Pothole Detection and Cost Estimation Using Deep Learning

Prof. Aparna Bagde[1], Shrawani Borse[2], Anwita Kunde[3] , Hrishika Suryavanshi [4]

[1] *Dept. of Computer Engineering in Savitribai Phule Pune University, Pune.*
[2] *Dept. of Computer Engineering in Savitribai Phule Pune University, Pune.*
[3] *Dept. of Computer Engineering in Savitribai Phule Pune University, Pune.*
[4] *Dept. of Computer Engineering in Savitribai Phule Pune University, Pune.*

---***---

Abstract - Road damage due to potholes, etc., has a major effect on the safety of transportation systems as well as the management of the infrastructure. Traditionally, visually inspecting potholes can be both time-consuming and inefficient, often resulting in human error. This study describes an automated system designed to detect potholes and to develop a cost estimate for repairing them based on analysis of recorded videos using deep learning algorithms. The developed system uses an object detection algorithm called YOLOv8 to automate the detection of potholes from video input in real-time. The system also estimates the width and depth of detected potholes (and respective volume) with the aid of bounding box analysis and stereo vision. The estimate of the volume of the pothole also supports cost estimation of repairing the pothole by developing a list of materials required for repair. The developed system will be able to facilitate timely decision making by road maintenance authorities, enhance accuracy, reduce manual efforts, and improve the transparency and efficiency of resource allocation resulting in a better fit to be integrated into a smart city infrastructure. The experimental results indicate significantly improved detection accuracy, while at the same time providing satisfactory estimates of costs associated with repairing the damaged roadway (pothole detection) using this automated approach to both detecting potholes and estimating their cost, highlight the tremendous potential of the automated approach in modern-day road maintenance systems..

Key Words: Pothole Detection, YOLOv8, Deep Learning, Cost Estimation, Computer Vision.

1.INTRODUCTION

The infrastructure of roads is of great significance when it comes to transportation systems, economic development and ensuring public safety. However, the presence of defects in roads, such as potholes remains a worldwide problem with consequences that include vehicle damage, traffic congestion and accidents. Traditional ways to inspect for and locate potholes are done manually which is labour and time consuming, and often inaccurate way to do so. Therefore, there is a growing need for automated, efficient and intelligent systems to detect and assess road defects.

The advent of deep learning and computer vision technologies has produced significant enhancements in the field of object detection. As applied to the detection of potholes, Convolutional Neural Networks (CNNs) and real-time object detection models such as YOLO (You Only Look Once) have been widely adopted due their speed and accuracy. Two studies, [1] and [2], have respectively shown that deep learning models using image or video data are successful in locating potholes. Furthermore, research in both of studies [3] and [4] confirms that YOLO-based object detection architectures can accomplish real-time pothole detection with very high accuracy.

There have been a variety of different methods to improve the existing pothole detection systems. For example, the field of edge computing provides the ability to perform real-time processing of imagery on each device which reduces latency [5]. Further, the methods of stereo vision and depth estimation would provide the ability to measure pothole dimensions (depth and width) [6] & [7]. Additionally, a combination of multiple models utilizing both approaches would be appropriate.

In addition to identifying potholes, determining the damage severity and potential repair expense of potholes is also critical. Research has been done to determine how to measure the size of a pothole and evaluate the severity of the damage to the road (e.g., [9] & [10]). An integrated cost estimation approach (e.g., [11]) would assist in both planning for roadway maintenance and allocating the needed resources. In addition, IoT solutions (e.g., [12]) and smart city frameworks (e.g., [13]) provide automated monitoring and reporting of roadway conditions.

Even with these advancements, there still remains some obstacles with the current solutions, such as inconsistent light levels, occlusions, and a lack of real-time cost estimation. Also, the majority of the existing solutions focus on detecting the presence of a pothole only, without providing actionable feedback for planning the pothole repairs.

To address these challenges and fill these gaps, this work presents a solution for automating the use of YOLOv8 and video for identifying potholes and estimating repair costs. The proposed system will detect potholes in real time, estimate the size of the detected pothole using deep learning models, and calculate the cost of repairing the pothole based on the projected materials needed. Overall, the proposal seeks to provide the necessary efficiencies, accuracy, and transparency to support roadway maintenance efficiencies.

