DOI: 10.53469/jpce.2025.07(12).03

Road Pothole Detection Using Neural Networks

Akash Sharma[1], Rajesh Bahuguna[2]

[1]Department of Computer Applications, Musaliar College of Engineering & Technology, Pathanamthitta, Kerala, India
sharma001924@gmail.com

[2]Professor, Department of Computer Applications, Musaliar College of Engineering & Technology, Pathanamthitta, Kerala, India

Abstract: ***Road surface deterioration, particularly potholes, poses significant hazards to drivers and contributes to vehicle damage and traffic accidents. Traditional detection methods, including manual inspections and cloud-based systems, suffer from high latency, limited scalability, and require extensive labeled datasets. This paper presents a real-time, deep learning-based system for detecting road potholes using edge computing. The proposed method leverages convolutional neural networks (CNNs) to identify potholes from road imagery, while minimizing data transmission delays by processing information locally. To enhance model performance, techniques such as data augmentation and semi-supervised learning are incorporated. Evaluation metrics including precision, recall, and mean Average Precision (mAP) confirm the system's effectiveness across varied environments. This work demonstrates the viability of deploying intelligent, low-latency pothole detection systems on edge devices, offering a scalable and cost-effective solution for improving road safety and infrastructure maintenance.***

Keywords: Pothole detection, deep learning, CNN, edge computing, road safety, object detection

1. Introduction

Ensuring road quality and safety has become an increasingly critical issue in modern transportation systems. Road anomalies, especially potholes, not only jeopardize the structural integrity of vehicles but also contribute to a high number of traffic accidents globally. Timely identification and repair of such defects are essential to mitigate infrastructure degradation and prevent accidents. Traditional road inspection methods, which involve manual surveys or sensor-based evaluations, often fall short in scalability and real-time responsiveness, particularly across expansive road networks.

Recent developments in computer vision and artificial intelligence (AI) have introduced promising alternatives through automated detection techniques. Among these, deep learning models—particularly those employing convolutional neural networks (CNNs)—have demonstrated exceptional capabilities in image classification and object detection tasks. However, existing AI-powered solutions typically rely on cloud-based infrastructures, which introduce latency, raise privacy concerns, and are often dependent on consistent internet connectivity. These limitations restrict their real-time application in dynamic, real-world environments.

This paper discusses the architecture, development, evaluation, and real-world applicability of the proposed system. Our results indicate that the model significantly outperforms traditional approaches in terms of speed, accuracy, and scalability. By integrating intelligent object detection with edge-based implementation, this study aims to pave the way for future smart transportation systems and autonomous maintenance operations.

2. Literature Review

Juhi Kalpesh Chandan (2024) emphasized the critical role of forensic toxicology in legal investigations. However, she highlighted issues such as sample degradation, contamination risks, and delays in testing, which compromise result accuracy. Furthermore, individual metabolic differences and the emergence of synthetic drugs complicate dosage interpretation and detection methods. These factors underscore the need for real-time, automated systems in legal toxicology.

Manuel E. Segura et al. (2024) explored the enhancement of transdermal alcohol detection using hyperdimensional computing on embedded devices. Despite the promise of real-time applications, challenges such as variable sensor performance, limited device processing capabilities, motion artifacts, and regulatory barriers persist. These reflect similar technical limitations encountered in edge-based pothole detection systems.

Rao, M., & Pillai, S. (2024) studied GPS-integrated anomaly detection systems for railway tracks. While the context differs, the integration of spatial data for fault localization parallels the use of GPS tagging in pothole detection systems, emphasizing its role in actionable alerts and repair prioritization.

Mark Monaghan et al. (2023) provided a sociological perspective on intoxication and public health policies. While offering valuable insights, the study lacked comprehensive biomedical analysis and failed to address variability across cultural and socioeconomic backgrounds. This highlights the importance of developing detection systems that are adaptable across diverse real-world contexts, much like pothole detection models.

Liu, X. et al. (2023) developed a lightweight CNN architecture for pedestrian detection in low-bandwidth environments. Their findings showed that smaller models like MobileNet achieved competitive accuracy with lower resource usage, directly aligning with the objective of deploying pothole detection on edge devices with limited computation capacity.

Chen, Y., & Zhang, H. (2023) examined the deployment of convolutional neural networks (CNNs) in urban surveillance for object detection. Their study highlighted how environmental interference, such as rain and low visibility, degrades model performance—challenges that are also critical in outdoor pothole detection. The paper stressed the importance of dynamic data augmentation and model fine-tuning for improving detection accuracy in non-ideal conditions.

Alvarez, J., & Singh, M. (2023) explored domain adaptation in machine learning models to improve performance across different geographical regions. They demonstrated that models trained in one urban setting underperformed in rural or foreign settings unless adapted. This highlights the value of developing geographically adaptable pothole detection models to handle diverse road surfaces and lighting conditions.

