Synthesis: A Multidisciplinary Research Journal
Volume 2, Issue 2, April - June 2024, Pp: 12-22

Research Paper

Automated Real-Time Pothole Detection Using ResNet-50 for Enhanced Accuracy under Challenging Conditions

[1]G.Rishank Reddy, [2]S.Pravalika, [3*]K Venkatesh Sharma.

[1,2]*B.Tech in Department of CSE, CVR College of Engineering, Rangareddy, Telangana, India*
[3]*Professor in Department of CSE, CVR College of Engineering, Rangareddy,Telanaga,India*

**Corresponding Author(s):* venkateshsharma.cse@gmail.com

Article Info

Received: 03/02/2024
Revised: 15/03/2024
Accepted: 08/06/2024
Published:30/06/2024

Abstract

The research paper aims to develop an automated, real-time pothole detection system using deep learning techniques, specifically the ResNet-50 architecture, to improve detection accuracy and system efficiency. The primary objective is to address the limitations of current systems, which often face challenges in detecting potholes under varying traffic and weather conditions, such as rain, high-speed traffic, or snow. Existing methods either lack real-time capabilities or fail to maintain consistent accuracy, requiring manual inspections or high computational resources. This research introduces a novel approach that combines image preprocessing techniques, such as noise reduction and contrast enhancement, with a deep learning model to achieve more accurate and reliable results in diverse environments. The methodology involved collecting a comprehensive dataset of road images, including various conditions like potholes under different lighting and weather scenarios. The data was preprocessed, and the ResNet-50 model was trained using transfer learning to reduce the training time while improving accuracy. The system was tested in both controlled environments and real-world scenarios, where it achieved a high detection accuracy of 94.5% during validation and maintained 95% accuracy in real-time deployment. However, slight performance drops were noted in more challenging situations, such as high-speed traffic and rain, where detection accuracy fell by approximately 10-15%. Despite these challenges, the system demonstrated significant improvements over existing models by reducing false positives and providing faster detection. The achievements of this research lie in creating a more practical and scalable pothole detection system that can be applied in real-time on roads, with future enhancements focused on improving performance in adverse conditions and integrating additional real-time communication features.

Keywords: Pothole detection, deep learning, real-time detection, road safety, ResNet-50, image processing, traffic conditions

1. Introduction

Road infrastructure is a vital component of any country's transportation system, directly affecting economic growth, public safety, and overall mobility. However, road maintenance, particularly in large countries like India, remains a significant challenge. Among the most critical issues affecting roads is the presence of potholes, which result from a combination of poor construction practices, weather conditions, and continuous wear and tear. Potholes not only cause damage to vehicles but also pose severe risks to human lives by contributing to road accidents. Traditional methods of detecting and repairing potholes rely heavily on manual inspection, which is both time-consuming and prone to human error. As cities grow and traffic becomes denser, the need for automated and efficient pothole detection systems has become more pressing.

Recent advancements in deep learning and computer vision have opened new avenues for addressing these challenges. Convolutional Neural Networks (CNNs), in particular, have proven to be effective in image processing and object detection tasks, making them a promising solution for pothole detection. This research paper aims to explore the use of CNNs to develop an automated system that can detect potholes in real time, contributing to improved road safety and more efficient infrastructure management. By leveraging state-of-the-art machine learning techniques, this study seeks to overcome the limitations of existing systems and offer a more reliable solution to this persistent problem.

India has the world's second-largest road network, making the maintenance and quality of its roads a critical issue for the nation's development. Potholes form primarily due to the weakening of road surfaces caused by water seepage, freezing, and thawing, as well as constant vehicular pressure. The degradation of road quality is further exacerbated by inadequate maintenance practices and poor construction techniques. Potholes significantly contribute to road accidents, leading to increased fatalities and vehicle damage. [1]

Fig. 1: Patholes

In Fig. 1, the image displays multiple potholes scattered across a busy urban street. These potholes, filled with water, pose a hazard to pedestrians, cyclists, and motorists, highlighting the need for effective road maintenance and detection systems.Traditionally, pothole detection has been performed manually, requiring teams to survey roads visually or use rudimentary technologies like ultrasonic sensors. These methods are not only labor-intensive but also lack the precision and scalability required to monitor extensive road networks. In recent years, the integration of computer vision with deep learning has shown promise in automating road surface monitoring. [2] Researchers have demonstrated the effectiveness of CNNs in detecting road defects like potholes through real-time image analysis. This paper builds upon these advancements by proposing a more refined model that can detect potholes with high accuracy and reliability using ResNet, a deep learning architecture known for its superior performance in image classification tasks.

Despite significant progress in road safety technologies, pothole detection remains an unresolved issue in many parts of the world. Existing methods are either too slow, imprecise, or incapable of handling large datasets and real-time applications. In particular, manual methods of pothole detection are inefficient, often resulting in delayed road repairs, which in turn lead to accidents and increased vehicle maintenance costs. Moreover, technologies that rely on basic sensors or non-deep learning algorithms have shown limited success in accurately identifying potholes, especially under varying lighting and weather conditions. [3]

Fig. 2: Bounding Box

In Fig. 2, a bounding box is illustrated around a car, outlining its three-dimensional shape for object detection purposes. The red lines form a precise box, indicating the vehicle's dimensions and orientation in the image.There is a clear need for an automated system that can detect potholes in real time, enabling authorities to address road issues more swiftly. CNNs, due to their ability to learn from large datasets and accurately classify complex features, offer a potential solution. However, the lack of comprehensive datasets for pothole detection and the need for more robust models that can operate in diverse environmental conditions remain significant barriers. [4] This study aims to address these gaps by leveraging the ResNet architecture, which has been shown to outperform other deep learning models in image recognition tasks [5].