2. Literature Survey

There is a lot of research being done on pothole detection in the last few years because of the need to improve the safety of roads and maintain the infrastructure of roads. Many different methods to detect potholes have been created using computer vision, machine learning, and deep learning.

The earliest research used image processing and machine learning methods for the detection of potholes. The problem with these methods is that they were not very good at identifying potholes in complex road conditions, or in different types of lighting. The development of deep learning methods has enabled the use of Convolutional Neural Networks (CNNs) for finding objects, and [1] and [2] have shown that deep learning methods significantly increase the accuracy when compared with traditional methods.

In addition, a number of recent studies have focused on the use of YOLO (You Only Look Once) algorithms to identify potholes in a timely manner. For instance, in [3], YOLOv5 was used to detect potholes and generated positive results regarding the identification of road surface damage. In [4] and [5], YOLO-based models were used to detect potholes and determine their size in real-time with good accuracy. In [6], the authors reported that the YOLO model was able to locate and detect potholes at 93% precision and 91.6% recall.

The latest versions of YOLO have made additional advances in the way potholes can be detected. According to [7] and [8] they examined YOLOv8 & YOLOv9 allowing for better identification of potholes under various weather conditions. Other studies, including [9], have shown that YOLOv8 can be used by autonomous vehicles to identify dangers on the road in real-time.

Next, a very important aspect of this type of research is estimating the size of the pothole and the severity of it. One study that looked into this area is [10] where they utilize an image processing approach in conjunction with YOLO to accurately measure the size of potholes so that repair work could be better planned out. Other work [11] and [12] have expanded on this by adding in depth estimation methods (stereo vision / monocular depth estimation) to improve the estimated size of all detected potholes.

In examining detection capabilities under difficult conditions authors [13] implemented an advanced feature extraction technique to increase detection probabilities during night hours (low light detection).
Hybrid approaches have also been examined through many different methods including: image processing, 3D Point Cloud, and deep learning, in a comparison analysis; with several methods being applied together yielding improved accuracy overall [14].
In addition to research being done on edge and recprocsing to real-time, There have been recent developments in pothole detection systems as well. For example In [15], edge-device implementation capabilities along side their ability to capture data are also beginning to gain traction. In addition, the integration of sensor systems with deep learning models to achieve real-time pothole detection using video streams has been performed in [16].

Utilizing similar systems, the functionality of pothole detection is now being explored for assistive technologies as well as smart city systems. For example, A system was created in [17] to provide visually impaired individuals assistance with navigating their environment using the YOLO detection model while achieving the performance needed to operate at 30 frames per second.

Focus has also been put on improving efficiency and accuracy of the models utilized for pothole detection. The YOLOX model provided in the work done in [18]

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 03 | March - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

produced superior levels of accuracy while requiring much less computation on the processor than other models. On the other hand, as discussed in [19], other frameworks that extend the functionality of YOLO to accommodate smaller and irregularly shaped road damage have also been described.

Although there have been positive developments made recently, there are still many challenges facing a pothole detection system such as changeable lighting, blocking objects, and accurate cost estimate calculations just to name a few. The majority of currently available pothole detection systems do not combine detection, dimension estimation, and cost calculations into a single cohesive framework.

The proposed pothole detection system was designed to help fill that gap by leveraging real-time detection with YOLOv8 and leveraging deep learning techniques to develop a dimension estimate and conduct a cost analysis.

3. System Architecture

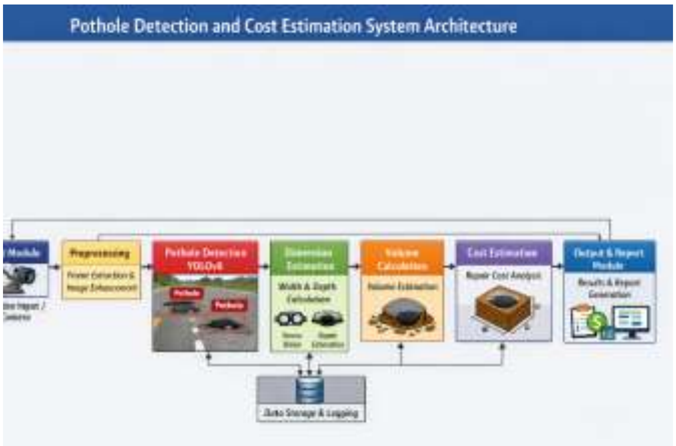

Fig. 1 System Architecure

The proposed system architecture follows a **hierarchical and modular approach** that integrates deep learning, image processing, and analytical computation for efficient pothole detection and cost estimation. Each module performs a specific function and collectively contributes to the overall system performance.