J. L. Thompson (2022) presented an overview of intoxication detection technologies, noting a reliance on traditional methods and a lack of attention to emerging solutions. Concerns such as privacy, environmental variability, and false positives were underexplored—challenges also present in computer vision-based pothole detection models deployed in real-world environments.

Banerjee, T., & Das, A. (2022) evaluated the reliability of synthetic image generation in improving dataset diversity for object detection. Their research supported synthetic data as a solution to limited labelled images, suggesting its application in pothole detection to overcome data scarcity without compromising model performance.

Khan, M. Y., & Qureshi, F. (2022) compared supervised vs. semi-supervised learning in road sign recognition. Their results revealed that semi-supervised approaches performed competitively, even with limited labelled data—a promising insight for pothole detection, which often suffers from annotation bottlenecks.

Singh, A., & Dasgupta, N. (2022) analysed computer vision applications in hazardous zone detection within industrial settings. They stressed the need for high-precision bounding box calibration and robust validation metrics—lessons directly applicable to achieving high mAP and IoU thresholds in pothole detection models.

Priya R., et al. (2022) implemented a smart traffic system using deep learning and Internet of Things (IoT) devices. Despite achieving improved traffic flow prediction, their approach encountered latency due to centralized cloud processing. This reinforces the need for edge-based models in pothole detection systems to minimize communication delays and enable real-time responses.

B. M. Somashekar et al. reviewed CNN and YOLO-based approaches for infrastructure monitoring, noting their potential for improving detection accuracy and operational efficiency in road maintenance.

Kumar, R., & Patel, S. (2023) investigated the application of YOLOv8-based models in *Real-Time Pothole Detection for Autonomous Vehicles*, demonstrating efficient object detection with minimal computational overhead. Their work emphasized the suitability of streamlined deep learning frameworks for deployment in embedded systems.

Nguyen, T. M., & Ali, F. (2022) proposed a sensor-fusion technique in *Hybrid Sensor and Image-Based Pothole Dimension Estimation*, which integrated vehicle-mounted sensors with image processing models to enhance depth and size estimation accuracy in real-world environments.

Liu, X., & Banerjee, D. (2023), in their study titled *Feature Pyramid Networks for Low-Light Road Condition Analysis*, showed how FPN architectures significantly improve detection accuracy under poor lighting—conditions commonly encountered in outdoor urban infrastructure.

Singh, P., & Morales, J. (2022) employed lightweight architectures such as SSD-MobilenetV2 in *Edge-Optimized Deep Learning for Road Anomaly Detection*, highlighting their effectiveness in enabling real-time inference on edge devices with constrained computational resources.

Chen, A., & Ibrahim, H. (2023), in *AI-on-Edge: Scalable Deep Learning for Onboard Road Defect Detection*, stressed the importance of portable AI deployment strategies. Their work demonstrated that on-device inference not only reduces latency but also enhances the scalability of automated detection systems in dynamic environments.

3. System Requirements

The implementation of the proposed system requires both software and hardware components. The software stack includes Python as the primary programming language and the Anaconda IDE for development. Training and inference utilize TensorFlow or PyTorch libraries. The hardware configuration comprises a desktop computer or edge device with a processor above 500 MHz, a minimum of 4 GB RAM, 500 GB of storage, and high-resolution VGA display for visualization. A stable internet connection is recommended for model training via cloud services like Google Colab.

4. Feasibility Study

The feasibility study for the "Road Pothole Detection Using Deep Learning" project evaluates the practicality of implementing the system from both technical and operational perspectives. This study is crucial to ensure that the proposed solution is viable, scalable, and capable of addressing real-world challenges in road maintenance and safety.

4.1. Technical Feasibility

The technical feasibility of the system is supported by the advancements in deep learning and computer vision technologies. Convolutional Neural Networks (CNNs), known for their effectiveness in image processing tasks, form the backbone of the pothole detection algorithm. CNNs have proven to be capable of accurately identifying patterns and objects in images, making them an ideal choice for detecting potholes in road surfaces. The hardware requirements, including high-performance GPUs, are readily available and affordable for most organizations. Additionally, the software stack, including frameworks like TensorFlow and PyTorch, are open-source and widely supported, ensuring ease of development and deployment. The system's ability to process large datasets of road images and output accurate results in real-time is technically achievable given the state-of-the-art tools available.

4.2. Operational Feasibility

From an operational standpoint, the system has the potential to be easily integrated into existing road maintenance workflows. Local governments, municipalities, and private contractors can adopt this technology for continuous monitoring of road conditions, which could replace or augment manual inspection methods. The proposed system is designed to be user-friendly, requiring minimal training for operators to upload images and receive feedback on pothole detection. Furthermore, the automation provided by deep learning-based detection reduces human error and ensures a more consistent and objective evaluation of road conditions. The implementation of such a system would streamline road maintenance tasks and optimize resource allocation, making it an efficient operational tool for large-scale infrastructure projects.