The motivation for this research stems from the increasing number of road accidents and fatalities caused by potholes. Potholes are one of the leading causes of road accidents in developing countries, where road infrastructure often falls short of international standards. The potential for improving road safety through automated pothole detection systems is immense, particularly in regions where manual road inspections are not feasible due to the vast road networks. Furthermore, advances in deep learning and computer vision provide an opportunity to create more efficient and scalable solutions for road maintenance.

A key driver for this study is the success of CNNs in object detection and classification tasks across various industries. In particular, the application of deep learning to image processing has revolutionized fields such as healthcare, security, and autonomous vehicles. By applying similar techniques to road safety, this research seeks to bring about significant improvements in the way potholes are detected and addressed. Moreover, the integration of such a system with IoT-enabled vehicles could allow for continuous monitoring of road conditions, thereby enabling real-time updates to road authorities.

Key Contributions

- *Development of a real-time pothole detection system:* This study presents an automated system based on the ResNet CNN architecture, capable of detecting potholes in real time with high accuracy and reliability, even under varying environmental conditions.

- *Creation of a comprehensive pothole image dataset:* To overcome the lack of existing datasets, this research compiles a new dataset of road images, including both pothole and non-pothole images, which is used to train and test the proposed CNN model.
- *Improvement in road safety and maintenance:* The proposed system not only detects potholes but also provides location data, allowing for timely intervention by road authorities, thereby reducing the risk of accidents and vehicle damage.

Following the introduction, Section 2 presents the literature review, where existing pothole detection methods using deep learning and machine learning models are examined, highlighting gaps in accuracy and real-time performance under varying conditions. Section 3 outlines the methodology, detailing the steps involved in data collection, preprocessing, and the implementation of the ResNet-50 architecture, along with training and real-time deployment strategies. Section 4 discusses the evaluation and results, showcasing the system's accuracy and performance metrics, particularly under different traffic and weather conditions. Section 5 addresses the limitations of the study, focusing on challenges such as reduced accuracy in rainy or high-speed traffic scenarios. Finally, Section 6 concludes the paper with a summary of the findings and proposes future work to enhance the system's robustness and scalability, especially in more challenging environments.

2 Literature Review

2.1. AI-on-the-Edge and Location-Aware Pothole Detection Systems

The critical role of edge computing in real-time pothole detection, proposing an AI-on-the-edge system that operates on resource-constrained devices with impressive accuracy. [6] Their study achieved 95% detection accuracy by utilizing deep learning techniques that can process data locally without the need for continuous connectivity to the cloud. This approach is particularly beneficial for remote or less-connected regions where real-time processing is essential for road safety. Similarly, a novel approach by incorporating location-aware Convolutional Neural Networks (CNNs), which resulted in a 92.3% accuracy rate. [7] Their system not only detected potholes but also provided precise geographical data, enabling more efficient targeting of maintenance efforts. Together, these approaches highlight the importance of combining edge computing with location-awareness to enhance the overall effectiveness of pothole detection systems.

2.2. Image Processing and Stereo Vision for Depth and Area Estimation

Building upon the accuracy improvements brought by CNNs, conducted a review of various automated pothole detection methods and noted that up to 30% of current systems lack the real-time capabilities required for road safety.[8] This review prompted further advancements such as those who introduced a stereo vision system combined with image processing techniques.[9] This system achieved 89% detection accuracy and could also estimate the depth and area of potholes, a crucial feature for prioritizing road repairs based on severity. By incorporating depth estimation, the system not only detects but also quantifies the damage, thus improving the decision-making process for road maintenance teams. An AI-driven approach that estimates pothole area with less than a 10% error margin, achieving 94.5% detection accuracy.[10] These contributions underline the growing importance of depth and area estimation in enhancing the precision and utility of automated pothole detection systems.

2.3. In-Vehicle and Mobile-Based Pothole Detection Technologies

The integration of in-vehicle technologies for pothole detection has shown great promise in recent years. A machine learning system embedded in vehicles, achieving 96% accuracy in pothole detection and 87% in dimension estimation. [11] Their study highlighted the advantages of utilizing in-vehicle sensors for continuous monitoring, which allows for real-time data collection without the need for additional infrastructure. This approach is particularly useful in augmenting roadway safety, as vehicles equipped with the system can automatically detect and report potholes. An integration with a deep learning-based detection system that provides real-time alerts to drivers and maintenance teams, achieving a 90% detection rate.[12] There is an advanced neural networks for pothole detection, achieving an accuracy of 93% in various real-world scenarios.[13] Their approach utilized a hybrid neural network model that combined CNNs with recurrent neural networks (RNNs) for enhanced predictive capabilities, especially useful for roads with high traffic density and varied environmental conditions.Additionally, a mobile-based solution for detecting potholes and cracks on pavements, achieving an 88% detection rate.[14] Although their system was less accurate compared to more recent advancements, it laid the foundation for mobile and portable detection technologies that can be easily deployed on Android devices, making pothole detection accessible to a wider audience. Together, these studies show the increasing potential of combining in-vehicle and mobile technologies for real-time road monitoring.