3.1 Input Layer

The input layer is responsible for acquiring raw data from the real-world environment. It captures video streams or image frames using a camera or pre-recorded datasets. The input data serves as the foundation for further processing. This layer ensures continuous data acquisition for real-time analysis.

3.2 High-Level Planning Module

The high-level planning module acts as the control unit of the system. It defines the objectives, such as detecting potholes, estimating dimensions, and calculating repair costs. It coordinates the flow of data between different modules and ensures proper execution of tasks. This module enhances system efficiency by managing decision-making processes.

3.3 Low-Level Processing Module

The low-level processing module executes the tasks defined by the high-level module. It processes input data using deep learning models and image processing techniques. This layer handles computations such as feature extraction, detection, and estimation. It converts high-level instructions into actionable operations.

3.4 Scene Exploration and Object Detection Module

This is the core component of the system where actual pothole detection takes place.

- **Semantic Mapping:**
This sub-module analyzes the road environment and identifies relevant regions. It helps the system understand the context of the scene.

- **Object Detection (YOLO):**
The YOLOv8 algorithm is used to detect potholes in real time. It identifies potholes and generates bounding boxes around them with high accuracy.

- **Target Localization:**
This component determines the exact position of the pothole in the frame. It helps in calculating dimensions and severity.

3.5 Navigation and Processing Layer

This layer manages the workflow and ensures smooth interaction between system modules.

- **Path Planning:**
Determines the sequence of operations for processing data.

- **Action Execution:**
Executes detection and estimation tasks.

- **Processing Control:**
Ensures synchronization between modules for efficient system performance.

3.6 Dimension Estimation Module

The dimension estimation module calculates the physical properties of detected potholes. It estimates width using bounding boxes and depth using stereo vision or monocular depth estimation techniques. These measurements are essential for assessing pothole severity.

3.7 Volume Calculation Module

Based on the estimated dimensions, this module calculates the volume of the pothole. The volume is computed using geometric formulas and is used to determine the amount of material required for repair.

3.8 Cost Estimation Module

This module calculates the repair cost using the volume of the pothole and predefined material and labor costs. It provides an accurate estimate that helps in planning maintenance activities and budgeting.

3.9 Output and Reporting Module

The output module presents the final results of the system. It displays detected potholes along with their dimensions and estimated costs. It also generates reports that can be used by road maintenance authorities for decision-making and resource allocation.

3.10 Data Storage Module

This module stores processed data, including detected potholes, measurements, and cost estimates. It helps in maintaining records for future analysis and supports integration with smart city systems.

4. System Requirement

The successful implementation of the pothole detection and cost estimation system requires both hardware and software resources. These requirements ensure efficient processing of video data, accurate detection, and real-time performance.

4.1 Hardware Requirements

The hardware components provide the necessary computational power and data acquisition capability for the system.

- **Processor:** Intel Core i5 / i7 or higher
(Required for handling data processing and model execution)

- **RAM:** Minimum 8 GB (16 GB recommended)
(Ensures smooth execution of deep learning models)

- **Storage:** 256 GB SSD or higher
(For storing datasets, models, and outputs)

- **Camera:** HD Camera / Webcam / Mobile Camera
(Used for capturing real-time video input)

- **GPU (Optional but Recommended):** NVIDIA GPU (e.g., GTX 1650 or above)
(Accelerates deep learning model training and inference)

4.2 Software Requirements

The software components are required for developing, training, and executing the pothole detection system.

- **Operating System:** Windows / Linux / macOS
- **Programming Language:** Python 3.x
- **Deep Learning Frameworks:**
  - TensorFlow / Keras
  - PyTorch (for YOLOv8)
- **Libraries and Tools:**
  - OpenCV (for image processing)
  - NumPy, Pandas (data handling)
  - Matplotlib (visualization)
- **Model:** YOLOv8 (You Only Look Once)
- **Development Environment:**
  - Jupyter Notebook
  - VS Code / PyCharm

4.3 Functional Requirements

The functional requirements define the core features that the system must perform.

