrics such as frame rate and allows users to adjust parameters like confidence thresholds. Unlike existing systems, the proposed solution offers high accuracy, real-time processing, and scalability. It can be deployed in various applications, including smart city infrastructure, autonomous vehicles, and road maintenance systems. By automating pothole detection, the system reduces manual effort, improves efficiency, and enhances road safety through timely identification and repair of road damages.

V. IMPLEMENTATION

The implementation of the pothole detection system is carried out using deep learning, computer vision, and a web-based interface. The system is developed using Python programming language with integration of frameworks such as Streamlit, OpenCV, and the Ultralytics YOLO model. The implementation process consists of multiple stages, including data preparation, model loading, real-time detection, tracking, and visualization. Initially, a pre-trained YOLO model is used, specifically trained on pothole datasets. The model weights are loaded from a trained file, enabling the system to perform object detection without retraining. YOLO is chosen due to its ability to perform real-time detection with high accuracy and efficiency. Studies have shown that YOLO-based systems can achieve real-time speeds of around 30 FPS while maintaining good detection accuracy. The system accepts both image and video inputs. For video input, OpenCV is used to capture frames from the video stream. Each frame is processed individually and passed to the YOLO model for detection. The model identifies potholes and generates bounding masks around detected regions. The detected results are then passed through the ByteTrack algorithm, which ensures object tracking across multiple frames, improving detection stability.

For image input, the uploaded image is processed directly without frame splitting. The YOLO model detects potholes and returns annotated images with highlighted regions. The use of segmentation-based detection allows the system to outline potholes precisely rather than using simple bounding boxes. The Streamlit framework is used to create an interactive web interface. Users can upload images or videos, select confidence thresholds, and view results in real time. The interface also displays metadata such as frame height, width, and frames per second (FPS), providing insights into system performance. Additionally, the system includes dynamic resizing of frames, allowing users to customize processing resolution. This feature helps in balancing performance and accuracy based on hardware capabilities. The confidence threshold slider enables users to control detection sensitivity, reducing false positives. The implementation also incorporates efficient memory handling and real-time updates. Frames are processed sequentially and displayed instantly, ensuring smooth visualization. The system continuously calculates FPS to evaluate processing speed. Overall, the implementation combines deep learning, real-time processing, and user interaction into a unified system capable of detecting potholes efficiently in practical scenarios.

