A Major Project Report On

# AUTOMATED ROAD DAMAGE DETECTION USING UAV IMAGES AND DEEP LEARNING TECHNIQUES

Submitted to CMREC (UGC Autonomous)

In Partial Fulfilment of the requirements for the Award of Degree

of

**BACHELOR OF TECHNOLOGY**

**IN**

**COMPUTER SCIENCE AND ENGINEERING (AI & ML)**

Submitted

By


| | |
|---|---|
| **GOLLA PRANEETH** | **(228R1A6681)** |
| **NAGILLA KAVYA** | **(228R1A66A7)** |
| **PINNAMSHETTY SHASHIDAR** | **(228R1A66B5)** |
| **THYRALA SHIVANI** | **(228R1A66C4)** |


Under the Esteemed guidance of

**Mrs. T. S. Suhasini**

Assistant Professor, Department of CSE(AI&ML)



**Department of Computer Science and Engineering (AI&ML)**

## CMR ENGINEERING COLLEGE

**(UGC AUTONOMOUS)**

(Accredited by NAAC & NBA, Approved by AICTE, New Delhi, Affiliated to JNTU, Hyderabad)
(Kandlakoya, Medchal Road, Medchal Malkajgiri Dist. Hyderabad-501 401)


**(2025-2026)**

# CMR ENGINEERING COLLEGE

## UGC AUTONOMOUS

*(Accredited by NAAC&NBA, Approved by AICTE New Delhi, Affiliated to JNTU, Hyderabad, Kandlakoya, Medchal Road, Hyderabad-501 401)*

## Department of Computer Science & Engineering (AI & ML)

## CERTIFICATE

This is to certify that the Major project entitled **"AUTOMATED ROAD DAMAGE DETECTION USING UAV IMAGES AND DEEP LEARNING TECHNIQUES"** is a bonafide work carried out by

| | |
|---|---|
| **GOLLA PRANEETH** | **(228R1A6681)** |
| **NAGILLA KAVYA** | **(228R1A66A7)** |
| **PINNAMSHETTY SHASHIDAR** | **(228R1A66B5)** |
| **THYRALA SHIVANI** | **(228R1A66C4)** |

in partial fulfillment of the requirement for the award of the degree of BACHELOR OF TECHNOLOGY in COMPUTER SCIENCE AND ENGINEERING (AI&ML) from CMR Engineering College, under our guidance and supervision.

The results presented in this Major project have been verified and are found to be satisfactory. The results embodied in this Major Project have not been submitted to any other university for the award of any other degree or diploma.

| **Internal Guide** | **Major Project Coordinator** | **Head of the Department** |
|---|---|---|
| Mrs. T. S. Suhasini | Mr. G. Venkateswarlu | Dr. Madhavi Pingili |
| Assistant Professor | Assistant Professor | Professor & HOD |
| Department of CSE (AI & ML) | Department of CSE (AI & ML) | Department of CSE (AI & ML) |

**External Examiner:**

# DECLARATION

This is to certify that the work reported in the present Major project entitled "**AUTOMATED ROAD DAMAGE DETECTION USING UAV IMAGES AND DEEP LEARNING TECHNIQUES"** is a record of bonafide work done by us in the Department of Computer Science and Engineering (AI & ML), CMR Engineering College. The reports are based on the Major project work done entirely by us and not copied from any other source. We submit our Major project for further development by any interested students who share similar interests to improve the Major project in the future. The results embodied in this Major project report have not been submitted to any other University or Institute for the award of any degree or diploma to the best of our knowledge and belief.


**GOLLA PRANEETH (228R1A6681)**
**NAGILLA KAVYA (228R1A66A7)**
**PINNAMSHETTY SHASHIDAR (228R1A66B5)**
**THYRALA SHIVANI (228R1A66C4)**

# ACKNOWLEDGEMENT

We are extremely grateful to **Dr. A. Srinivasula Reddy**, Principal and **Dr. Madhavi Pingili**, Professor & HOD, Department of CSE(AI&ML), CMR Engineering College for their constant support.

We are extremely thankful to **Mrs. T. S. Suhasini,** Assistant Professor, Internal Guide, Department of CSE(AI&ML), for her constant guidance, encouragement and moral support throughout the Major project.

We will be failing in duty if we do not acknowledge with grateful thanks to the authors of the references and other literatures referred in this Major Project.

We thank **Mr. G. Venkateswarlu,** Major Project Coordinator for his constant support in carrying out the project activities and reviews.

We express our thanks to all staff members and friends for all the help and co-ordination extended in bringing out this Major project successfully in time.

Finally, We are very much thankful to our parents who guided us for every step.

**GOLLA PRANEETH (228R1A6681)**
**NAGILLA KAVYA (228R1A66A7)**
**PINNAMSHETTY SHASHIDAR (228R1A66B5)**
**THYRALA SHIVANI (228R1A66C4)**

# CONTENTS


| TOPIC | PAGE NO |
|---|---|
| **ABSTRACT** | **I** |
| **LIST OF FIGURES** | **II** |
| **1. INTRODUCTION** | **1** |
| 1.1. Introduction and Objectives | 1 |
| 1.2. Project Objectives | 1 |
| 1.3. Purpose of the project | 2 |
| 1.4. Problem Statement | 2 |
| 1.5. Existing System with Disadvantages | 3 |
| 1.6. Proposed System with Advantages | 4 |
| 1.7. Input and Output Design | 6 |
| **2. LITERATURE SURVEY** | **8** |
| **3. SOFTWARE REQUIREMENTS ANALYSIS** | **13** |
| 3.1. Modules and their Functionalities | 13 |
| 3.2. Functional Requirements | 14 |
| 3.3. Non-Functional Requirements | 14 |
| 3.4. Feasibility Study | 15 |
| **4. SYSTEM SPECIFICATIONS** | **16** |
| 4.1. Software requirements | 16 |
| 4.2. Hardware requirements | 16 |

**5. SOFTWARE DESIGN** **17**

5.1. System Architecture 17

5.2. Dataflow Diagrams 18

5.3. UML Diagrams 19

**6. CODING AND IMPLEMENTATION** **28**

6.1. Source Code 28

6.2. Implementation 79

**7. SYSTEM TESTING** **82**

7.1. Types of System Testing 82

7.2. Test Strategies 83

7.3. Sample Test Cases 86

**8. OUTPUT SCREENS** **89**

**9. CONCLUSION** **94**

**10. FUTURE ENHANCEMENTS** **95**

**11. REFERENCES** **96**

# ABSTRACT

Road infrastructure plays a critical role in ensuring safe and efficient transportation, but defects such as cracks, potholes, and surface damage can compromise safety and increase maintenance costs. Traditional inspection methods are largely manual, time-consuming, and prone to human error, making them unsuitable for large-scale monitoring. Existing approaches, including basic image processing and sensor-based techniques, suffer from limitations such as low accuracy, restricted coverage, high costs, and difficulty handling environmental variations like lighting and weather. These challenges highlight the need for an automated, reliable, and scalable solution for road damage detection.

To overcome these issues, this project proposes a deep learning-based system(ResNet101 & CNN) that uses UAV (drone) images for detecting and classifying road surface defects. The system employs a Convolutional Neural Network (ResNet-101) with transfer learning for efficient feature extraction and improved accuracy, even with limited data. An image preprocessing pipeline—comprising resizing, normalization, noise reduction, and data augmentation—enhances robustness and generalization. Additionally, a Flask-based web application allows users to upload images and obtain real-time predictions, ensuring scalability and ease of use. This approach reduces manual effort, improves detection accuracy, and supports efficient large-scale infrastructure monitoring.

**Keywords:** Road Damage Detection; UAV Imaging; Deep Learning; ResNet-101; Image Preprocessing; Transfer Learning; Computer Vision; Flask Deployment; Smart Infrastructure; Automated Inspection.

# LIST OF FIGURES


| S.NO | FIGURE NO | DESCRIPTION | PAGENO |
|---|---|---|---|
| 1. | 1.6.1 | Block diagram of proposed system | 5 |
| 2. | 5.1 | System Architecture | 17 |
| 3. | 5.2.1 | Data Flow diagram | 19 |
| 4. | 5.3.1 | Sequence diagram | 21 |
| 5. | 5.3.2 | Use case diagram | 22 |
| 6. | 5.3.3 | Activity diagram | 23 |
| 7. | 5.3.4 | Class diagram | 24 |
| 8. | 5.3.5 | Collaboration Diagram | 25 |
| 9. | 5.3.6 | Component Diagram | 26 |
| 10. | 5.3.7 | Deployment Diagram | 27 |
| 11. | 7.3.1 | User Login | 87 |
| 12. | 7.3.2 | Image Processing | 87 |
| 13. | 7.3.3 | Redirection To Web app | 88 |
| 14. | 8.2 | Confusion Matrix | 89 |
| 15. | 8.3 | Output Screens | 90 |
| 16. | 8.4 | Output Screens | 90 |
| 17. | 8.5 | Output Screen | 91 |
| 18. | 8.6 | Output Screens | 91 |
| 19. | 8.7 | Output Screens | 92 |
| 20. | 8.8 | Output Screens | 92 |
| 21. | 8.9 | Output Screens | 93 |

# LIST OF TABLES


| S.NO | TABLE NO | DESCRIPTION | PAGE NO |
|---|---|---|---|
| 1 | 2.1 | Literature Review Summary | 12 |
| 2 | 7.3 | Sample Test Cases | 86 |
| 3 | 8.1 | Classification Report | 89 |

# 1. INTRODUCTION

## 1.1 Introduction and Objectives

The rapid expansion of transportation networks and urban infrastructure has significantly increased the demand for efficient road maintenance and monitoring systems. Roads serve as a backbone for economic development, trade, and daily commuting, making their proper maintenance essential for ensuring safety and efficiency. However, road surfaces are continuously exposed to environmental factors such as temperature variations, rainfall, and heavy traffic loads, which result in various types of damage including potholes, longitudinal cracks, transverse cracks, and surface wear. These damages not only degrade transportation quality but also lead to accidents, vehicle damage, and increased operational costs.

Traditionally, road inspection is carried out manually by field personnel or through specialized survey vehicles equipped with sensors. While widely used, these methods suffer from several limitations, including high operational costs, time consumption, and lack of scalability. Manual inspections are often subjective and prone to human error, resulting in inconsistent outcomes, and they can also pose safety risks in high-traffic or hazardous environments. To address these challenges, there is a growing shift toward automated road damage detection systems. Advances in computer vision and artificial intelligence have enabled the development of intelligent models capable of analyzing visual data and accurately identifying infrastructure defects, offering faster, more reliable, and scalable solutions compared to traditional approaches.

## 1.2 Project Objectives

The main objective of this research is to develop an automated road damage classification system using CNN and ResNet101, with the following key goals:

1. Develop a high-accuracy CNN and ResNet101 to classify road damage into four categories: good, poor, satisfactory, and very poor.
2. Utilize transfer learning with ResNet101 to leverage pre-trained features and improve classification accuracy.
3. Optimize the model with techniques such as mixed precision training, gradient clipping, and learning rate scheduling to enhance performance.

4. Deploy the trained model for real-world applications, enabling users to classify road damage from uploaded images.
5. Reduce computational overhead, ensuring the system is scalable and efficient for practical use in smart infrastructure monitoring.

## 1.3 Purpose of the project

The primary purpose of this project is to develop an automated and intelligent system for detecting and classifying road damages such as cracks and potholes using deep learning techniques. The system aims to reduce the dependency on manual inspection methods by providing a faster, more accurate, and scalable solution for monitoring road conditions.

Another key objective is to leverage UAV (drone) imagery along with Convolutional Neural Networks, specifically ResNet-101, to efficiently analyze large road areas and identify defects under varying environmental conditions. By incorporating image preprocessing and transfer learning, the project ensures improved performance even with limited datasets.

Additionally, the project focuses on creating a user-friendly and deployable solution through a Flask-based web application, enabling real-time predictions. Overall, the purpose is to enhance road safety, minimize maintenance costs, and support smart infrastructure management through an efficient and reliable automated detection system.