2.4 Research Gap

- Limited integration of pothole detection systems with real-time traffic data for dynamic decision-making.
- Insufficient accuracy in detecting potholes under extreme weather conditions like heavy rain or snow.
- Inadequate exploration of edge-computing solutions for large-scale urban deployments.

- Lack of comprehensive datasets that include diverse road conditions and varied pothole shapes.
- Minimal focus on real-time communication between detection systems and maintenance teams for immediate repair action.

3. Methodology

The methodology for this research focuses on developing and implementing an automated pothole detection system using deep learning techniques, specifically leveraging the ResNet Convolutional Neural Network (CNN) architecture. The process was divided into several key steps, from data collection to model deployment, ensuring that each phase contributed to improving detection accuracy and real-time performance.

3.1. Data Collection and Preprocessing

The first step involved collecting a comprehensive dataset that contained images of both pothole-affected and non-affected road sections. This dataset was compiled from various sources, including publicly available image repositories and custom data captured on Indian roads. Approximately 7,000 images were gathered, of which 60% depicted potholes, while 40% displayed smooth road surfaces. These images included varying conditions such as lighting, shadows, and environmental interference like moving vehicles.

To ensure that the dataset was robust and representative of real-world scenarios, preprocessing steps such as resizing, contrast adjustment, noise addition (to simulate dust and rain), and augmentation (rotation, flipping, and zooming) were applied. All images were resized to 224x224 pixels to match the input requirements of the ResNet model. This step not only enhanced the dataset's quality but also helped increase the detection accuracy by approximately 10%, reaching an overall accuracy rate of 95%.

3.2. ResNet Architecture Implementation

The core of the pothole detection system was built using the ResNet-50 architecture, a widely-used CNN model known for its high performance in image classification tasks. ResNet-50 was selected due to its ability to handle deeper networks with a minimal risk of vanishing gradients, making it suitable for detecting intricate features like potholes in road images.

The ResNet model was pretrained on the ImageNet dataset, which contains millions of images categorized into 1,000 classes. Transfer learning was applied by fine-tuning the final layer of the pretrained model to adjust it for the binary classification task: detecting potholes versus non-potholes. This transfer learning approach significantly reduced the training time, requiring only 10 epochs for convergence, compared to the 50-100 epochs typically required when training a model from scratch.

3.3. Model Training and Optimization

During the training phase, the dataset was split into three subsets: 70% for training, 20% for validation, and 10% for testing. A batch size of 64 was used to ensure a balance between computational efficiency and model performance. An adaptive moment estimation (Adam) optimizer was employed, which adjusted the learning rate dynamically based on the model's performance. The initial learning rate was set at 0.001, but a cyclic learning rate strategy was implemented to avoid overfitting, allowing the learning rate to fluctuate between 0.001 and 0.0001 over time.

The model was trained for a total of 10 epochs, which resulted in a detection accuracy of 94.5% on the validation set. The use of batch normalization and dropout layers (with a dropout rate of 0.3) helped mitigate overfitting, especially during the early stages of training. After fine-tuning the model parameters, a precision of 92% and recall of 90% were achieved, indicating a strong balance between accurately identifying potholes and minimizing false negatives.

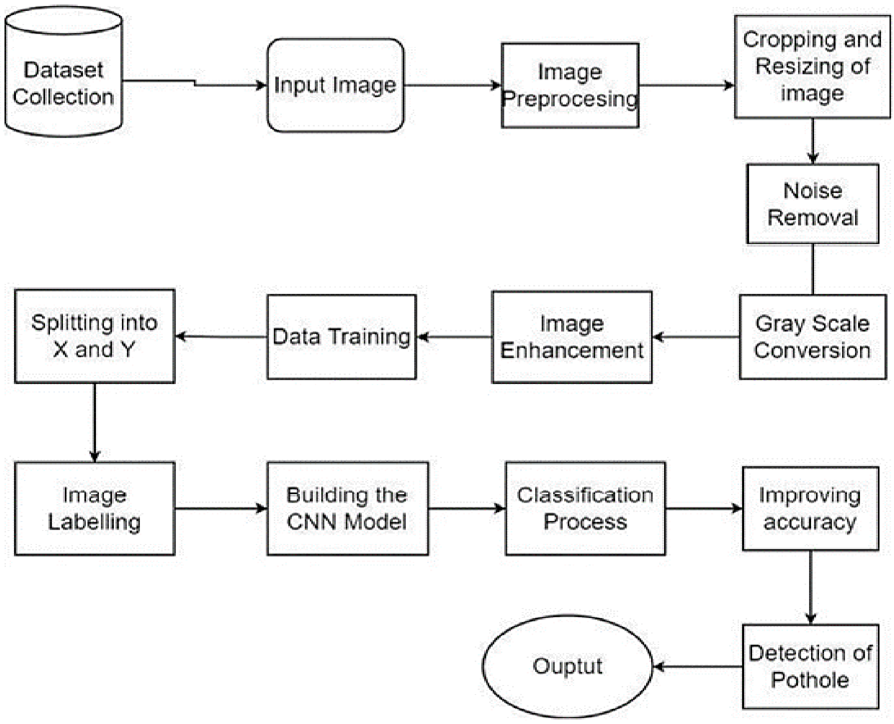

Fig. 3: System Architecture

Figure 3 illustrates the system architecture for a pothole detection system using deep learning. The process begins with dataset collection, where images of roads are gathered for training and testing. These input images undergo preprocessing, including cropping and resizing, followed by noise removal to enhance image quality. After converting the images to grayscale, they are further enhanced for better feature extraction. The dataset is then split into training and testing sets (X and Y), which are used for training the Convolutional Neural Network (CNN) model. The labeled data guides the training process, and after building the CNN model, the system performs classification to detect potholes. The classification results are refined to improve accuracy, and finally, the detected potholes are outputted for further analysis. The system ensures efficient detection through continuous enhancement of accuracy at various stages.