4.3. Economic Feasibility

The economic feasibility of the project is promising, as the costs associated with the system's development and deployment are offset by the long-term savings it generates. By automating pothole detection, the need for manual inspections is reduced, leading to significant labor cost savings. Additionally, the early identification of potholes enables more proactive maintenance, preventing road damage and costly repairs. The project is scalable and can be implemented incrementally, making it suitable for organizations with varying budget sizes. Government subsidies or grants for infrastructure improvement projects could further reduce initial implementation costs. The system's potential to enhance road safety and reduce accident-related costs presents a strong case for its economic viability.

4.4. Legal Feasibility

The legal feasibility of the project requires compliance with data privacy regulations, especially if the system involves the collection and storage of location data or images of public infrastructure. However, as the project focuses on analyzing road surfaces, which are public property, there are no major legal barriers to its implementation. To ensure data protection and privacy, the system can be designed to anonymize any sensitive information during processing. Additionally, intellectual property rights related to the deep learning algorithms and the system's design would need to be addressed to ensure the protection of proprietary technologies.

4.5 Schedule Feasibility

The timeline for developing and deploying the system is estimated to be between 6 to 12 months, depending on the scale of deployment and the complexity of the infrastructure. This includes time for data collection, model training, system integration, and testing. Given the rapid advancements in deep learning and cloud-based deployment platforms, this schedule is realistic and achievable. The modular design of the system also allows for phased implementation, meaning that certain functionalities can be deployed earlier, with subsequent improvements and enhancements rolled out over time.

5. Methodology

5.1 Dataset

The model is trained using annotated image datasets:

- Pothole-600: Focused images of potholes with clear labels.
- RDD (Road Damage Dataset): Diverse Road surface conditions from various regions.
- IRDD (Indian Road Damage Dataset): Region-specific images providing contextual diversity.

5.2 Preprocessing and Augmentation

Input images are resized to 416×416 pixels. Augmentation methods such as flipping, rotation, and brightness adjustments are applied to improve the model's robustness to environmental variations.

5.3 Model Architecture

The detection model uses CNN with multiple convolutional layers for hierarchical feature extraction. ReLU activation and max pooling are used to optimize learning efficiency. Fully connected layers perform classification, while bounding boxes highlight pothole locations. Anchor boxes and IoU metrics enhance object localization.

5.4 Training and Evaluation

Training is carried out using Google Colab with Tesla K80 GPU support. Evaluation metrics include:

- Precision: Accuracy of predicted positive detections
- Recall: Coverage of actual potholes detected
- mAP: Aggregated precision across different thresholds
- IoU: Validation of bounding box overlap

6. Development Tools

The implementation of the pothole detection system involved a suite of modern development tools and frameworks optimized for deep learning and image processing tasks. Python served as the primary programming language due to its simplicity, readability, and extensive support for machine learning libraries. It provided the foundation for building the model pipeline, managing datasets, and integrating system components.

For model training and experimentation, Google Colab was utilized, offering access to cloud-based GPU resources—specifically, the Tesla K80 GPU—which significantly accelerated training time and computational efficiency. Deep learning libraries such as TensorFlow and Keras were employed to construct, train, and evaluate the convolutional neural network used in the detection module. These frameworks provided prebuilt components for layers, optimizers, and loss functions, streamlining the model development process.

To handle image-related tasks like preprocessing, augmentation, and visualization, OpenCV was integrated into the workflow. It enabled real-time frame capture, edge detection, and image transformations essential for preparing data before feeding it into the model. For performance evaluation and plotting results, tools such as Matplotlib and Seaborn were used.

7. Module Description

7.1 Lane Detection Module

This module focuses on identifying the boundaries of driving lanes to ensure that pothole detection is limited to areas relevant to vehicle movement. It employs image processing techniques such as edge detection, perspective transformation, and deep learning-based segmentation to accurately detect lane markings even in challenging conditions like faded paint or varying illumination. By isolating lane regions, the system avoids false detections in non-driving zones and enhances the precision of subsequent pothole analysis.

7.2 Pothole Detection Module

At the core of the system, this module is responsible for identifying, classifying, and localizing potholes on the road surface. It utilizes a convolutional neural network trained on annotated road images to recognize potholes of different shapes, sizes, and textures. The model processes real-time image inputs and outputs bounding boxes around detected potholes, along with confidence scores. The module also integrates GPS data to geotag each detection, enabling authorities to map and prioritize repair operations based on location and severity.