## 1.4 Problem Statement

Despite significant advancements in deep learning, road damage detection remains a challenging problem due to several factors. Road surfaces vary widely based on geographic location, climate conditions, and construction materials, making it difficult for models to generalize effectively. In addition, road damage appears in diverse forms such as cracks, potholes, surface wear, and structural failures, which further increases the complexity of accurate classification. Variations in lighting and weather conditions, including differences between day and night or the presence of rain, fog, and shadows, can also negatively impact image quality and detection performance. Furthermore, many existing models require high computational power, making real-time implementation difficult, especially on edge devices and mobile platforms.To address these challenges, the proposed system introduces an automated road damage classification approach using Convolutional Neural Networks (CNN) with ResNet-101. By leveraging transfer learning and deep feature extraction, the system

aims to improve accuracy, robustness, and efficiency, making it suitable for real-world road monitoring applications.

## 1.5 Existing System with Disadvantages

The traditional approach to road damage detection relies heavily on manual inspection methods, where human operators physically survey roads and highways. This approach presents several limitations, including high labor intensity and time consumption, as personnel must travel extensively to identify damages and experts must manually analyze the severity and type of defects. It also involves significant financial and logistical costs due to the requirement of specialized sensors, vehicles, and trained personnel, along with continuous resource allocation for regular monitoring. Additionally, safety concerns arise as inspections are often conducted in hazardous conditions and active traffic zones, increasing the risk of accidents.

Traditional methods also suffer from limited accuracy and coverage, as they may fail to detect small or emerging cracks and are inefficient for large-scale monitoring.Furthermore, the lack of real-time data results in delayed identification of critical damage, leading to slower maintenance response, inefficient decision-making, and prolonged road deterioration.

## Disadvantages

- Manual inspection is slow, labor-intensive, and highly dependent on human involvement.
- High operational costs arise from the need for trained staff, vehicles, and continuous field monitoring.
- Inspectors face safety risks when working on damaged or high-traffic road sections.
- Traditional methods fail to detect small cracks or early damage stages, reducing overall detection accuracy.
- Limited coverage during each inspection cycle results in delayed reporting and slower maintenance response.

## 1.6 Proposed System

The proposed system leverages Unmanned Aerial Vehicles (UAVs) combined with deep learning techniques for automated road damage detection, aiming to enhance efficiency, accuracy, and scalability.

High-resolution drone cameras are used to capture images of road surfaces from multiple angles and heights, enabling rapid large-area scanning and reducing the need for on-site human intervention.

The system employs advanced object detection models CNN to automatically identify and classify different types of road damage, using a combined dataset from the RDD2022 dataset for effective training.

To improve detection accuracy, the best-performing model, ResNet101 achieves a mean average precision of 73.2% (mAP@0.5), while a modified CNN integrated with a Transformer Prediction Head further enhances object localization and classification, especially in high-density scenarios.

Additionally, the system supports automated decision-making by sending real-time notifications to maintenance teams along with geographical coordinates of detected damage, enabling faster response and efficient repair scheduling.

Furthermore, UAV deployment across multiple locations ensures real-time monitoring and scalability, making the system suitable for large-scale and continuous road infrastructure assessment.

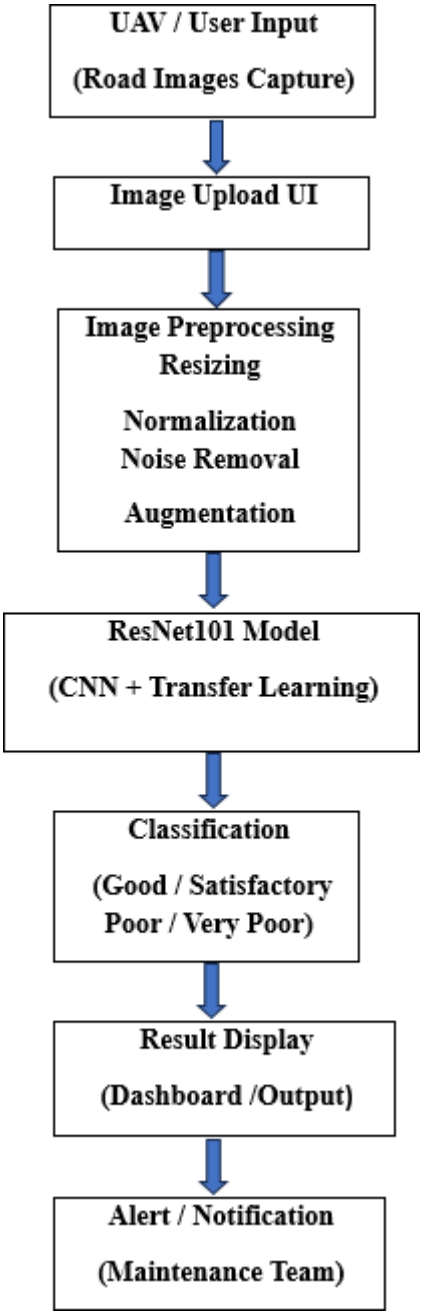


**Figure 1.6.1: Block diagram of proposed system**

## Advantages

- Enhanced reliability: Improves prediction consistency through ensemble-based decision making rather than relying on a single model.
- Robust processing capability: Effectively handles noisy, complex, and diverse input patterns using a unified preprocessing approach.
- Improved performance on imbalanced data: Achieves better accuracy on imbalanced datasets through integrated class-balancing techniques.
- Scalable architecture: Designed to support extension to multi-class and multi-label detection scenarios.

## 1.7 Input and Output Design

### 1.7.1 Input Design

Input design serves as the link between the user and the information system, focusing on developing specifications and procedures for accurate data preparation so that information is available in a usable form for processing. It can be achieved through automated data entry from printed documents or manual input provided by users. A well-structured input design helps minimize errors, avoid delays, reduce unnecessary processing steps, and ensure data security and privacy. It also involves identifying what data should be provided as input, determining how the data should be arranged or coded, designing user-friendly interactions to guide users during input, and implementing effective validation and error-handling mechanisms to maintain data accuracy and system reliability.

**Objectives of Input Design:**

- Converting user-oriented data into system-usable format: Input design ensures that user descriptions are transformed into structured data suitable for processing, thereby avoiding errors and ensuring accurate data collection.
- Creating user-friendly data entry screens: The system is designed to handle large volumes of data efficiently and provides an easy-to-use interface that minimizes user errors, while allowing users to manage and view records effectively.
- Ensuring data validity and accuracy: Input design incorporates validation mechanisms to prevent incorrect data entry, uses appropriate messages and prompts to guide users, and ensures that the input layout is simple, clear, and efficient.

### 1.7.2 Output Design

A high-quality output effectively meets user requirements and presents information in a clear and understandable manner. Outputs serve as the communication medium between the system and its users, conveying the results of system processing to both users and other systems. Output design determines how information is displayed for immediate use as well as how it is formatted for printed reports or documents. The effectiveness of output design plays a crucial role in supporting user decision-making by ensuring that the information provided is relevant, accurate, and easy to interpret.

Output design should follow well-defined principles to ensure clarity and usability. It must be organized and systematically developed so that the system generates meaningful and useful outputs. Appropriate methods for presenting information should be selected, such as reports, documents, dashboards, or visual displays, depending on user needs. Additionally, outputs should be formatted in a clear, concise, and user-friendly manner to enhance readability and ensure that users can easily understand and interpret the information provided.

The primary objectives of output design include conveying information efficiently by providing details about past activities, current status, and future projections. It should also signal important events by alerting users to critical issues, opportunities, or situations requiring immediate attention. Furthermore, outputs should help trigger and confirm actions by guiding users on necessary steps to be taken and providing confirmation that those actions have been successfully completed.

# 2. LITERATURE SURVEY

1. **Chen W., Liu H., Zhang Q., "Real-Time Road Damage Detection Using YOLOv8 for Intelligent Transportation Systems," IEEE Access, vol. 13, pp. 4567–4578, 2025.** This paper proposes a real-time road damage detection system using the YOLOv8 model. The approach focuses on improving detection speed and accuracy while reducing computational complexity. The model efficiently identifies road defects such as cracks and potholes under various environmental conditions. Experimental results show improved performance compared to earlier YOLO versions and traditional CNN models.

2. **Patel R., Kumari S., "Adaptive Road Surface Damage Detection Using Vision Transformers and UAV Imagery," IEEE Transactions on Intelligent Transportation Systems,vol.26,no.3,pp.1124–1138,2025.**
   This paper proposes a UAV-based road monitoring system using Vision Transformer (ViT) models for large-area infrastructure inspection. The authors show that transformer encoders capture long-range texture variations more effectively than CNNs. The method yields superior accuracy in identifying fine cracks and potholes across complex surfaces. Experimental results demonstrate significant improvements over existing CNN-based baselines.

3. **Rasheed S., Khan M., Ali F., "UAV-Based Road Damage Detection Using YOLO Deep Learning Approach," International Journal of Innovative Research in Technology, vol. 11,no.1,pp.102–110,2025.**
   This paper presents a UAV-based system for detecting road damages using YOLO architecture. The approach enables real-time detection of cracks and potholes from aerial images. The system improves efficiency in large-scale monitoring and reduces manual inspection effort. Results indicate high detection accuracy and practical applicability in smart city environments.

4. **Jiang Y., Wang X., Li Z., "Improved YOLOv8 with Attention Mechanism for Road Damage Detection," Springer Nature, vol. 14, no. 2, pp. 210–225, 2024.** This paper introduces an enhanced YOLOv8 model integrated with attention mechanisms to improve feature extraction. The model achieves better precision and recall in detecting road damages compared to traditional object detection methods. It effectively handles complex

backgrounds and varying lighting conditions. Experimental results show improved performance over Faster R-CNN and SSD models.

5. **Cheng C., Zhao L., Sun Y., "Multiscale Feature Fusion-Based Road Damage Detection Using Improved YOLOv5," Journal of Remote Sensing, vol. 18, no. 4, pp. 334–348, 2024.**
   This paper proposes an improved YOLOv5 model with multiscale feature fusion for UAV-based road damage detection. The approach enhances the detection of both small and large defects by combining global and local features. The model performs well on high-resolution aerial images. Results demonstrate increased accuracy and robustness

6. **Silva L.A., Pereira M., Costa R., "Automated Road Damage Detection Using UAV Images and Deep Learning Techniques," IEEE Conference on Smart Infrastructure, pp. 89–96, 2023.**
   This paper presents an automated framework combining UAV imagery with deep learning models for road damage detection. The system reduces manual inspection effort and improves detection efficiency. It performs well across diverse environmental conditions. Experimental results confirm the effectiveness of the proposed method..

7. **Zhang Y., Liu J., Chen K., "Road Damage Detection Using YOLOv3 with Multi-Level Attention Mechanism," Expert Systems with Applications, vol. 202, pp. 117–129, 2022.**
   This paper proposes a YOLOv3-based model enhanced with a multi-level attention mechanism for better feature extraction. The approach improves detection accuracy for cracks and potholes in complex road environments. It effectively captures fine-grained features from images. Results show significant improvement over baseline YOLOv3.

8. **Arya D., Maeda H., Ghosh S., "RDD2022: A Multi-Country Road Damage Detection Dataset for Deep Learning," arXiv preprint arXiv:2209.08538, 2022.**
   This paper introduces a large-scale dataset for road damage detection collected from multiple countries. The dataset supports training and evaluation of deep learning models under diverse conditions. It improves generalization capabilities of models. The study highlights the importance of standardized datasets in this domain.

9. **Arya D., Maeda H., Ghosh S., "Deep Learning-Based Road Damage Detection System for Multiple Countries," IEEE International Conference on Big Data, pp. 1–8, 2021.** This paper presents a deep learning-based system for detecting road damages across multiple countries. The approach uses transfer learning to handle variations in road conditions and environments. The system achieves reliable performance across different datasets. Results demonstrate scalability and robustness.