VI. ALGORITHMS

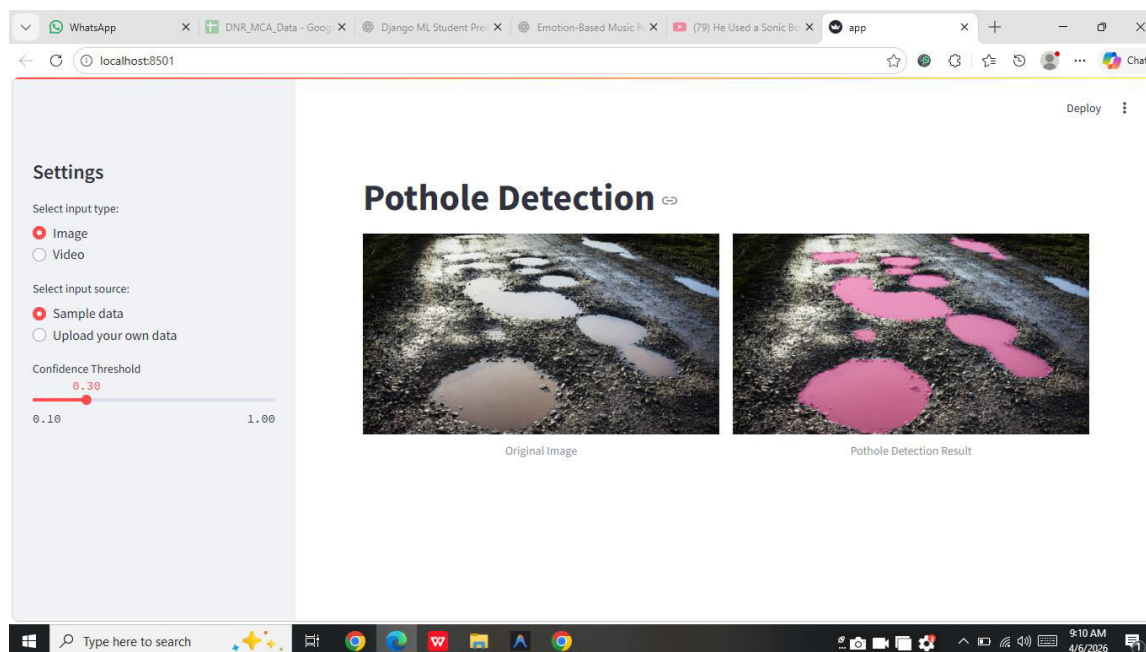
The proposed pothole detection system primarily relies on the YOLO (You Only Look Once) algorithm, along with tracking and preprocessing techniques. YOLO is a single-stage object detection algorithm that divides the input image into a grid and predicts bounding boxes and class probabilities simultaneously. Unlike traditional multi-stage detectors, YOLO processes the entire image in one pass, making it extremely fast and suitable for real-time applications. Advanced versions such as YOLOv8 further improve detection accuracy and segmentation capabilities. The algorithm workflow begins with input preprocessing, where images or video frames are resized and normalized. The YOLO model then extracts features using convolutional layers and predicts bounding boxes along with confidence scores. Non-Maximum Suppression (NMS) is applied to eliminate duplicate detections. In this project, segmentation-based YOLO is used, which enhances detection by identifying the exact shape of potholes. This improves accuracy in complex road conditions. The ByteTrack algorithm is used for object tracking in video streams. It assigns unique IDs to detected objects and tracks them across frames. This helps in maintaining consistency and avoiding repeated detections. Additionally, frame processing and visualization algorithms are implemented using OpenCV. These include frame resizing, color conversion, and annotation rendering. Overall, the combination of YOLO for detection and ByteTrack for tracking ensures high accuracy, real-time performance, and robustness in diverse environments.