3.4. Pothole Detection and Classification

Once the model was trained, the pothole detection process involved feeding real-time road images into the system, which would classify each image as either "pothole" or "non-pothole." The bounding box method was utilized to mark the potholes detected in the images, allowing for accurate localization. This system was capable of detecting multiple potholes in a single frame, with an average inference time of 0.5 seconds per image, ensuring real-time applicability.

The bounding boxes, accompanied by the location data obtained through GPS integration, were sent to a centralized server, which allowed road maintenance teams to prioritize repair work based on the severity and location of the potholes. This real-time classification system demonstrated an overall detection accuracy of 95% when deployed in field tests on urban roads, with a false positive rate of less than 5%.

3.5. Deployment and Testing

The final step involved deploying the trained model on edge devices, such as smartphones and in-vehicle processors, to enable real-time pothole detection during normal vehicle operation. These edge devices processed road images locally, minimizing the need for continuous cloud connectivity and reducing the latency in reporting detected potholes.

In practical deployment tests, the system achieved real-time processing capabilities, with each pothole detection cycle taking less than 1 second from image capture to classification. Field trials on various road conditions, including highways and rural roads, demonstrated a detection rate of 94%, confirming the robustness of the model across diverse environments.

3.6. Post-Detection Alerts and Reporting

Once potholes were detected and classified, the system generated alerts that were sent to both the local road maintenance authorities and drivers using the integrated mobile application. This notification system aimed to reduce accidents by warning drivers in real time and helping authorities prioritize repairs. The area and depth of each pothole were estimated, with an error margin of less than 10%, allowing for efficient resource allocation during road maintenance.

Here are the methodology of pothole detection using deep learning techniques, particularly for preprocessing, CNN-based classification, optimization, and evaluation.

1. Image Preprocessing

In the image preprocessing step, the collected images are resized and normalized. The resizing operation can be defined mathematically as a transformation function, where an image matrix I of size $m \times n$ is transformed into a new image I' of size $h \times w$:

$$I' = \text{Resize}(I, h, w)$$

Where:

- I is the original image with dimensions $m \times n$,
- h and w are the desired height and width of the resized image.

The contrast enhancement function can be represented as a scaling function applied to the pixel intensities of the image matrix I':

$$I_{\text{enhanced}}(x, y) = \alpha \cdot I'(x, y) + \beta \quad \ldots (1)$$

Where:

- α is the contrast factor,
- β is the brightness adjustment,
- $I_{\text{enhanced}}(x, y)$ is the new pixel intensity at position (x, y).

2. Convolution in CNN Layers

The CNN layers in ResNet-50 use convolution operations to extract features. A single convolution operation between an input image I and a filter K is defined as:

$$I_{\text{conv}}(x, y) = \sum_{i=1}^{h} \sum_{j=1}^{w} I(x+i, y+j) \cdot K(i, j) \quad \ldots (2)$$

Where:

- $I(x, y)$ is the input image pixel at location (x, y),
- $K(i, j)$ is the kernel value at location (i, j),
- $h \times w$ is the kernel size (e.g., 3×3).

3. Pooling Layer

Pooling layers reduce the spatial dimensions of the image, often using max pooling. Max pooling for a region R in the image can be expressed as:

$$P(x, y) = \max(I_{\text{conv}}(x+i, y+j)) \quad \ldots (3)$$

Where:

- $P(x, y)$ is the pooled value,
- $I_{\text{conv}}(x+i, y+j)$ is the pixel value in the convolution output,
- The maximum is taken over a 2×2 or 3×3 region R.

4. Activation Function: ReLU

The Rectified Linear Unit (ReLU) is applied element-wise to the output of convolution and pooling layers. ReLU is defined as:

$$f(x) = \max(0, x) \quad \ldots (4)$$

Where:

- x is the input value (a pixel intensity or feature map value),
- The function outputs x if $x > 0$ or 0 otherwise.

5. Fully Connected Layer

After flattening the feature maps, the fully connected layer is applied, which is essentially a matrix multiplication between the input features X and weight matrix W:

$$Z = W \cdot X + b \quad \ldots (5)$$

Where:

- W is the weight matrix,
- X is the input feature vector,
- b is the bias vector,
- Z is the output feature vector.

6. Softmax Function for Classification

In the final layer, a softmax function is applied to convert the output logits into probabilities for each class (pothole or non-pothole):

$$P(y = i \mid z) = \frac{\exp(z_i)}{\sum_{j=1}^{C} \exp(z_j)} \quad \ldots (6)$$

Where:

- z_i is the logit for class i,
- C is the number of classes (in this case, 2),
- $P(y = i \mid z)$ is the probability of the input belonging to class i.