10. **Doshi K., Yilmaz Y., "Road Damage Detection Using Deep Ensemble Learning and YOLOv4,"arXivpreprintarXiv:2011.00728,2020.**
This paper proposes a deep ensemble learning approach combined with YOLOv4 for road damage detection. The method improves classification accuracy by combining multiple models. It effectively detects different types of road damages. Experimental results show enhanced performance compared to single-model approaches.

| FocusedArea / Title | Key Findings | Reference |
|---|---|---|
| YOLOv8 for Real-Time Road Damage Detection | Improves detection speed and accuracy while reducing computational cost. Efficiently detects cracks and potholes under differentenvironmental conditions. | Chen W., Liu H., Zhang Q., IEEE Access, 2025 |
| Vision Transformer-Based Road Damage Detection using UAV | Captures long-range texture variations better than CNNs. Achieves higher accuracy in detecting fine cracks and potholes. | Patel R., Kumari S., IEEE TITS, 2025 |
| UAV-Based Road Damage Detection using YOLO] | Enables real-time detection of road damages using aerial imagery. Reduces manual effort and improves large-scale monitoring efficiency. | Rasheed S., Khan M., IJIRT, 2025 |
| YOLOv8 with Attention Mechanism | Enhances feature extraction using attention modules. Improves precision, recall, and performance in complex environments. | Jiang Y., Wang X., Springer, 2024 |
| YOLOv5 with Multiscale Feature Fusion | Combines global and local features for better detection of small and large defects. Improves robustness in UAV images.. | Cheng C., Zhao L., Journal of Remote Sensing, 2024 |
| UAV + Deep Learning Framework | Reduces manual inspection effort and improves detection efficiency across diverse conditions | Silva L.A., Pereira M., IEEE Conference, 2023 |

| FocusedArea / Title | Key Findings | Reference |
|---|---|---|
| YOLOv3 with Multi-Level Attention | Improves feature extraction and detection accuracy for cracks and potholes in complex environments.. | Zhang Y., Liu J., Expert Systems, 2022 |
| RDD2022 Dataset for Road Damage Detection | Provides large-scale multi-country dataset for training and benchmarking models. Improves generalization capability. | Arya D., Maeda H., arXiv, 2022 |
| Multi-CountryRoad Damage Detection using Deep Learning | Uses transfer learning to handle variations across countries. Achieves scalable and robust performance. | Arya D., Maeda H., IEEE Big Data, 2021 |
| Deep Ensemble Learning with YOLOv4 | Combines multiple models to improve classification accuracy. Outperforms single-model approaches. | Doshi K., Yilmaz Y., arXiv, 2020 |

**Table no. 2**.1: Literature Review Summary

# 3. SOFTWARE REQUREIMENTS ANALYSIS

## 3.1 Modules and Their Functionalities

### 3.1.1. Data Analysis

Data analysis involves examining the structure, quality, and distribution of the dataset used for road damage detection. The dataset includes road surface images categorized into four classes—good, satisfactory, poor, and very poor. These images vary widely in terms of brightness, angle, texture, weather conditions, and camera quality.

Initial inspection often reveals significant class imbalance, where certain damage categories appear less frequently than others. The dataset may also contain noise such as shadows, uneven lighting, blurred regions, and background distractions. Understanding these characteristics is essential for designing effective preprocessing strategies, optimizing augmentation techniques, and ensuring robust classification across diverse road conditions.

### 3.1.2 Data Preprocessing

Data preprocessing is a crucial step in preparing raw UAV images for input into deep learning models, ensuring that the data is clean, consistent, and suitable for effective training. In this stage, all images are resized to a uniform resolution to maintain consistency and reduce computational complexity, while pixel intensity values are normalized to stabilize and accelerate the learning process. Noise and irrelevant background elements are removed to enhance image clarity and focus on important features related to road damage. Additionally, data augmentation techniques such as rotation, flipping, brightness adjustment, and cropping are applied to increase dataset diversity and improve model robustness. Overall, these preprocessing steps significantly improve the model's generalization capability and strengthen its ability to accurately detect road damages under varying real-world conditions such as changes in lighting, weather, and road surface textures.

### 3.1.3 Machine Learning Algorithm for prediction

The proposed system utilizes a deep learning-based approach using ResNet101, a powerful convolutional neural network pretrained on ImageNet. Through transfer learning, the model extracts rich feature representations from road images, capturing cracks, surface irregularities, and texture variations with high precision.

The model is fine-tuned on the road damage dataset to classify images into four categories. Training

optimization techniques—such as gradient clipping, learning rate scheduling, mixed precision training, and early stopping—are applied to enhance stability. The final classification layer outputs the predicted road damage class with high accuracy. This framework supports real-time deployment, enabling automated inspection for road authorities and transportation agencies.

## 3.2 Functional Requirements

The functional requirements describe the core operations that the road damage detection system performs to achieve its intended objectives. They define how the system receives and processes input images, applies preprocessing and deep learning techniques, and generates accurate classification results. These requirements ensure that the system operates reliably and efficiently while supporting real-time analysis and integration with monitoring applications. The following points summarize the primary functional capabilities implemented within the framework.

- The system shall accept input images of road surfaces captured through UAVs or uploaded by users and route them to the preprocessing module.
- The system shall perform image preprocessing operations such as resizing, normalization, and enhancement to prepare data for model analysis.
- The system shall detect and classify road damages (such as cracks and potholes) using a deep learning model based on ResNet-101.
- The system shall present the detection results in a clear, structured, and interpretable format through the user interface.

## 3.3 Non-Functional Requirements

Non-functional requirements describe the quality attributes and performance expectations that govern how the road damage detection system operates. Rather than focusing on specific features, they define how the system behaves under different conditions and how efficiently it delivers results. These requirements ensure reliability, usability, scalability, and consistency throughout the system's operation and future maintenance. The following points summarize the key non-functional characteristics implemented in the system.

- The system shall ensure high reliability and stability during all stages of image processing and damage detection.
- The system shall support scalable performance as the volume of UAV images and variety of road conditions increase.

- The system shall maintain efficient processing with minimal latency in generating detection results.
- The system shall preserve data security and confidentiality for all user-uploaded images and system outputs.

## 3.4 Feasibility Study

The feasibility of the project is analyzed in this phase, and a business proposal is put forth with a general plan for the project and cost estimates. During system analysis, the feasibility study ensures that the proposed system is viable and not a burden to the company.
Understanding the major system requirements is essential for feasibility analysis.

### 3.4.1 Economic Feasibility

The proposed system is cost-effective compared to traditional road inspection methods that require manual labor and specialized equipment. UAV-based image collection reduces the need for continuous field inspections, saving both time and manpower. Additionally, the use of open-source technologies such as Python and Flask minimizes development and deployment costs, making the system affordable for large-scale infrastructure monitoring.

### 3.4.2 Technical Feasibility

The system is technically feasible as it is built using widely used technologies such as Python, Flask, and deep learning frameworks. The ResNet-101 model is well-suited for image classification and can effectively detect road damages like cracks and potholes from UAV images. Image preprocessing techniques such as resizing, normalization, and augmentation ensure better model performance under varying conditions. The system can be developed and executed on standard computing devices with GPU support for training, making implementation practical.

### 3.4.3 Social Feasibility

The proposed road damage detection system is socially beneficial as it enhances public safety and improves infrastructure maintenance by automatically identifying road defects such as cracks and potholes using UAV imagery and the ResNet-101 model. It reduces the dependency on manual inspections, thereby minimizing human exposure to traffic risks and hazardous conditions. The system enables faster detection and repair of road damages, leading to smoother transportation, accidents.

# 4. SYSTEM SPECIFICATIONS

## 4.1 Software Requirements

The software requirements define the essential tools, frameworks, and platforms necessary for designing and developing the automated road damage detection system. The project is implemented using a stable deep learning environment that supports image preprocessing, transfer learning, and model evaluation. These tools ensure smooth experimentation, reliable performance, and scalability for future deployment stages.

1. Operating System: Windows 10 / Windows 11 / Linux (Ubuntu)
2. Programming Language: Python 3.x
3. Core Libraries: TensorFlow / Keras, OpenCV, NumPy, Pandas, Matplotlib, Scikit-learn
4. Development Environment: Jupyter Notebook / VS Code / PyCharm
5. Documentation Tools: MS Word

## 4.2 Hardware Requirements

The hardware requirements specify the minimum computational resources necessary for executing deep learning models, image processing tasks, and dataset handling operations. Standard computing hardware is suitable for the design and testing phase; however, systems with enhanced specifications can significantly improve training speed and model performance. Additional resources may be used during large-scale deployment or cloud integration.

1. Processor: Intel Core i5 / i7 or equivalent
2. Memory (RAM): Minimum 8 GB
3. Storage: 250 GB SSD / HDD
4. Display: 14" or higher
5. Optional: NVIDIA GPU for faster training and real-time inference
6. Optional: Stable internet connectivity for downloading libraries, datasets.

# 5. SOFTWARE DESIGN

## 5.1 System Architecture

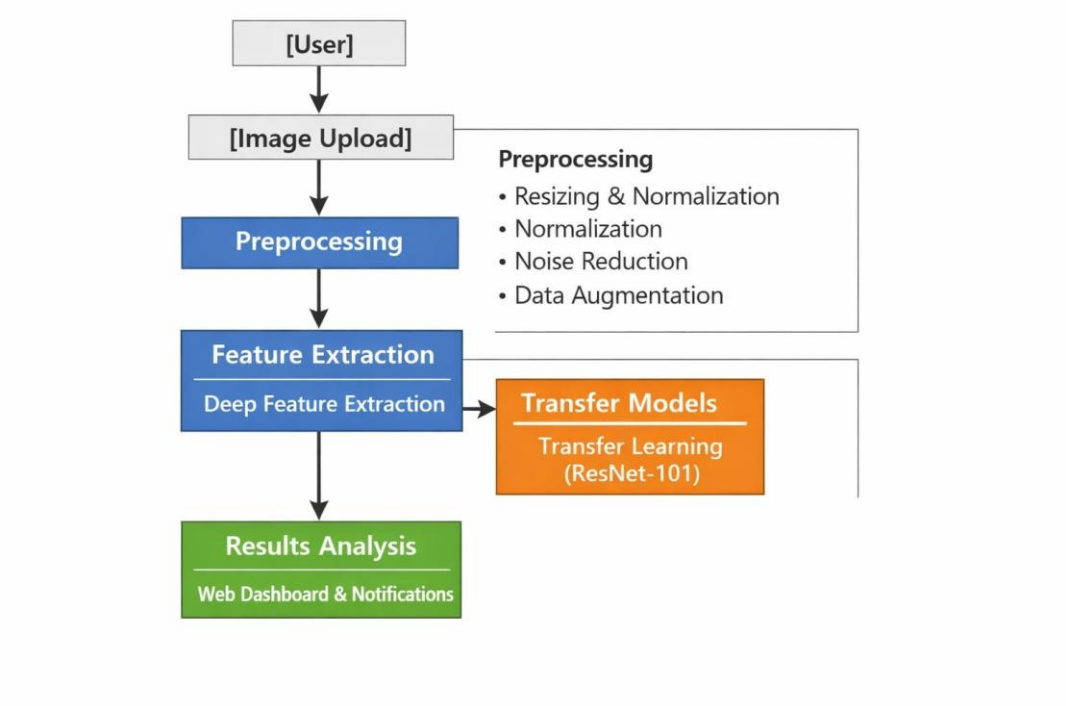


Figure: 5.1 System architecture

The system architecture establishes a structured pipeline for automated road damage detection using deep learning. The process begins with the User / Road Inspector, who captures road surface images and submits them through the Image Upload & User Interface module. This interface validates the image format and forwards the data to the Image Preprocessing Module, where operations such as resizing, normalization, noise reduction, and augmentation are performed. These steps standardize the input resolution, improve image quality, and artificially expand the dataset, ensuring that the model can handle variations in lighting, angle, texture, and road conditions.

After preprocessing, the refined images are passed to the ResNet101 CNN Model – Transfer

Learning block. Here, deep convolutional layers extract high-level features representing cracks, potholes, and surface irregularities. The fine-tuned model generates prediction scores that are interpreted by the Classification & Confidence module, which assigns each image to one of the four classes: Good, Satisfactory, Poor, or Very Poor. The final decision, along with its confidence value, is presented to the user through the Result Display – Dashboard / Console, enabling quick assessment of road conditions.

To support continuous improvement, the architecture includes persistent storage components: the Road Image Dataset, which maintains training and evaluation images, and Model Weights & Logs, which stores parameters, checkpoints, and performance metrics. An Admin module interacts with these repositories to manage the dataset and trigger Retrain / Update Model operations whenever new labeled data becomes available. This modular and extensible design ensures that the system remains scalable, maintainable, and ready for integration into real-time road monitoring workflows.

## 5.2 Dataflow Diagram

The Data Flow Diagram Level 0 (DFD–0) provides a high-level view of the information flow in the proposed road damage detection system. At this level, the entire functionality is represented as a single process, "Road Damage Classification System," interacting with external entities and primary data stores. The diagram begins with the User / Road Inspector, who captures road surface images and submits them to the system for analysis.