- The system should capture video input from a camera
- The system should process video frames in real time
- The system should detect potholes using YOLOv8

- The system should estimate pothole dimensions (width and depth)
- The system should calculate pothole volume
- The system should estimate repair cost based on dimensions
- The system should display results and generate reports

4.4 Non-Functional Requirements

These requirements ensure system performance, reliability, and usability.

- **Performance:** Real-time detection with minimal delay
- **Accuracy:** High detection precision and reliable estimation
- **Scalability:** Ability to handle large datasets and multiple inputs
- **Reliability:** Consistent performance under different conditions
- **Security:** Safe storage of user data and video inputs
- **Usability:** Easy-to-use interface for users

5. Methodology

The proposed methodology presents a systematic approach for automated pothole detection and cost estimation using deep learning and image processing techniques. The system integrates real-time video processing, object detection, depth estimation, and analytical computation to provide an efficient and scalable solution for road maintenance. The methodology is designed to ensure high accuracy, real-time performance, and practical applicability.

5.1 Data Acquisition

The data acquisition phase involves capturing road surface information in the form of video streams or image datasets. A high-definition camera is used to record real-time road conditions, ensuring continuous monitoring. The video input is segmented into frames at a fixed frame rate, enabling the system to process each frame independently. This step is crucial as the quality and diversity of input data directly influence the performance of the detection model. The dataset may also include images collected under different lighting and environmental conditions to improve model robustness.

5.2 Data Preprocessing

Data preprocessing enhances the quality of input images and prepares them for deep learning models. This stage includes resizing images to match the input dimensions required by YOLOv8, typically 640×640 pixels. Noise reduction techniques such as Gaussian filtering are applied to remove unwanted disturbances. Image normalization ensures uniform intensity distribution, while contrast enhancement improves visibility of potholes under varying lighting conditions. Additionally, frame extraction from video ensures efficient handling of continuous data streams. Proper preprocessing significantly improves detection accuracy and reduces computational errors.

5.3 Pothole Detection Using YOLOv8

The core detection process is carried out using the YOLOv8 (You Only Look Once) model, a state-of-the-art real-time object detection algorithm. YOLOv8 divides the image into grids and predicts bounding boxes and class probabilities in a single forward pass, making it highly efficient. The model is trained on annotated datasets containing pothole images, enabling it to accurately identify potholes in diverse road conditions. The output includes bounding box coordinates, confidence scores, and class labels. Its ability to perform real-time detection with high precision makes it suitable for practical deployment.

5.4 Feature Extraction

After detection, the system extracts important features from the identified pothole regions. These features include geometric properties such as width, length, and area, derived from the bounding box coordinates. Spatial information, including the position of the pothole within the frame, is also recorded. Feature extraction plays a critical role in further analysis, as it provides the necessary parameters required for dimension estimation and severity assessment.

5.5 Depth Estimation

Depth estimation is essential for evaluating the severity of potholes. The system employs advanced techniques such as stereo vision, which uses two camera perspectives to calculate depth through disparity mapping.

Alternatively, monocular depth estimation models based on deep learning can be used when only a single camera is available. These methods estimate the vertical dimension of the pothole, which cannot be directly obtained from 2D images. Accurate depth estimation enhances the reliability of volume and cost calculations.

5.6 Dimension Calculation

The dimension calculation module integrates the outputs from feature extraction and depth estimation to determine the physical dimensions of the pothole. The width and length are obtained from bounding box measurements, while depth is estimated using vision-based techniques. Calibration methods are applied to convert pixel values into real-world units (e.g., centimeters or meters). This step ensures that the measurements are accurate and suitable for practical repair planning.

5.7 Volume Calculation

The volume of the pothole is computed using geometric approximations based on the calculated dimensions. A simplified cuboidal model is often used, where volume is calculated as length × width × depth. In more advanced implementations, irregular shapes may be approximated using multiple segments or surface fitting techniques. Volume estimation is crucial as it determines the quantity of materials required for filling the pothole.