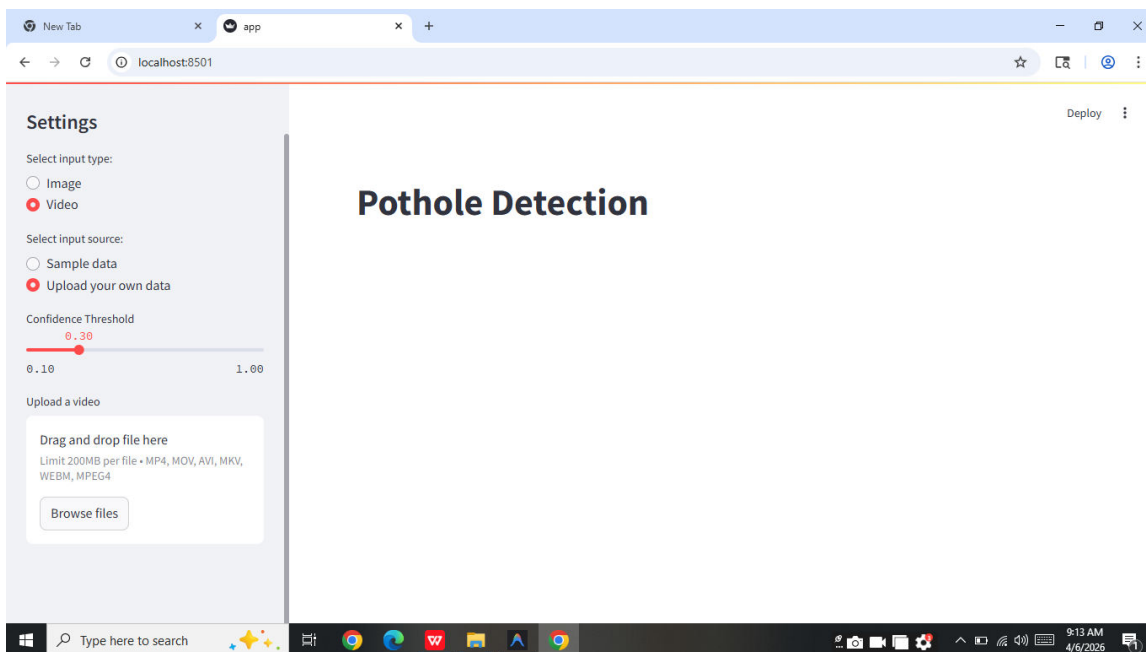
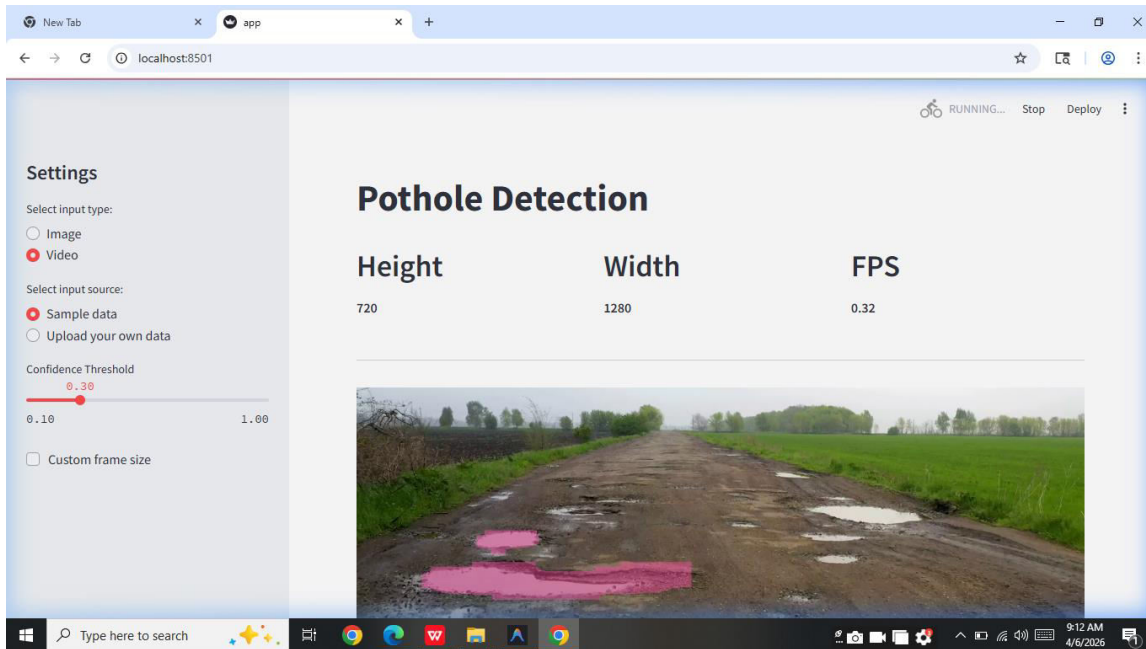
VII. SYSTEM DESIGN

The system design of the pothole detection system follows a modular architecture, ensuring scalability, flexibility, and efficiency. The system is divided into multiple components, including input handling, processing, detection, tracking, and output visualization. The first component is the input module, which allows users to select either image or video input. The system supports both sample data and user-uploaded files. For