The Data Flow Diagram Level 1 (DFD–1) expands the main Road Damage Classification System process into a series of interconnected sub-processes that depict the detailed flow of data within the application. The accepted image is then passed to the Preprocessing Module, which performs resizing, normalization, noise reduction, and augmentation operations to standardize the input and enhance the robustness of the model. The preprocessed image is subsequently routed to the Model Inference (ResNet101 CNN) process. Here, the system loads pretrained weights from the Model Weights & Logs data store and performs feature extraction and classification. The resulting prediction scores are then transferred to the Result Generation & Storage process, where the final road condition label— Good, Satisfactory, Poor, or Very Poor—is derived and formatted.

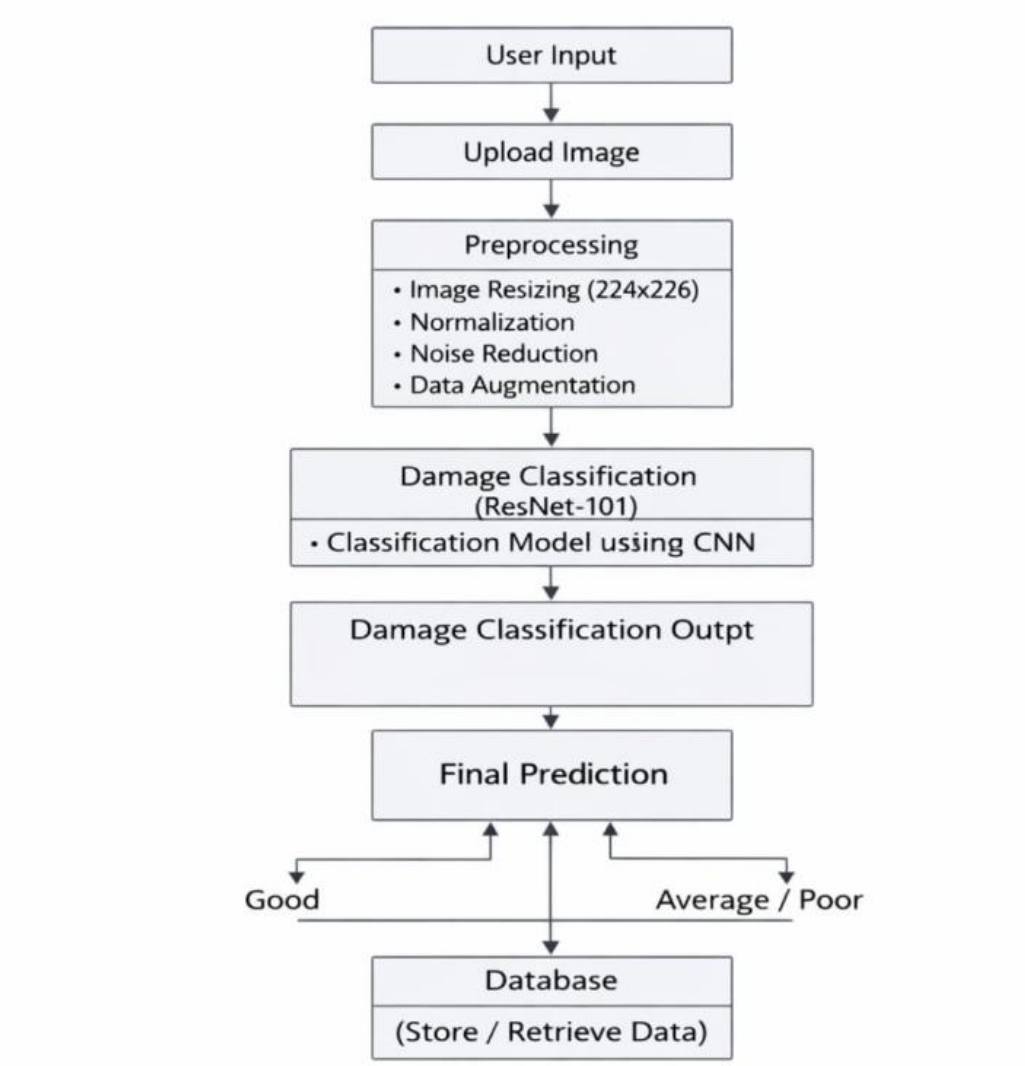


Figure: 5.2.1 Data flow diagram

## 5.3 UML Diagrams

UML (Unified Modeling Language) is a standardized visual language used for specifying, constructing, and documenting the artifacts of software systems. It provides a common notation for developers, analysts, and stakeholders to describe both the static structure and dynamic behavior of an application. In this project, UML is used to represent the design of the automated road damage detection system, capturing how users interact with the system, how components communicate, and how data flows through different modules such as image preprocessing and the ResNet101-based

classifier.

UML is maintained by the Object Management Group (OMG) and is closely associated with object-oriented analysis and design. The diagrams are broadly divided into behavioral diagrams, which describe how the system behaves over time, and structural diagrams, which describe the organization of classes, components, and deployment elements.

Goals of UML:

- Provide an expressive visual modeling language for developing and exchanging meaningful system models.
- Establish a common, formal basis for understanding and documenting system design.
- Support object-oriented analysis and design principles in a standardized way.
- Encourage best practices, modularity, and reusability during system development.

Types of UML Diagrams Used in This Project:

1. Sequence Diagram
2. Use Case Diagram
3. Activity Diagram
4. Class Diagram
5. Collaboration Diagram
6. Component Diagram
7. Deployment Diagram

### 5.3.1 Sequence Diagram

The sequence diagram represents the interaction between different components of the Automated Road Damage Detection System in a time-ordered manner. It shows how messages are passed between various system elements to complete the damage detection process. The sequence begins with the UAV capturing road images and transmitting them to the system. The system receives the image and forwards it to the ImageProcessor for preprocessing operations such as resizing and normalization.

After preprocessing, the processed image is sent to the Model, which loads the trained ResNet101 model and performs damage detection. The model analyzes the image and returns the detection results.

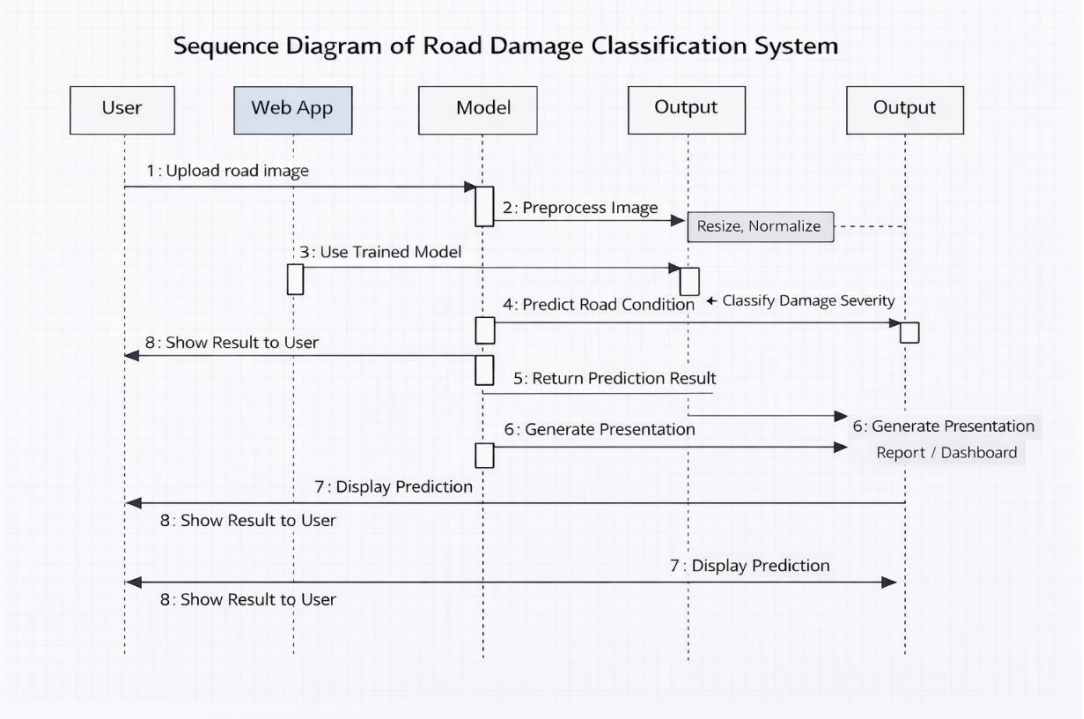


Figure: 5.3.1 Sequence Diagram

Once preprocessing is completed, the System passes the refined image to the Model, which loads the pretrained ResNet101 model and performs the prediction. The Model sends the predicted class label and confidence score to the Result, which formats the output and optionally stores the result and image details in the Database. Finally, the Result sends the final response back to the User Interface, which displays the predicted road condition as Good, Satisfactory, Poor, or Very Poor to the User. This sequence clearly shows the order of message passing between objects and highlights how the system components collaborate to complete a single road damage classification request.

### 5.3.2 Use case Diagram

The use case diagram represents the functional requirements of the Automated Road Damage Detection System and illustrates the interactions between different actors and the system.

The primary actors involved in the system are the User, the UAV System, and the Administrator. The UAV system is responsible for capturing road images and transmitting them to the main system. The user interacts with the system to upload images (if required), initiate the detection process, and view the results. The administrator manages system operations such as dataset handling and model maintenance.

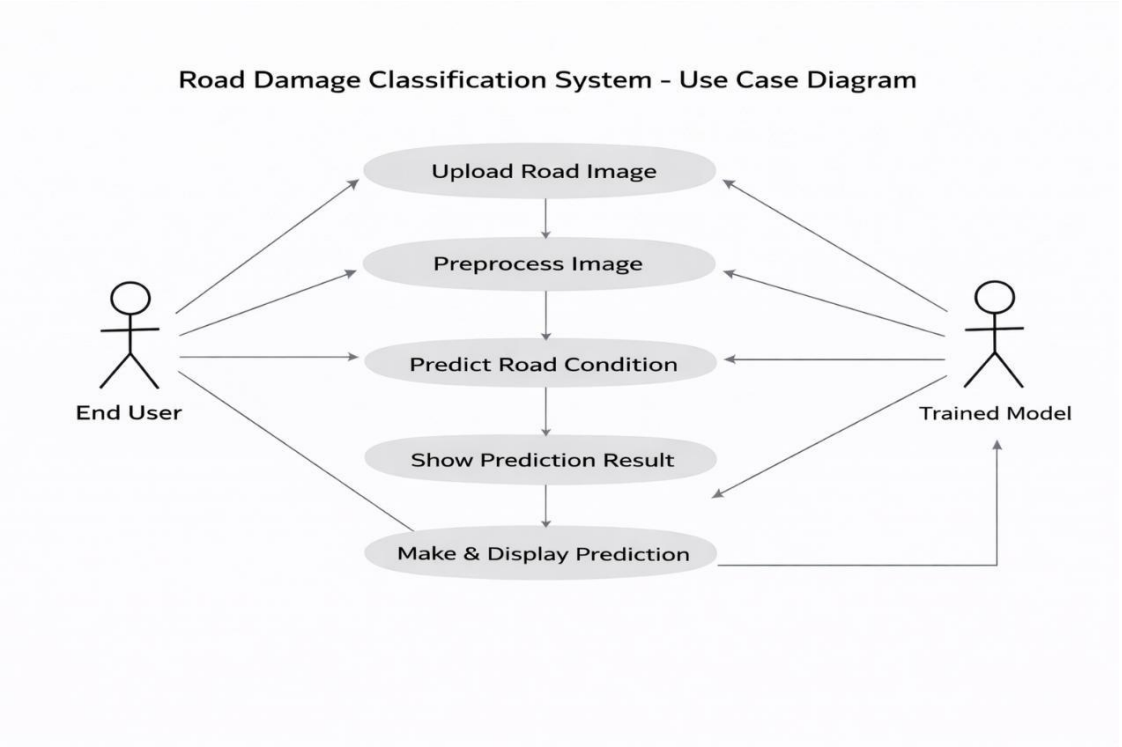


Figure: 5.3.2 Use case Diagram

An additional actor, the Admin, manages system maintenance tasks. The Admin can curate and update the road image dataset, initiate model retraining with newly labeled data, and monitor overall system performance. Through these use cases, the diagram clearly represents how each actor interacts with the system to achieve specific goals—capturing images, classifying road conditions, maintaining datasets, and improving model accuracy. This visual model helps in understanding the functional boundaries of the system and ensures that all critical user requirements are addressed in the design.