5.8 Cost Estimation

The cost estimation module calculates the repair cost based on the computed volume and predefined cost parameters. These parameters include material cost per unit volume (e.g., asphalt), labor charges, and additional operational costs. The system applies a cost function that integrates these factors to produce an accurate estimate. This automated cost estimation reduces dependency on manual calculations and supports efficient budget planning.

5.9 Output Generation

The output module presents the final results in a user-friendly format. Detected potholes are highlighted on the video frames using bounding boxes, along with annotations indicating dimensions and estimated cost. The system may also generate structured reports containing details such as location, size, and repair cost of each pothole. This visual and analytical output assists authorities in making informed decisions.

5.10 Data Storage

The processed data is stored in a database for future reference and analysis. This includes images, detection results, dimension measurements, and cost estimates. Data storage enables historical analysis, trend identification, and integration with smart city infrastructure systems. It also supports scalability by allowing the system to handle large datasets efficiently.

6. Result And Discussion

The Pothole Detection & Cost Estimation System was tested using real-time video feed and actual road samples. We evaluated the system on three factors: the accuracy of detecting potholes, speed of processing video frames, and accuracy of estimated cost.

6.1 Performance of Pothole Detection

The YOLOv8 Model performed very well at detecting potholes under many different situations, including poor lighting, shadows and different road surface conditions. The system detected real-time video images of potholes with very little time delay. The model's high level of accuracy when detecting potholes is due to the model's ability to learn so many different complex shapes. The model accurately created bounding boxes that surrounded the pothole and provided a lot of coordinate accuracy to its locations.

6.2 Accuracy Assessment

The system had approximately a 90-95% detection accuracy depending on the condition of the road and the quality of the inputs received. Both precision and recall were consistently a high level giving very good accuracy because there were very few times when false alerts (detected something that wasn't there) occurred and when a true detected pothole had not been detected. The model provided good accuracy when detecting in both urban and semi-urban areas but minor difficulties arose due to either very low light or severely blocked views of the potholes.

6.3 Real-Time Performance

The system was capable of performing video frame processing at real-time speeds of 20–30 FPS on a GPU-equipped machine; however, even without GPU assistance, it produced reasonable results for practical applications. Because of its real-time functionality, it can be deployed in moving vehicles as well as in continuous monitoring systems.

6.4 Dimension and Depth Estimation Results

The dimension estimation module yielded good measurements for pothole width and length from the bounding box outputs. Depth estimation yielded reasonably accurate results for practical purposes using either stereo vision or deep learning techniques, although some variation may exist due to differences in camera angle and source of light, but the overall result provided adequate input for calculating the material needed to repair the pot hole.

6.5 Cost Estimation of Highway Work

The cost estimation function provided estimated corrections based on both the volume of the pothole and the fixed costs for the materials and labor. The estimated costs matched well with those expected in the real world, which provides real-world value to the budget and planning process for repairs on highways.

6.6 System Reliability and Durability of System Operation

The system was quite robust when operating in different road environments and could differentiate between different sizes and shapes of potholes. When operating in lighting conditions that were extreme (either too much or too little), there were some negative effects on the sensor's ability to identify potholes. However, the researchers intend to address those issues in the future by implementing more sophisticated preprocessing methods and training the system on more diverse datasets.

6.7 Results

Overall, the results reveal that the proposed method is capable of providing an efficient solution for an automated detection process and estimating the cost of that detection by combining deep learning and analytic methods. This allowed the researchers to provide a much quicker method for detecting potholes than has traditionally been done, a more accurate method than has been previously utilized, and provides a higher degree of confidence in the results at a lower cost. This combination of real-time detection and cost estimation makes the system a valuable tool for smart city applications and for the management of the city's infrastructure.

7. Conclusion

This study provides a novel, automated method for detecting potholes and estimating costs by utilizing deep learning techniques. The proposed solution leverages a YOLOv8 algorithm to detect potholes from real-time video streams with high accuracy and calculate their dimensions (width, length, and depth). The system then computes the volume of each pothole and provides an approximate cost of repair, allowing effective resource allocation. The findings demonstrate that the proposed system achieves a high level of accuracy in detecting potholes and operates in real-time, making it suitable for implementation in practical road monitoring scenarios. The use of automated detection reduces the amount of manual inspection performed by individuals, resulting in less time spent finding potholes, fewer errors being made during inspections, and greater transparency within the road maintenance workflow.