video input, frames are extracted using OpenCV, while images are directly processed. The second component is the preprocessing module. In this stage, input data is resized, normalized, and converted into a suitable format for the YOLO model. Preprocessing ensures that the input data is consistent and improves detection performance. The core component is the detection module, which uses the YOLO deep learning model. The model processes input frames and detects potholes by generating bounding boxes or segmentation masks. YOLO's ability to perform detection in a single pass makes it highly efficient for real-time applications. Following detection, the tracking module is implemented using ByteTrack. This module assigns unique identifiers to detected potholes and tracks them across consecutive frames. Tracking improves system reliability by reducing duplicate detections and maintaining object continuity. The visualization module is responsible for displaying results to the user. Detected potholes are highlighted using masks or bounding boxes. The system also displays real-time metrics such as FPS, frame dimensions, and detection confidence. The user interface is developed using Streamlit, providing an interactive and user-friendly platform. Users can upload files, adjust confidence thresholds, and view results instantly. The interface is designed to be intuitive, making it accessible to non-technical users. The system also includes performance monitoring components, which calculate frame processing speed and update metrics dynamically. This helps in evaluating system efficiency. From a deployment perspective, the system can be integrated into smart transportation systems, surveillance cameras, or vehicle-mounted devices. The modular design allows easy integration with additional features such as GPS tagging and cloud storage. Overall, the system design ensures efficient data flow, real-time processing, and scalability, making it suitable for practical applications in road monitoring and maintenance.

SYSTEM DESIGN IMAGES





VIII. CONCLUSION

The pothole detection system developed in this project demonstrates the effectiveness of deep learning and computer vision techniques in addressing real-world infrastructure challenges. By leveraging advanced algorithms such as YOLO and ByteTrack, the system achieves accurate and real-time detection of potholes in both images and video streams. One of the key advantages of the proposed system is its ability to operate in real time, making it suitable for applications such as smart city monitoring, autonomous vehicles, and road maintenance systems. The use of segmentation-based detection further enhances accuracy by precisely identifying pothole regions. Compared to traditional methods, the proposed system significantly reduces manual effort and improves detection reliability. It overcomes limitations of sensor-based and image processing techniques by utilizing deep learning models capable of handling complex road conditions and varying lighting environments. The integration of a user-friendly interface using Streamlit makes the system accessible to a wide range of users. It allows easy interaction, real-time visualization, and customization of detection parameters. Experimental results and existing research indicate that deep learning-based approaches can achieve high accuracy and efficiency in pothole detection. Some advanced YOLO-based models have reported accuracy levels exceeding 90%, demonstrating their effectiveness in real-world scenarios. In conclusion, the proposed pothole detection system provides a scalable, efficient, and accurate solution for automated road monitoring. Future enhancements may include integration with GPS systems for location tracking, deployment on edge devices, and incorporation of additional sensors for improved accuracy.

REFERENCES

1. Bhavana, N. et al., "POT-YOLO: Real-Time Road Pothole Detection using YOLOv8," *IEEE Sensors Journal*, 2024.
2. Zanevych, Y. et al., "Evaluation of Pothole Detection Performance under Low-Light Conditions," *Sustainability*, 2024.
3. Zhong, J. et al., "YOLOv8 and Point Cloud Fusion for Pothole Detection," *Scientific Reports*, 2025.
4. Ravichandra Reddy, J. et al., "Real-Time Pothole Detection Using YOLO," *ICOFE Conference*, 2024.
5. Yurdakul, M. et al., "Enhanced YOLOv8 Model for Accurate Pothole Detection," 2025.
6. Rout, N.K. et al., "Improved Pothole Detection Using YOLOv7 and ESRGAN," 2023.
7. Kyem, B.A. et al., "Deep Learning for Pavement Distress Detection," 2024.
8. Tang, J. et al., "BsS-YOLO for Road Crack Detection," 2024.
9. Li, L. et al., "Road Pothole Detection using Mask R-CNN," *IEEE ITS*, 2024.
10. Younes, O. et al., "YOLO-based Object Detection Systems," 2024.
11. Redmon, J. et al., "You Only Look Once: Unified Object Detection," *IEEE CVPR*.

12. Bochkovskiy, A. et al., "YOLOv4: Optimal Speed and Accuracy," 2020.
13. Jocher, G. et al., "Ultralytics YOLOv8 Documentation," 2024.
14. Zhang, X. et al., "Deep Learning for Road Damage Detection," 2023.
15. Chen, L. et al., "Real-Time Object Detection for Smart Transportation," 2024.