### 5.3.3 Activity Diagram

The class diagram represents the structural design of the Automated Road Damage Detection System by illustrating the key classes, their attributes, methods, and relationships.The main class in the system is the System class, which acts as the central controller. It is responsible for receiving input images and displaying the final detection results.The UAV class handles image acquisition by capturing road images and transmitting them to the system. The ImageProcessor class is responsible for preprocessing tasks such as resizing and normalizing the images before they are fed into the model.

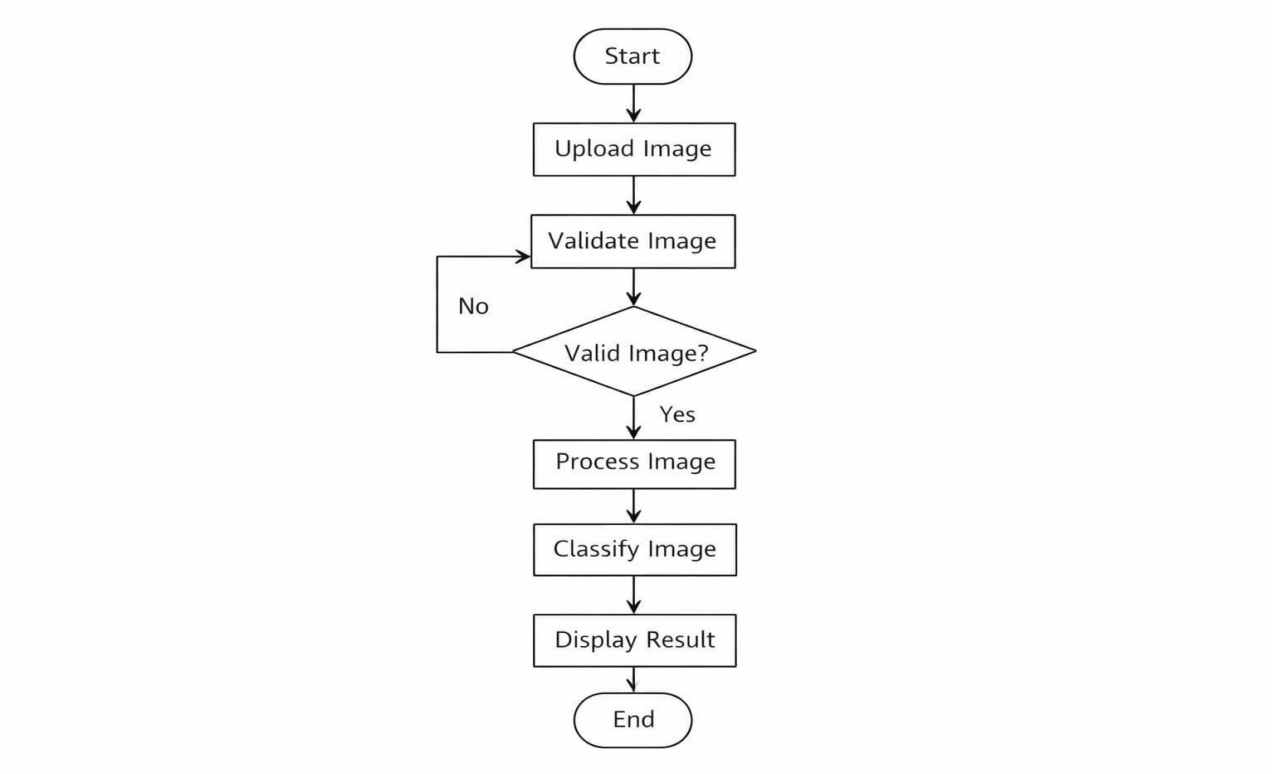


Figure: 5.3.3 Activity Diagram

After preprocessing, the activity diagram proceeds to Load ResNet101 Model and then to Perform Damage Classification, where the trained model predicts the road condition category. A Generate Result activity converts the model output into a user-readable form, including the predicted class such as Good, Satisfactory, Poor, or Very Poor. The next activity is Display Result to User, where the classification outcome is shown on the interface. Optionally, the flow may pass through Store Image and Result in Database for maintaining inspection history.

### 5.3.4 Class Diagram

The class diagram represents the static structure of the automated road damage detection system by showing the main classes, their attributes, methods, and relationships. In this project, the central class is the RoadDamageSystem class, which coordinates the overall workflow of image acquisition, preprocessing, model inference, and result generation. It collaborates with classes such as User, Admin, ImageManager, PreprocessingModule, ModelHandler, and ResultManager to complete the operations required for road damage classification.

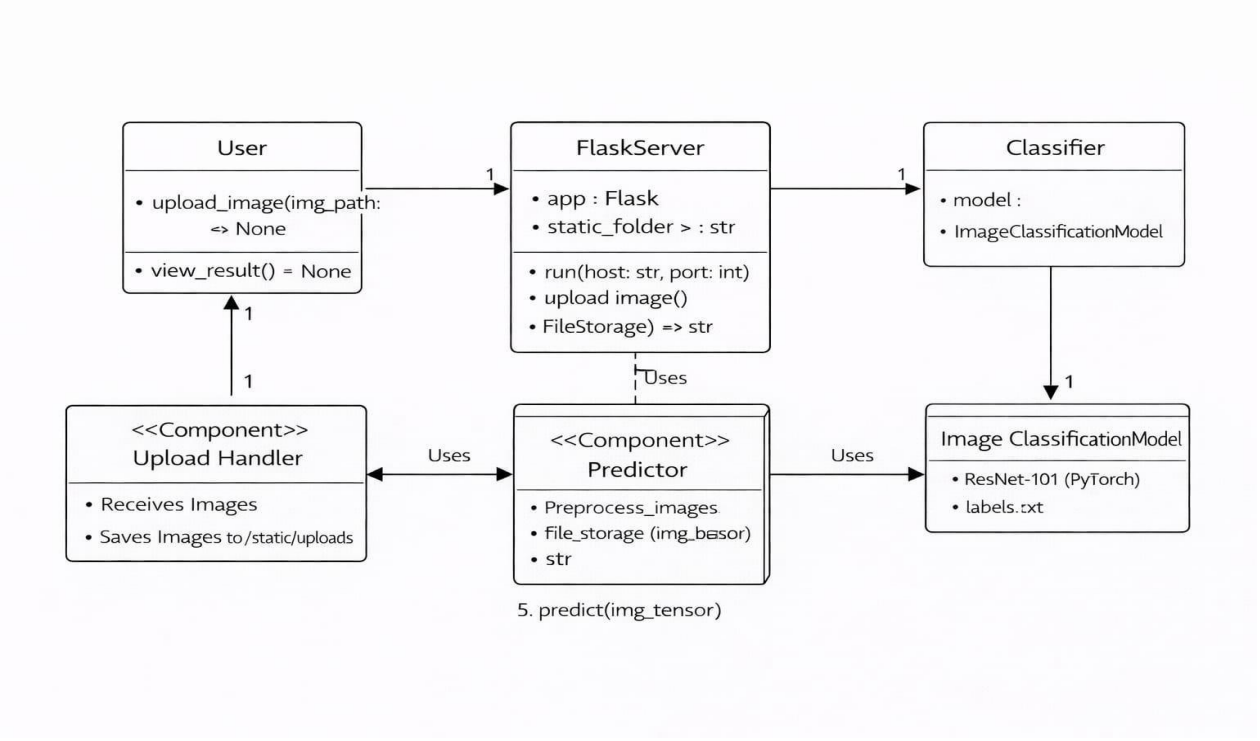


Figure: 5.3.4 Class Diagram

The User class represents the road inspector who uploads road images and views the classification results, while the Admin class extends the User with additional privileges such as managing datasets and retraining the model. The ImageManager class handles functions like uploadImage(), validateImage(), and storeImage(). The PreprocessingModule class includes methods such as resizeImage(), normalizeImage(), and augmentImage() to prepare data for the model. The ModelHandler class is responsible for loading the pretrained ResNet101 model, performing predictDamageLevel(), and updating model weights when required.

### 5.3.5 Collaboration Diagram

The collaboration diagram illustrates how different components of the Automated Road Damage Detection System interact with each other to accomplish the damage detection process. Unlike the sequence diagram, which focuses on the order of interactions, the collaboration diagram emphasizes the relationships and communication between system components. In this system, the UAV interacts with the main system by capturing and sending road images. The system coordinates with the ImageProcessor to perform preprocessing operations such as resizing and normalization of images.The processed data is then communicated to the ModelHandler, which utilizes the trained ResNet101 model to detect road damages. The results generated by the model are passed to the DetectionResult component, where the output is structured.

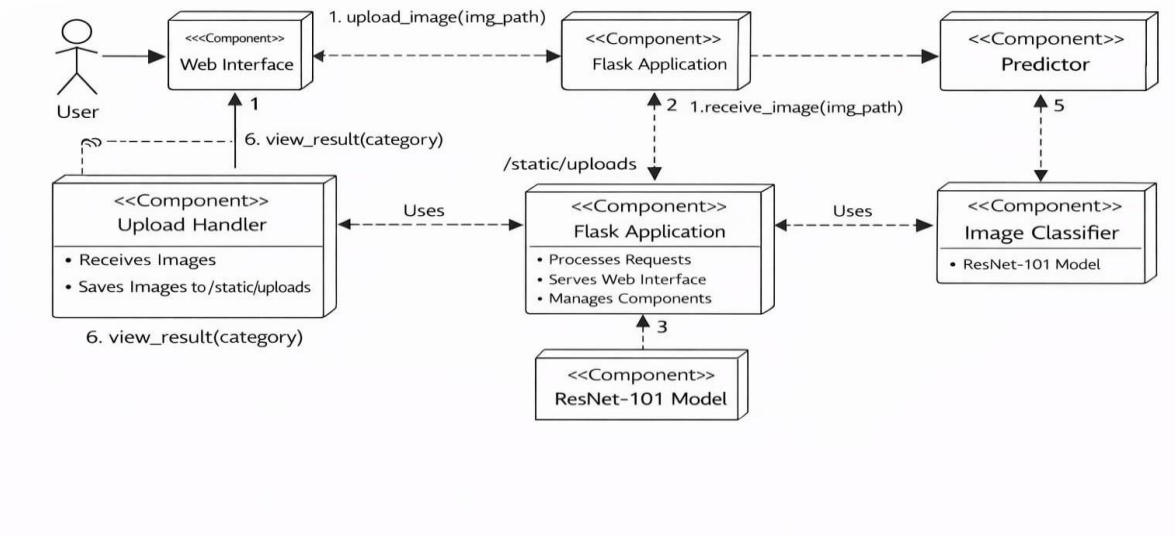


Figure: 5.3.5 Collaboration Diagram

After the image has been preprocessed, the PreprocessingModule returns it to the ImageManager, which then sends another message to the ModelHandler to classify the image using the pretrained ResNet101 model. The ModelHandler processes the image, predicts the corresponding damage level, and forwards the result to the ResultManager. The ResultManager may interact with the Database object to store the image, prediction label, and timestamp for future reference. Finally, the ResultManager sends the formatted output back to the User Interface, which displays the predicted road condition to the User.

### 5.3.6 Component Diagram

The component diagram represents the high-level architecture of the Automated Road Damage Detection System by illustrating the major system components and their interactions. The system is divided into several key components, including the UAV module, frontend interface, backend server, preprocessing module, deep learning model, dataset repository, and result display module. The UAV component is responsible for capturing road images and sending them to the backend system. The frontend component provides an interface for users to interact with the system, such as uploading images and viewing results. The backend server acts as the core processing unit, coordinating various modules.