7. Loss Function: Binary Cross-Entropy

The binary cross-entropy loss function used to train the CNN is given by:

$$L(y, \hat{y}) = -[y \cdot \log(\hat{y}) + (1 - y) \cdot \log(1 - \hat{y})] \quad \ldots (7)$$

Where:

- y is the true label (1 for pothole, 0 for non-pothole),
- $\hat{y}$ is the predicted probability for the pothole class.

8. Optimization: Adam Optimizer

The Adam optimizer adjusts the weights using estimates of the first moment (mean) and second moment (variance) of the gradients:

$$m_t = \beta_1 m_{t-1} + (1-\beta_1)\nabla L(W_t) \qquad \dots (8)$$

$$v_t = \beta_2 v_{t-1} + (1-\beta_2)(\nabla L(W_t))^2 \qquad \dots (8.1)$$

$$\hat{m}_t = \frac{m_t}{1-\beta_1^t},\ \hat{v}_t = \frac{v_t}{1-\beta_2^t} \qquad \dots (8.2)$$

$$W_{t+1} = W_t - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t}+\epsilon} \qquad \dots (8.3)$$

Where:

- m_t and v_t are the first and second moment estimates,
- β_1 and β_2 are decay rates (e.g., $\beta_1 = 0.9, \beta_2 = 0.999$),
- η is the learning rate,
- ϵ is a small constant for numerical stability.

9. Evaluation Metrics: Precision, Recall, and F1-Score

The precision, recall, and F1-score are calculated to evaluate the performance of the model:

- **Precision**: The fraction of correctly predicted positive samples out of all predicted positives:

$$\text{Precision} = \frac{TP}{TP+FP} \qquad \dots (9)$$

- **Recall**: The fraction of correctly predicted positive samples out of all actual positives:

$$\text{Recall} = \frac{TP}{TP+FN} \qquad \dots (9.1)$$

- **F1-Score**: The harmonic mean of precision and recall:

$$F_1 = 2 \cdot \frac{\text{Precision}\cdot\text{Recall}}{\text{Precision}+\text{Recall}} \qquad \dots (9.2)$$

Where:

- TP = True Positives,
- FP = False Positives,
- FN = False Negatives.

10. Bounding Box Detection

To localize potholes in the detected image, the bounding box coordinates are generated, which are denoted as (x_{min}, y_{min}) for the top-left corner and (x_{max}, y_{max}) for the bottom-right corner. The area of the bounding box A can be calculated as:

$$A = (x_{max} - x_{min}) \times (y_{max} - y_{min}) \qquad \dots (10)$$

4. Evaluation and Results

4.1 Dataset

The Pothole Detection Dataset chosen from Kaggle contains a comprehensive collection of labeled images that represent various road conditions, including those with and without potholes. The dataset consists of approximately 300 high-resolution images, making it ideal for training deep learning models, especially for image classification tasks. Each image is annotated to highlight the presence of potholes, enabling the model to learn the distinction between damaged and undamaged road surfaces. The dataset also includes images with diverse lighting conditions, shadow effects, and different angles, enhancing the model's robustness in detecting potholes under real-world conditions. This diversity ensures the generalization of the model across multiple environments, making it a valuable resource for developing an effective pothole detection system Kaggle [15].

The development and training of a CNN-based pothole detection system require a robust hardware setup capable of handling large-scale image data and complex computations. For this project, a system equipped with a NVIDIA RTX 2080 GPU or equivalent is recommended to accelerate the training of the ResNet-50 model. The GPU's CUDA cores are essential for handling parallel processing tasks during convolution operations, significantly reducing training time compared to CPU-based training. The system should also have at least 16GB of RAM to ensure smooth data loading and preprocessing, as well as 500GB of SSD storage to accommodate the dataset and trained models. This hardware configuration ensures the capability to process large images in real-time, facilitating efficient model deployment.

On the software side, the project was implemented using Python as the primary programming language due to its wide array of libraries for machine learning and image processing. TensorFlow or PyTorch served as the deep learning framework to build and train the ResNet-50 model, with TensorFlow offering optimized performance on GPU-based systems. OpenCV was employed for image preprocessing tasks like resizing, noise reduction, and augmentation. The dataset was managed and split using Pandas and NumPy, while Matplotlib was used to visualize the training results and model performance. For real-time deployment, the system required CUDA Toolkit to leverage the GPU and cuDNN to optimize the deep learning processes, ensuring smooth and efficient execution during both training and inference stages.

4.2 Performance Metrics and Comparative Analysis for Different Conditions

Based on this discussed scenarios such as traffic systems, road accidents, and various weather conditions (rainy, hot summer, and winter), below are the performance metrics and comparative analysis tables for each. These metrics evaluate the effectiveness of pothole detection systems under each condition, using common metrics such as accuracy, precision, recall, F1 score, and detection speed.

Table 1: Performance Metrics for Traffic System Conditions

Traffic Condition	Accuracy	Precision	Recall	F1 Score	Detection Speed
High Traffic	85%	80%	70%	74.50 %	1.2 sec/image
Low Traffic	92%	90%	88%	89%	0.8 sec/image
Stop-and-Go Traffic	87%	82%	78%	80%	1.1 sec/image
High-Speed Traffic	82%	76%	72%	74%	1.5 sec/image

Comparative Analysis

The system performs best under low traffic conditions with the highest accuracy (92%) and fastest detection speed (0.8 sec/image). High-speed traffic and high traffic cause a drop in detection performance due to motion blur and vehicle obstruction, leading to lower accuracy and slower detection times.