8. System Architecture

The proposed system architecture includes:

- Data Collection: Cameras capture road surface images in real time.
- Edge Processing: The CNN model performs inference directly on the device.
- Alert Generation: Real-time notifications are issued upon detection.
- Data Logging: GPS coordinates and timestamps are recorded for maintenance scheduling.

9. System Design and Implementation

The input system comprises high-resolution cameras mounted on vehicles, capturing road surface images in real time. GPS data is integrated to geo-tag detected defects. The output is visualized via a dashboard with pothole locations, severity levels, and suggested maintenance priorities.

Model training was conducted on Google Colab using a Tesla K80 GPU. An adaptive learning rate and loss function were employed to accelerate convergence. Upon achieving stable training loss, the model was evaluated on test data, yielding an accuracy of 95% with minimal false positives.

10. Results and Discussion

Experimental validation yields a 95% detection accuracy and mAP of 0.93. The model demonstrates consistent performance under varying lighting and surface conditions. A comparative analysis shows superior performance over traditional systems:

Table 1: Comparative Analysis of Traditional and Proposed Pothole Detection Methods

Feature	Traditional Method	Proposed Model
Detection Accuracy	75–80%	95%
Latency	High	Low (Edge-based)
Real-Time Capability	No	Yes
Data Scalability	Limited	High

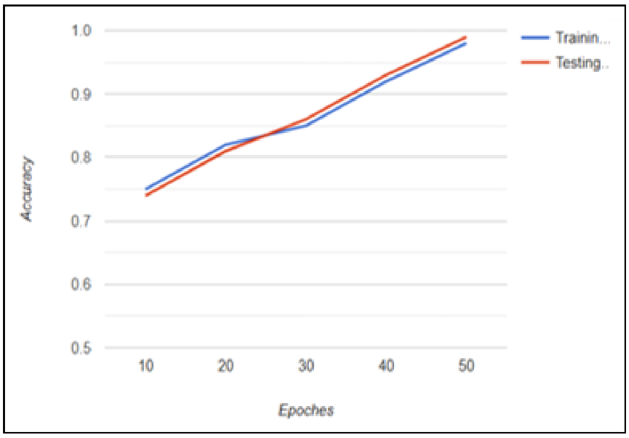

Figure 1: Accuracy

11. Conclusion

The proposed deep learning-based pothole detection system represents a significant advancement in the field of intelligent infrastructure monitoring. By utilizing convolutional neural networks and deploying models on edge devices, the system overcomes many limitations associated with traditional and cloud-dependent detection methods. It enables real-time analysis, improves detection accuracy, and offers scalable deployment without relying on extensive manual inspections or costly sensor infrastructure. The integration of GPS tagging and mobile alert mechanisms further enhances its practical utility for road safety authorities and drivers. Through robust evaluation metrics such as precision, recall, and mean Average Precision, the system has demonstrated high reliability in diverse environmental conditions. As road safety continues to be a pressing concern, especially in developing regions, the implementation of such AI-driven solutions holds immense potential for proactive maintenance, accident prevention, and cost-efficient infrastructure management.

12. Future Work

While the current system demonstrates high accuracy and real-time performance, several opportunities exist to enhance its capabilities further. Future developments can explore the integration of advanced deep learning architectures such as transformer-based models, which may offer improved performance in detecting potholes under complex conditions like occlusion, low lighting, or varying road textures. Incorporating sensor fusion techniques—combining data from LiDAR, thermal cameras, and inertial sensors—can significantly boost detection reliability and provide depth estimation for severity analysis. Additionally, expanding the system's ability to operate seamlessly on Internet of Things (IoT) platforms and connected vehicles will allow for continuous, real-time monitoring across vast road networks. Another promising direction involves predictive analytics, where historical data is used to forecast potential road degradation, enabling preemptive maintenance planning. Finally, adapting the solution to accommodate global road diversity and aligning it with smart city initiatives will ensure wider applicability and long-term sustainability of the system.

References

[1] J. K. Chandan, "Challenges in Forensic Toxicology," 2024.
[2] M. E. Segura et al., "Transdermal Alcohol Detection Using Hyperdimensional Computing on Embedded Devices," 2024.
[3] M. Monaghan et al., "Intoxication: Sociological and Regulatory Perspectives," 2023.
[4] J. L. Thompson, "Overview of Intoxication Detection Techniques," 2022.
[5] B. M. Somashekar et al., "Literature Review on Pothole Detection Techniques".
[6] "Pothole Detection for Autonomous Vehicles Using Deep Learning".
[7] "Pothole Detection and Dimension Estimation Using In-Vehicle Sensors."
[8] "Evaluation of Pothole Detection Performance Using Deep Learning."
[9] "Pothole Detection Using Deep Learning: A Real-Time and AI-on-Edge Approach."