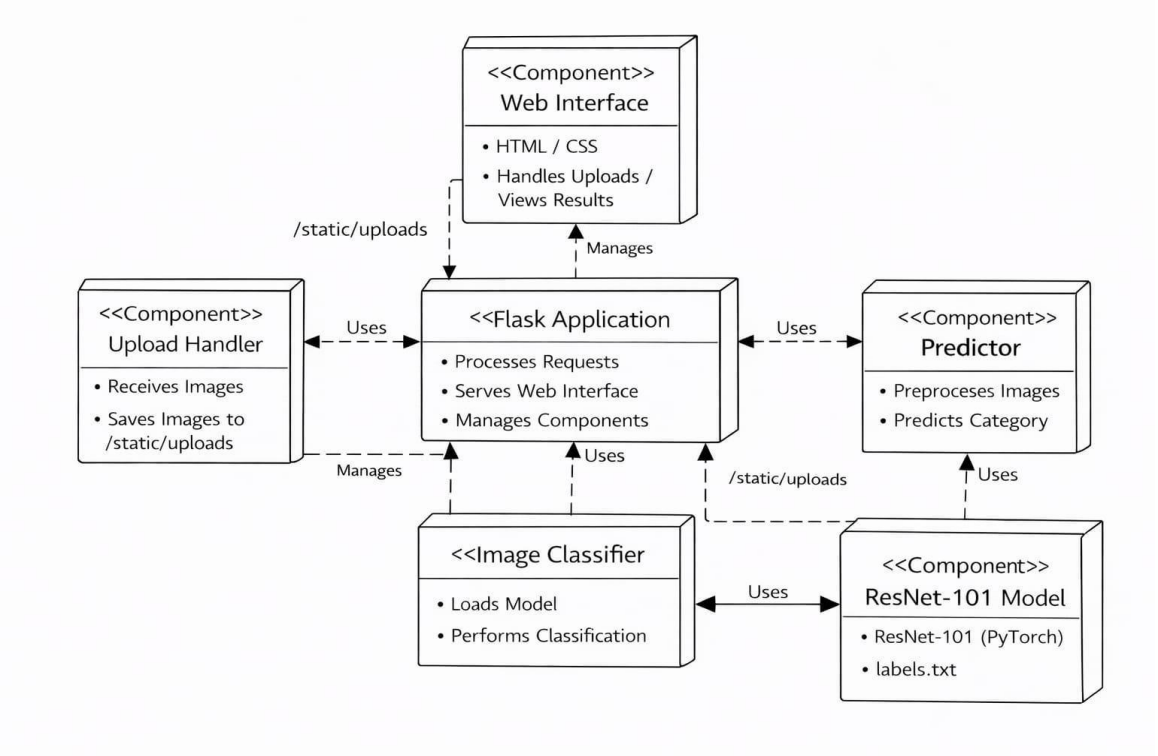


Figure: 5.3.6 Component Diagram

The User Interface Component handles user interactions such as uploading road images and viewing classification results. It is connected to the Image Processing Component, which encapsulates image validation and preprocessing operations like resizing, normalization, and augmentation. The processed image is then passed to the Model Component, which contains the trained ResNet101 deep learning model and the logic for performing prediction. After classification, the output is sent to the Result Management Component, which formats the result, generates reports, and coordinates with the Data Storage Component to store images, predicted labels, user details, and logs. The Admin or Maintenance Tools Component can also interact with the Model and Data Storage components to update datasets and retrain the model when new data becomes available.

### 5.3.7 Deployment Diagram

The deployment diagram represents the physical deployment of the automated road damage detection system on various hardware nodes. It shows how the software components are mapped onto the underlying infrastructure used during development and, optionally, during real-time usage. In this project, the system is primarily deployed on a Client Machine used by the Road Inspector or User and, if required, on a Server / Cloud Node for hosting the trained deep learning model and storing large datasets. This diagram helps in understanding where each part of the application runs and how different devices communicate during road damage detection.

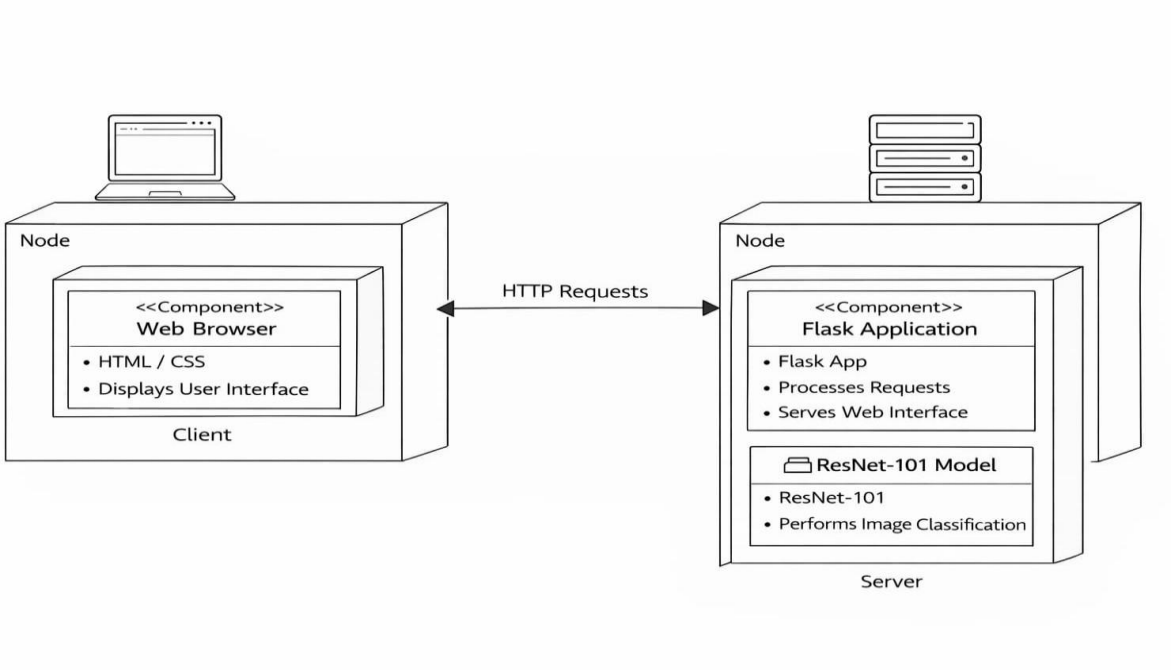


Figure: 5.3.7 Deployment Diagram

On the Client Node (User System), the User Interface Application is deployed, which allows the user to upload road images and view classification results. This node runs the front-end interface and may also execute lightweight preprocessing tasks. The Application Server / Processing Node hosts the core modules such as the Image Processing Service, ResNet101 Model Service, and Result Management Service, which perform image preprocessing, deep learning–based classification, and result generation. A separate Database Server node (or cloud storage) is used to store road images, model weights, logs, and inspection history. Communication between these nodes occurs over a network connection, enabling the client system to send image data to the processing server and receive predicted road condition labels in response.

# 6. CODING AND IMPLEMENTATION

## 6.1 Source Code:

**drawIOAppForChrome.html:**

```
<!DOCTYPE html>
<html lang="en">
  <head>
    <meta charset="UTF-8" />
    <meta name="viewport" content="width=device-width, initial-scale=1.0" />
    <title>drawio</title>
    <style>
      body,
      html {
        margin: 0;
        padding: 0;
        height: 100%;
        overflow: hidden;
      }
      #app {
        width: 100%;
        height: 100%;
      }
    </style>
  </head>
  <body style="height:100%;width:100%">
    <iframe
      frameborder="0"
      id="app"
      src="http://draw.io?embed=1&ui=atlas&spin=1&proto=json&lang="
    ></iframe>
    <script>
      var $app=document.getElementById('app');
```

```
        </script>
    </body>
</html>
```

**model.ipynb :**

```
import os
import glob
import random
import shutil
import matplotlib.pyplot as plt
import cv2
import numpy as np
from sklearn.model_selection import train_test_split
import torch
from torchvision import transforms
from torch.utils.data import DataLoader, Dataset
from PIL import Image

# Set dataset path
dataset_path = "/kaggle/input/road-damage-classification-and-assessment/sih_road_dataset"

# Define class names (folder names)
classes = ["good", "poor", "satisfactory", "very_poor"]

# Create a list to store image paths and their corresponding labels
image_paths = []
labels = []

# Collect all images and assign labels
for idx, class_name in enumerate(classes):
    class_path = os.path.join(dataset_path, class_name)
    image_files = glob.glob(os.path.join(class_path, "*.*"))  # Load all image formats
```

```
    for img_path in image_files:
        image_paths.append(img_path)
        labels.append(idx)  # Label is index of class in `classes` list

# Shuffle dataset
data = list(zip(image_paths, labels))
random.shuffle(data)
image_paths, labels = zip(*data)

# Split dataset into train (70%), validation (20%), and test (10%)
train_paths, temp_paths, train_labels, temp_labels = train_test_split(
    image_paths, labels, test_size=0.3, stratify=labels, random_state=42
)
val_paths, test_paths, val_labels, test_labels = train_test_split(
    temp_paths, temp_labels, test_size=0.33, stratify=temp_labels, random_state=42
)  # 33% of 30% ≈ 10%

print(f"Total images: {len(image_paths)}")
print(f"Train size: {len(train_paths)}, Validation size: {len(val_paths)}, Test size: {len(test_paths)}")
```

```
Total images: 2074
Train size: 1451, Validation size: 417, Test size: 206
```

```
# Define Image Transformations
image_transforms = {
    "train": transforms.Compose([
        transforms.Resize((224, 224)),  # Resize images
        transforms.RandomHorizontalFlip(),
        transforms.RandomRotation(15),
        transforms.ToTensor(),
```

```
        transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])
    ]),
    "val": transforms.Compose([
        transforms.Resize((224, 224)),
        transforms.ToTensor(),
        transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])
    ]),
    "test": transforms.Compose([
        transforms.Resize((224, 224)),
        transforms.ToTensor(),
        transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])
    ])
}

# Define Custom Dataset Class
class RoadDamageDataset(Dataset):
    def __init__(self, image_paths, labels, transform=None):
        self.image_paths = image_paths
        self.labels = labels
        self.transform = transform

    def __len__(self):
        return len(self.image_paths)

    def __getitem__(self, idx):
        img_path = self.image_paths[idx]
        label = self.labels[idx]

        # Read Image
        image = cv2.imread(img_path)
        image = cv2.cvtColor(image, cv2.COLOR_BGR2RGB)  # Convert BGR to RGB
        # Convert to PIL Image (Fix for torchvision transforms)
        image = Image.fromarray(image)
```

```
        # Apply transformations
        if self.transform:
            image = self.transform(image)

        return image, label

# Create datasets
train_dataset = RoadDamageDataset(train_paths, train_labels, transform=image_transforms["train"])
val_dataset = RoadDamageDataset(val_paths, val_labels, transform=image_transforms["val"])
test_dataset = RoadDamageDataset(test_paths, test_labels, transform=image_transforms["test"])

# Create Data Loaders
batch_size = 32
train_loader = DataLoader(train_dataset, batch_size=batch_size, shuffle=True)
val_loader = DataLoader(val_dataset, batch_size=batch_size, shuffle=False)
test_loader = DataLoader(test_dataset, batch_size=batch_size, shuffle=False)

# Function to visualize sample images
def visualize_sample_images(data_loader, class_names):
    images, labels = next(iter(data_loader))
    images = images.numpy().transpose((0, 2, 3, 1))  # Convert to numpy array
    fig, axes = plt.subplots(1, 5, figsize=(15, 5))

    for i in range(5):
        img = images[i]
        img = (img * 0.229 + 0.485)  # De-normalization
        axes[i].imshow(img)
        axes[i].set_title(class_names[labels[i]])
        axes[i].axis("off")

    plt.show()
# Visualize training samples
visualize_sample_images(train_loader, classes)
```


```
import torch
import torch.nn as nn
import torch.optim as optim
import torchvision.models as models
from torchsummary import summary

# Check for available GPUs
device = torch.device("cuda" if torch.cuda.is_available() else "cpu")
print(f"Using device: {device}")

# Load Pretrained ResNet101
resnet101 = models.resnet101(weights=models.ResNet101_Weights.DEFAULT)

# Freeze Initial Layers (First 50 layers)
for param in list(resnet101.parameters())[:150]:  # Freeze first 150 layers
    param.requires_grad = False

# Modify the Final Classification Layer
num_features = resnet101.fc.in_features
resnet101.fc = nn.Sequential(
    nn.Linear(num_features, 512),  # Extra Dense Layer
    nn.ReLU(),
    nn.Dropout(0.5),  # Regularization
    nn.Linear(512, 4)  # Output 4 classes
)
```

```
# Move Model to GPU
resnet101 = resnet101.to(device)

# Enable Multi-GPU Training if multiple GPUs are available
if torch.cuda.device_count() > 1:
    print(f"Using {torch.cuda.device_count()} GPUs for training.")
    resnet101 = nn.DataParallel(resnet101)

# Print Model Summary
summary(resnet101, (3, 224, 224))

# Define Loss Function with Label Smoothing (Prevents Overconfidence)
loss_function = nn.CrossEntropyLoss(label_smoothing=0.1)

# Define Optimizer - AdamW (Better than Adam)
optimizer = optim.AdamW(resnet101.parameters(), lr=1e-4, weight_decay=1e-4)

# Learning Rate Scheduler - Cosine Annealing (For better convergence)
scheduler = optim.lr_scheduler.CosineAnnealingLR(optimizer, T_max=10, eta_min=1e-6)

# Function for Gradient Clipping (Prevents Exploding Gradients)
def clip_gradient(optimizer, grad_clip):
    for group in optimizer.param_groups:
        for param in group["params"]:
            if param.grad is not None:
                param.grad.data.clamp_(-grad_clip, grad_clip)

print("Model is ready for training!")
```