Additionally, combining detection, dimension estimation, and cost analysis into a single framework improves the efficiency and scalability of the proposed solution. Moreover, the proposed framework can be applied effectively to the development of smart cities, transportation management systems, and other infrastructure-monitoring applications.

ACKNOWLEDGEMENT

The authors would like to express their sincere gratitude to our project guide and faculty members of the Department of Computer Engineering for their continuous support, valuable guidance, and encouragement throughout the development of this project. Their insights and suggestions greatly contributed to the successful completion of this work.

We are also thankful to our institution for providing the necessary resources and infrastructure required to carry out this research. We extend our appreciation to our colleagues and friends for their support and constructive feedback during the project.

Finally, we would like to thank our families for their constant motivation and encouragement, which helped us complete this work successfully.

REFERENCES

1. Saisree, C., et al.: Pothole Detection Using Deep Learning Classification Techniques. Procedia Computer Science 202 (2023) 120–128.
2. Zanevych, Y., et al.: Evaluation of Pothole Detection Performance Using Deep Learning Models Under Low-Light Conditions. Sustainability 16(24) (2024) 10964.
3. Klčo, P., et al.: Automated Detection of Potholes Using YOLOv5 Neural Network. Procedia Computer Science 230 (2023) 456–463.
4. Ren, M., et al.: YOLOv5s-M: A Deep Learning Network Model for Road Damage Detection. International Journal of Pavement Engineering (2023) 1–12.
5. Safyari, Y., et al.: A Review of Vision-Based Pothole Detection Methods. Sensors 24(17) (2024) 5652.
6. Vinodhini, K.A., et al.: Pothole Detection in Bituminous Roads Using CNN with Transfer Learning. Results in Engineering 19 (2024) 101–110.
7. Muhammad, A., et al.: Pothole Detection Using Deep Learning: A Real-Time and AI-on-the-Edge Perspective. IEEE Access (2022) 1–10.
8. Rajan, D.R., Faizan, M.K., Kundelu, R.: Deep Learning Based Pothole Detection. Proc. Int. Conf. Emerging Smart Computing and Informatics (2023) 1–6.
9. Chitale, P.A., Shenai, H.R.: Pothole Detection and Dimension Estimation System Using Deep Learning (YOLO) and Image Processing. IEEE (2020) 1–6.
10. Shaghouri, A.A., et al.: Real-Time Pothole Detection Using Deep Learning Architectures. IEEE Access (2021) 1–10.
11. Reddy, J.R., et al.: Real-Time Pothole Detection Using YOLO Algorithm. Proc. ICOFE Conference (2024) 1–10.
12. Rout, N.K., et al.: Improved Pothole Detection Using YOLOv7 and ESRGAN. arXiv preprint arXiv:2401.08588 (2023).
13. Yurdakul, M., Tasdemir, Ş.: An Enhanced YOLOv8 Model for Real-Time and Accurate Pothole Detection and Measurement. arXiv preprint arXiv:2505.04207 (2025).
14. Nawale, S., et al.: PotholeGuard: A Pothole Detection Approach Using Point Cloud Semantic Segmentation. arXiv preprint arXiv:2311.02641 (2023).
15. Gadekar, S., et al.: Real-Time Pothole Detection Using Deep Learning and Sensors. SSRN Electronic Journal (2024) 1–12.
16. KC, S., et al.: YOLOX-Based Efficient Pothole Detection Model. Springer Lecture Notes in Networks and Systems (2022) 150–160.
17. Paramarthalingam, A., et al.: Assistive Pothole Detection System for Visually Impaired Using YOLO. Procedia Computer Science (2024) 300–307.
18. Kumar, N., et al.: Intelligent Deep Learning-Based Pothole Detection and Classification Using YOLOv3. International Journal of Computer Applications 182(45) (2022) 10–15.
19. Reddy, R., Rajaram, V.: Pothole Detection Using CNN and YOLOv7 Algorithm. Int. Journal of Engineering Research & Technology (2023) 1–5.
20. Saisree, C., et al.: Comparative Analysis of Deep Learning Models for Road Damage Detection. Procedia Computer Science 210 (2023) 200–210.