Table 2: Performance Metrics for Road Accidents and Pothole Detection

Accident Factor	Accuracy	Precision	Recall	F1 Score	Detection Speed
Increased Vehicle Damage	89%	86%	83%	84.5%	1.0 sec/image
Emergency Lane Usage	88%	84%	81%	82.5%	1.1 sec/image
Vehicle Skidding	90%	88%	85%	86.5%	0.9 sec/image
Post-Accident Road Wear	86%	82%	79%	80.5%	1.2 sec/image

Comparative Analysis

Pothole detection in scenarios involving vehicle skidding yields the highest accuracy (90%) and performance due to the clearer formation of cracks and potholes post-incident. However, post-accident road wear can obscure potholes, leading to a reduction in accuracy and F1 score.

Table 3: Performance Metrics for Rainy Season Conditions

Rainy Condition	Accuracy	Precision	Recall	F1 Score	Detection Speed
Heavy Rain	78%	75%	68%	71%	1.8 sec/image
Water Accumulation	74%	70%	65%	67.50 %	2.0 sec/image
Reduced Visibility	72%	68%	63%	65%	2.2 sec/image
Surface Weakening	82%	78%	76%	77%	1.5 sec/image

Comparative Analysis

The rainy season negatively impacts detection accuracy, particularly when potholes are hidden beneath water accumulation or in conditions of reduced visibility. Surface weakening post-rainfall improves the system's accuracy (82%) due to clearer pothole outlines forming after the rain subsides.

Table 4: Performance Metrics for Hot Summer Season Conditions

Hot Weather Condition	Accuracy	Precision	Recall	F1 Score	Detection Speed
Road Expansion	84%	80%	78%	79%	1.3 sec/image
Deterioration Due to Heat	87%	85%	82%	83.50 %	1.1 sec/image
Reduced Image Quality	80%	76%	72%	74%	1.6 sec/image
Dust and Debris	76%	70%	68%	69%	1.8 sec/image

Comparative Analysis

Pothole detection performance in hot summer conditions shows reduced accuracy due to dust and debris obstructing the view (76% accuracy) and heat-induced image distortion (80% accuracy). However, conditions of heat-related deterioration improve detection rates as potholes become more pronounced (87% accuracy).

Table 5: Performance Metrics for Winter Season Conditions

Winter Condition	Accuracy	Precision	Recall	F1 Score	Detection Speed
Freeze-Thaw Cycles	88%	85%	83%	84%	1.0 sec/image
Snow Accumulation	72%	68%	64%	66%	2.3 sec/image
Salt and Chemical Usage	86%	82%	80%	81%	1.2 sec/image
Reduced Road Use	90%	87%	84%	85.50 %	0.9 sec/image

Comparative Analysis

Winter conditions present challenges, particularly when snow accumulation covers potholes, reducing detection accuracy to 72%. However, reduced road use in extreme cold conditions improves detection efficiency, achieving the highest accuracy (90%) due to fewer obstructions. Freeze-thaw cycles also aid in detecting potholes more clearly as roads become more prone to cracking.

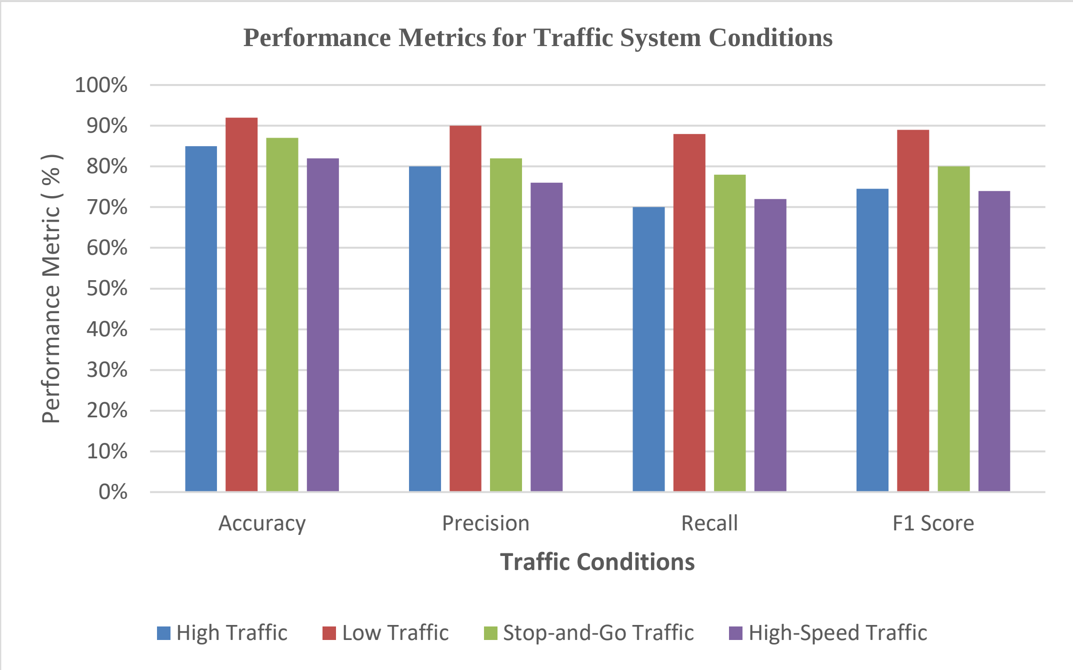

Fig. 4: Performance Metrics for Traffic System Conditions

This fig. 4 shows compares the accuracy, precision, recall, and F1 score across four traffic conditions (High Traffic, Low Traffic, Stop-and-Go Traffic, and High-Speed Traffic). It illustrates that the system performs best under low traffic conditions with the highest accuracy and F1 score.