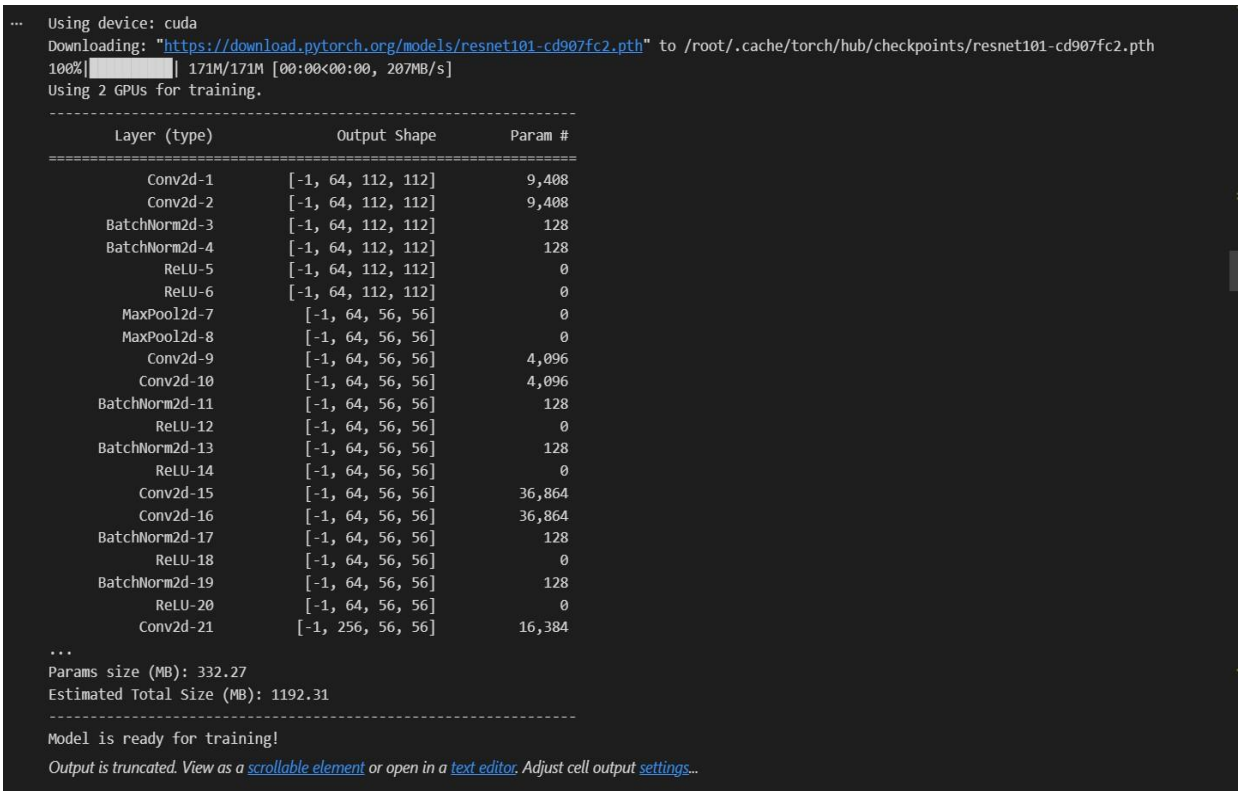

```
Using device: cuda
Downloading: "https://download.pytorch.org/models/resnet101-cd907fc2.pth" to /root/.cache/torch/hub/checkpoints/resnet101-cd907fc2.pth
100%|██████████| 171M/171M [00:00<00:00, 207MB/s]
Using 2 GPUs for training.
----------------------------------------------------------------
        Layer (type)               Output Shape         Param #
================================================================
            Conv2d-1        [-1, 64, 112, 112]           9,408
            Conv2d-2        [-1, 64, 112, 112]           9,408
       BatchNorm2d-3        [-1, 64, 112, 112]             128
       BatchNorm2d-4        [-1, 64, 112, 112]             128
              ReLU-5        [-1, 64, 112, 112]               0
              ReLU-6        [-1, 64, 112, 112]               0
         MaxPool2d-7          [-1, 64, 56, 56]               0
         MaxPool2d-8          [-1, 64, 56, 56]               0
            Conv2d-9          [-1, 64, 56, 56]           4,096
           Conv2d-10          [-1, 64, 56, 56]           4,096
      BatchNorm2d-11          [-1, 64, 56, 56]             128
             ReLU-12          [-1, 64, 56, 56]               0
      BatchNorm2d-13          [-1, 64, 56, 56]             128
             ReLU-14          [-1, 64, 56, 56]               0
           Conv2d-15          [-1, 64, 56, 56]          36,864
           Conv2d-16          [-1, 64, 56, 56]          36,864
      BatchNorm2d-17          [-1, 64, 56, 56]             128
             ReLU-18          [-1, 64, 56, 56]               0
      BatchNorm2d-19          [-1, 64, 56, 56]             128
             ReLU-20          [-1, 64, 56, 56]               0
           Conv2d-21         [-1, 256, 56, 56]          16,384
...
Params size (MB): 332.27
Estimated Total Size (MB): 1192.31
----------------------------------------------------------------
Model is ready for training!
Output is truncated. View as a scrollable element or open in a text editor. Adjust cell output settings...
```

```
import time
import copy
from tqdm import tqdm

# Define training loop
def train_model(model, train_loader, val_loader, criterion, optimizer, scheduler, num_epochs=25, grad_clip=1.0):
    since = time.time()

    best_model_wts = copy.deepcopy(model.state_dict())
    best_acc = 0.0
    early_stopping_counter = 0
    patience = 5  # Stop training if no improvement after 5 epochs

    scaler = torch.cuda.amp.GradScaler()  # Mixed Precision Training

    for epoch in range(num_epochs):
        print(f'\nEpoch {epoch+1}/{num_epochs}')
```

```
print('-' * 50)

# Each epoch has a training and validation phase
for phase in ['train', 'val']:
    if phase == 'train':
        model.train()  # Set model to training mode
        dataloader = train_loader
    else:
        model.eval()  # Set model to evaluation mode
        dataloader = val_loader

    running_loss = 0.0
    running_corrects = 0
    total_samples = 0

    # Iterate over data
    loop = tqdm(dataloader, desc=f"{phase} Phase", leave=False)
    for inputs, labels in loop:
        inputs, labels = inputs.to(device), labels.to(device)

        # Forward pass with Mixed Precision
        with torch.cuda.amp.autocast():
            outputs = model(inputs)
            loss = criterion(outputs, labels)

        _, preds = torch.max(outputs, 1)
        total_samples += labels.size(0)
        running_corrects += torch.sum(preds == labels.data)
        running_loss += loss.item() * inputs.size(0)

        if phase == 'train':
            optimizer.zero_grad()
            scaler.scale(loss).backward()  # Scale loss for mixed precision
            scaler.unscale_(optimizer)  # Unscale gradients before clipping
```

```
                torch.nn.utils.clip_grad_norm_(model.parameters(), grad_clip)  # Gradient Clipping
                scaler.step(optimizer)  # Optimizer step
                scaler.update()  # Update the scaler

            # Adjust learning rate after an epoch (for train phase)
            if phase == 'train':
                scheduler.step()

            # Calculate epoch loss and accuracy
            epoch_loss = running_loss / total_samples
            epoch_acc = running_corrects.double() / total_samples

            print(f'{phase} Loss: {epoch_loss:.4f} Acc: {epoch_acc:.4f}')

            # Deep copy the model if validation accuracy improves
            if phase == 'val' and epoch_acc > best_acc:
                best_acc = epoch_acc
                best_model_wts = copy.deepcopy(model.state_dict())
                early_stopping_counter = 0  # Reset early stopping counter
            elif phase == 'val':
                early_stopping_counter += 1  # Increment if no improvement

            # Early stopping if validation accuracy doesn't improve for 'patience' epochs
            if early_stopping_counter >= patience:
                print("Early stopping triggered. Stopping training.")
                break

    time_elapsed = time.time() - since
    print(f'\nTraining complete in {time_elapsed // 60:.0f}m {time_elapsed % 60:.0f}s')
    print(f'Best Validation Acc: {best_acc:.4f}')

    # Load best model weights
    model.load_state_dict(best_model_wts)
    return model
```

```
# Set Number of Epochs
num_epochs = 30  # You can adjust based on performance

# Train Model
trained_model = train_model(
    model=resnet101,
    train_loader=train_loader,
    val_loader=val_loader,
    criterion=loss_function,
    optimizer=optimizer,
    scheduler=scheduler,
    num_epochs=num_epochs
)

# Save Best Model
torch.save(trained_model.state_dict(), "best_resnet101_model.pth")
print("Best model saved as best_resnet101_model.pth")
```

```
<ipython-input-10-1ea23fb810df>:14: FutureWarning: `torch.cuda.amp.GradScaler(args...)` is deprecated. Please use `torch.amp.GradScaler('cuda'
  scaler = torch.cuda.amp.GradScaler()  # Mixed Precision Training

Epoch 1/30
----------
train Phase:   0%|          | 0/46 [00:00<?, ?it/s]<ipython-input-10-1ea23fb810df>:39: FutureWarning: `torch.cuda.amp.autocast(args...)` is de
  with torch.cuda.amp.autocast():

train Loss: 0.8423 Acc: 0.7664

val Loss: 0.4703 Acc: 0.9496

Epoch 2/30
----------

train Loss: 0.4752 Acc: 0.9449

val Loss: 0.4067 Acc: 0.9760

Epoch 3/30
----------

train Loss: 0.4239 Acc: 0.9717

val Loss: 0.4138 Acc: 0.9712

Epoch 4/30
----------

train Loss: 0.4005 Acc: 0.9855

val Loss: 0.4078 Acc: 0.9712

Epoch 5/30
```

```
----------

train Loss: 0.3849 Acc: 0.9890

val Loss: 0.4107 Acc: 0.9712

Epoch 6/30
----------

train Loss: 0.3746 Acc: 0.9924

val Loss: 0.4106 Acc: 0.9760

Epoch 7/30
----------

train Loss: 0.3692 Acc: 0.9945

val Loss: 0.4108 Acc: 0.9760
Early stopping triggered. Stopping training.

Epoch 8/30
----------

train Loss: 0.3664 Acc: 0.9966
Early stopping triggered. Stopping training.

Epoch 9/30
----------

train Loss: 0.3664 Acc: 0.9959
Early stopping triggered. Stopping training.

Epoch 10/30
----------
```

```
----------

train Loss: 0.3671 Acc: 0.9959
Early stopping triggered. Stopping training.

Epoch 11/30
----------

train Loss: 0.3631 Acc: 0.9986
Early stopping triggered. Stopping training.

Epoch 12/30
----------

train Loss: 0.3657 Acc: 0.9972
Early stopping triggered. Stopping training.

Epoch 13/30
----------

train Loss: 0.3619 Acc: 0.9993
Early stopping triggered. Stopping training.

Epoch 14/30
----------

train Loss: 0.3693 Acc: 0.9931
Early stopping triggered. Stopping training.

Epoch 15/30
----------
```

```
----------

train Loss: 0.3660 Acc: 0.9952
Early stopping triggered. Stopping training.

Epoch 16/30
----------

train Loss: 0.3646 Acc: 0.9972
Early stopping triggered. Stopping training.

Epoch 17/30
----------

train Loss: 0.3660 Acc: 0.9952
Early stopping triggered. Stopping training.

Epoch 18/30
----------

train Loss: 0.3674 Acc: 0.9952
Early stopping triggered. Stopping training.

Epoch 19/30
----------

train Loss: 0.3692 Acc: 0.9952
Early stopping triggered. Stopping training.

Epoch 20/30
----------
```

```
train Loss: 0.3664 Acc: 0.9952
Early stopping triggered. Stopping training.

Epoch 21/30
----------

train Loss: 0.3628 Acc: 0.9959
Early stopping triggered. Stopping training.

Epoch 22/30
----------

train Loss: 0.3619 Acc: 0.9972
Early stopping triggered. Stopping training.

Epoch 23/30
----------

train Loss: 0.3610 Acc: 0.9972
Early stopping triggered. Stopping training.

Epoch 24/30
----------

train Loss: 0.3596 Acc: 0.9979
Early stopping triggered. Stopping training.

Epoch 25/30
```

```
train Loss: 0.3584 Acc: 0.9979
Early stopping triggered. Stopping training.

Epoch 26/30
----------

train Loss: 0.3588 Acc: 0.9972
Early stopping triggered. Stopping training.

Epoch 27/30
----------

train Loss: 0.3562 Acc: 0.9993
Early stopping triggered. Stopping training.

Epoch 28/30
----------

train Loss: 0.3599 Acc: 0.9972
Early stopping triggered. Stopping training.

Epoch 29/30
----------

train Loss: 0.3557 Acc: 0.9993
Early stopping triggered. Stopping training.

Epoch 30/30
----------

train Loss: 0.3560 Acc: 0.9993
Early stopping triggered. Stopping training.

Training complete in 20m 46s
Best Validation Acc: 0.9760
Best model saved as best_resnet101_model.pth
```

```
from sklearn.metrics import accuracy_score, classification_report, confusion_matrix
import seaborn as sns
import numpy as np

# Load the Best Model
model_path = "best_resnet101_model.pth"
resnet101.load_state_dict(torch.load(model_path))
resnet101.eval()  # Set model to evaluation mode

# Function to Evaluate the Model on the Test Set
def evaluate_model(model, test_loader, class_names):
    y_true = []
    y_pred = []

    with torch.no_grad():
        for inputs, labels in tqdm(test_loader, desc="Evaluating on Test Set"):
            inputs, labels = inputs.to(device), labels.to(device)

            # Forward Pass
            outputs = model(inputs)
            _, preds = torch.max(outputs, 1)

            y_true.extend(labels.cpu().numpy())
            y_pred.extend(preds.cpu().numpy())

    # Compute Accuracy
    accuracy = accuracy_score(y_true, y_pred)
    print(f"\nTest Accuracy: {accuracy:.4f}")

    # Generate Classification Report
    print("\nClassification Report:")
    print(classification_report(y_true, y_pred, target_names=class_names))
```

```
    # Generate Confusion Matrix
    cm = confusion_matrix(y_true, y_pred)
    plt.figure(figsize=(8, 6))
        sns.heatmap(cm, annot=True, fmt="d", cmap="Blues", xticklabels=class_names, yticklabels=class_names)
    plt.xlabel("Predicted")
    plt.ylabel("Actual")
    plt.title("Confusion Matrix")
    plt.show()

    return y_true, y_pred

# Evaluate on Test Set
y_true, y_pred = evaluate_model(resnet101, test_loader, classes)
```

```
<ipython-input-11-ee5ddb0622c5>:7: FutureWarning: You are using `torch.load` with `weights_only=False` (the current default value), which uses
  resnet101.load_state_dict(torch.load(model_path))
Evaluating on Test Set: 100%|██████████| 7/7 [00:06<00:00,  1.03it/s]

Test Accuracy: 0.9515

Classification Report:
              precision    recall  f1-score   support

        good       0.99      1.00      0.99        84
        poor       0.89      0.87      0.88        39
satisfactory       0.90      0.92      0.91        51
   very_poor       1.00      0.97      0.98        32

    accuracy                           0.95       206
   macro avg       0.95      0.94      0.94       206
weighted avg       0.95      0.95      0.95       206
```