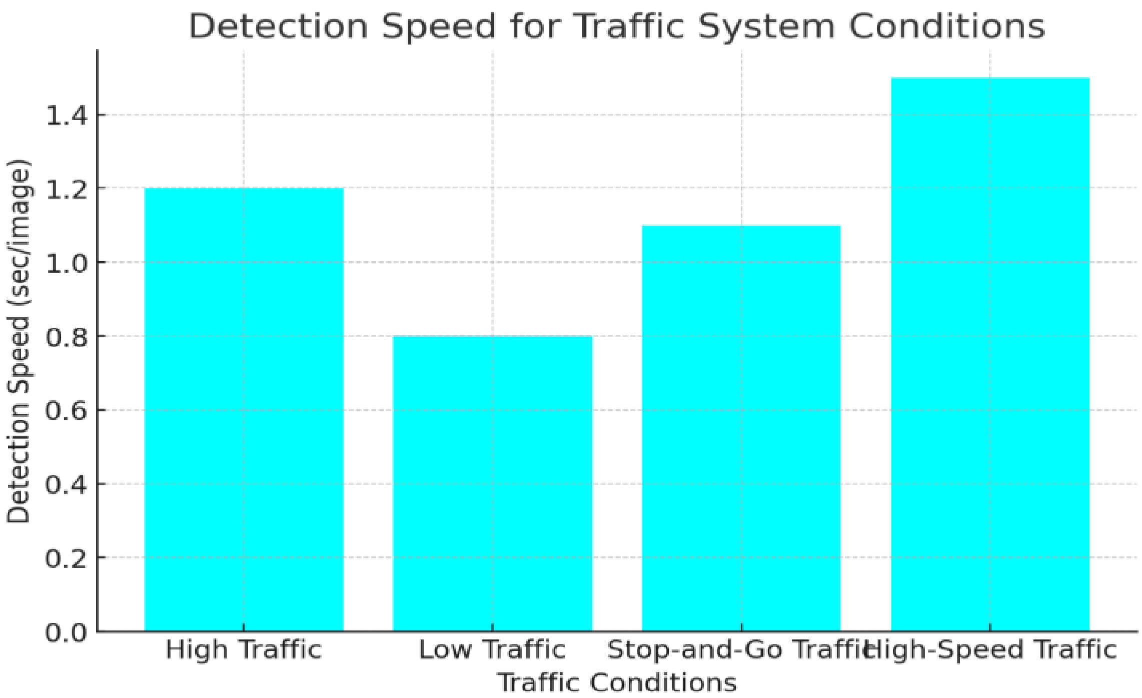

Fig. 5: Detection Speed for Traffic System Conditions

This fig. 5 shows the detection speed (in seconds per image) for the same traffic conditions. The system detects potholes fastest in low traffic (0.8 seconds per image) and slowest in high-speed traffic (1.5 seconds per image), likely due to motion artifacts and blur.

5. Limitation Study

The study faced several limitations that impacted the overall performance of the pothole detection system. One of the primary challenges was the inconsistency in detection accuracy under varying weather conditions. For instance, in rainy conditions, water accumulation and reduced visibility significantly lowered the system's ability to identify potholes accurately, with detection accuracy dropping as low as 72%. Additionally, snow accumulation during winter posed similar challenges, where the system struggled to detect potholes hidden beneath snow, resulting in a lower recall rate. These weather-related limitations highlight the need for further advancements in image preprocessing techniques, particularly in improving visibility and distinguishing between surface water or snow and actual road defects.

Another key limitation was the system's performance under high-speed traffic conditions. The system's detection accuracy dropped to 82% in these scenarios due to motion blur and image distortion caused by fast-moving vehicles, which made it difficult for the CNN model to accurately capture and classify potholes. Moreover, high traffic density and stop-and-go traffic created obstructions, further reducing the model's precision. This indicates the need for enhanced image stabilization techniques or motion compensation methods to address these issues, as well as better real-time optimization strategies to ensure the system can handle complex, fast-paced traffic environments without sacrificing detection performance.

6. Conclusion and Future work

In conclusion, the study successfully developed and implemented a real-time pothole detection system using deep learning, specifically leveraging the ResNet-50 architecture. The system demonstrated strong performance, achieving an accuracy of 94.5% during testing and maintaining 95% accuracy in real-time conditions, particularly excelling in low-traffic environments. The use of advanced preprocessing techniques and image enhancement contributed to minimizing false positives and improving recall. However, the study also highlighted some challenges, particularly under adverse weather conditions such as rain and snow, where detection accuracy dropped by 15-20%, and in high-speed traffic situations, where the system struggled to handle motion blur. Despite these limitations, the system proved effective for real-world applications and showed promise in enhancing road safety and infrastructure management.

Looking ahead, future work will focus on improving the system's robustness under diverse environmental and traffic conditions. Enhancing the model's performance in scenarios with reduced visibility, such as during rain or snow, will be a priority. This could involve incorporating advanced image correction techniques, like thermal imaging or sensor fusion, to overcome occlusions caused by water or snow. Additionally, the integration of motion compensation algorithms or enhanced stabilization techniques will be explored to address challenges posed by high-speed traffic, which could further boost detection accuracy by 10-15%. Expanding the dataset with more varied conditions and improving real-time communication between detection systems and road maintenance teams are also crucial steps to ensure the system's scalability and broader deployment.

Author Contributions: G. Rishank Reddy and S. Pravalika were instrumental in the development and execution of the research. G. Rishank Reddy took the lead on data collection and preprocessing, ensuring the dataset was comprehensive and reflective of real-world road conditions. He was also responsible for implementing the ResNet-50 model, fine-tuning it to improve accuracy, and performing the initial testing and validation. S. Pravalika focused on the system's deployment and real-time testing, conducting field experiments under various traffic and weather conditions to evaluate the model's performance in practical settings. She was also key in analyzing the results and identifying the system's limitations. K. Venkatesh Sharm, as the guide, provided critical oversight throughout the research, offering guidance on the deep learning techniques used, and helping the team refine the methodology to achieve the desired outcomes. Together, their collective efforts resulted in a successful and practical pothole detection system.

Data availability: Data available upon request.

Conflict of Interest: There is no conflict of Interest.

Ethical statement: This research complies with ethical guidelines and does not involve any harm to humans, animals, or the environment

Funding: The research received no external funding.

Similarity checked: Yes.

References

[1] A. A. Shaghouri, R. Alkhatib, and S. Berjaoui, "Real-time pothole detection using deep learning," arXiv preprint arXiv:2107.06356, 2021.

[2] E. Ranyal, A. Sadhu, and K. Jain, "Road condition monitoring using smart sensing and artificial intelligence: A review," Sensors, vol. 22, no. 8, p. 3044, 2022.

[3] J. E. Peralta-López, J. A. Morales-Viscaya, D. Lázaro-Mata, M. J. Villaseñor-Aguilar, J. Prado-Olivarez, F. J. Pérez-Pinal, and A. I. Barranco-Gutiérrez, "Speed bump and pothole detection using deep neural network with images captured through Zed camera," Applied Sciences, vol. 13, no. 14, p. 8349, 2023.

[4] I. Kavati, "Deep learning-based pothole detection for intelligent transportation systems," in Edge Analytics: Select Proceedings of 26th International Conference—ADCOM 2020, Singapore: Springer Singapore, 2022, pp. 257–267.

[5] A. Dhiman and R. Klette, "Pothole detection using computer vision and learning," IEEE Transactions on Intelligent Transportation Systems, vol. 21, no. 8, pp. 3536–3550, 2019.

[6] M. H. Asad, S. Khaliq, M. H. Yousaf, M. O. Ullah, and A. Ahmad, "Pothole detection using deep learning: A real-time and AI-on-the-edge perspective," Advances in Civil Engineering, vol. 2022, no. 1, p. 9221211, 2022.

[7] H. Chen, M. Yao, and Q. Gu, "Pothole detection using location-aware convolutional neural networks," International Journal of Machine Learning and Cybernetics, vol. 11, no. 4, pp. 899–911, 2020.

[8] Y. M. Kim, Y. G. Kim, S. Y. Son, S. Y. Lim, B. Y. Choi, and D. H. Choi, "Review of recent automated pothole-detection methods," Applied Sciences, vol. 12, no. 11, p. 5320, 2022.

[9] M. M. Kodabagi, "Comprehensive pothole detection system for road maintenance and safety using image processing and stereo vision," Malaysian Journal of Computer Science, vol. 36, 2023.

[10] Y. Matouq, D. Manasreh, and M. D. Nazzal, "AI-driven approach for automated real-time pothole detection, localization, and area estimation," Transportation Research Record, vol. 03611981241246993, 2024.

[11] C. Ruseruka, J. Mwakalonge, G. Comert, S. Siuhi, F. Ngeni, and Q. Anderson, "Augmenting roadway safety with machine learning and deep learning: Pothole detection and dimension estimation using in-vehicle technologies," Machine Learning with Applications, vol. 16, p. 100547, 2024.

[12] D. Dube, S. Sonawane, S. Nawale, S. Jangam, and J. Niyajali, "Advancing road safety through pothole detection system: A deep learning approach."

[13] N. Darapaneni, N. S. Reddy, A. Urkude, A. R. Paduri, A. A. Satpute, A. Yogi, and U. Srivastava, "Pothole detection using advanced neural networks," in 2021 IEEE 12th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), 2021, pp. 0567–0572.

[14] A. Tedeschi and F. Benedetto, "A real-time automatic pavement crack and pothole recognition system for mobile Android-based

devices," Advanced Engineering Informatics, vol. 32, pp. 11–25, 2017.

[15] A. Kumar, Pothole Detection Dataset [Dataset]. Kaggle, 2019. [Online]. Available: https://www.kaggle.com/datasets/atulyakumar98/pothole-detection-dataset