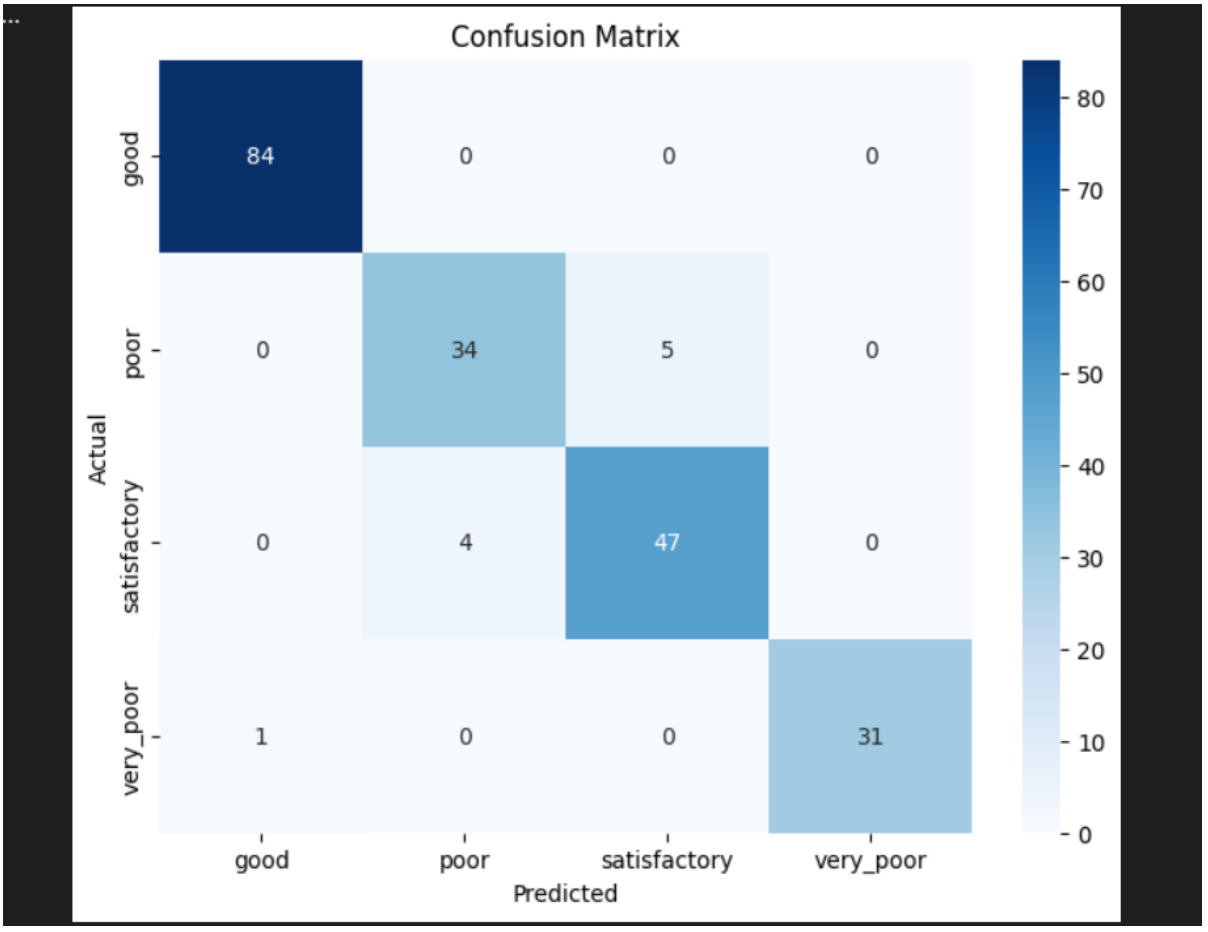


**predict.py :**

```
import torch
import torchvision.transforms as transforms
import torchvision.models as models
from PIL import Image
import os
import torch.nn as nn

# Define class names
classes = ["good", "poor", "satisfactory", "very_poor"]

# Function to Load Model
def load_model(model_path, device):
    model = models.resnet101(weights=None)  # Load ResNet101 structure
    num_features = model.fc.in_features
```

```
    # Match the exact FC structure used during training
    model.fc = nn.Sequential(
        nn.Linear(num_features, 512),
        nn.ReLU(),
        nn.Dropout(0.5),
        nn.Linear(512, len(classes))
    )

    # Load model checkpoint correctly
    state_dict = torch.load(model_path, map_location=device)

    # If model was trained using DataParallel, remove "module." prefix
    if "module." in list(state_dict.keys())[0]:  # Check if 'module.' exists
        from collections import OrderedDict
        new_state_dict = OrderedDict()
        for k, v in state_dict.items():
            name = k.replace("module.", "")  # Remove 'module.' prefix
            new_state_dict[name] = v
        state_dict = new_state_dict

    model.load_state_dict(state_dict)
    model.to(device)
    model.eval()  # Set model to evaluation mode
    return model

# Define Image Transformation (Same as Training)
transform = transforms.Compose([
    transforms.Resize((224, 224)),
    transforms.ToTensor(),
    transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])
])
```

```
# Function to Predict Class for a Given Image
def predict_image(model, image_path, device):
    image = Image.open(image_path).convert("RGB")  # Open image
    image = transform(image).unsqueeze(0)  # Apply transformations & add batch dimension
    image = image.to(device)

    # Run inference
    with torch.no_grad():
        output = model(image)
        _, predicted_class = torch.max(output, 1)  # Get class with highest probability
        predicted_label = classes[predicted_class.item()]

    return predicted_label

# Set Device (Use GPU if available)
device = torch.device("cuda" if torch.cuda.is_available() else "cpu")

# Load Model
model_path = "best_resnet101_model.pth"  # Ensure this file is in the same directory as predict.py
if not os.path.exists(model_path):
    print(f"Error: Model file '{model_path}' not found.")
    exit()

model = load_model(model_path, device)

# **Directly specify the image path**
image_path = r"sih_road_dataset\poor\poor_042.jpg"  # Change this path accordingly

if not os.path.exists(image_path):
    print(f"Error: Image file '{image_path}' not found.")
    exit()
# Predict Image Class
prediction = predict_image(model, image_path, device)
print(f"Predicted Class: {prediction}")
```

**index.html :**

```
<!DOCTYPE html>
<html lang="en">
<head>
   <meta charset="UTF-8">
   <meta name="viewport" content="width=device-width, initial-scale=1">
   <title>Road Damage Classification</title>
<linkrel="stylesheet"
href="https://cdn.jsdelivr.net/npm/bootstrap@5.3.0/dist/css/bootstrap.min.css">
<script src="https://cdn.jsdelivr.net/npm/bootstrap@5.3.0/dist/js/bootstrap.bundle.min.js"></script>
   <style>
      * {
         margin: 0;
         padding: 0;
         box-sizing: border-box;
      }

      body {
         background: linear-gradient(135deg, #667eea 0%, #764ba2 50%, #f093fb 100%);
         min-height: 100vh;
         font-family: 'Segoe UI', Tahoma, Geneva, Verdana, sans-serif;
         overflow-x: hidden;
      }

      /* Animated background particles */
      .particles {
         position: fixed;
         top: 0;
         left: 0;
         width: 100%;
         height: 100%;
         pointer-events: none;
         z-index: 0;
```

```
}
.particle {
    position: absolute;
    border-radius: 50%;
    animation: float 15s infinite ease-in-out;
    opacity: 0.3;
}

@keyframes float {
    0%, 100% { transform: translateY(0) translateX(0) rotate(0deg); }
    25% { transform: translateY(-100px) translateX(50px) rotate(90deg); }
    50% { transform: translateY(-50px) translateX(100px) rotate(180deg); }
    75% { transform: translateY(-150px) translateX(25px) rotate(270deg); }
}

/* Navigation */
nav {
    background: rgba(255, 255, 255, 0.1) !important;
    backdrop-filter: blur(20px);
    border-bottom: 1px solid rgba(255, 255, 255, 0.2);
    box-shadow: 0 8px 32px rgba(0, 0, 0, 0.1);
    transition: all 0.3s ease;
}

nav:hover {
    background: rgba(255, 255, 255, 0.15) !important;
}

.navbar-brand {
    font-size: 1.5rem;
    font-weight: 700;
    color: #1e293b;
    text-shadow: 0 2px 4px rgba(0, 0, 0, 0.1);
    transition: all 0.3s ease;
```

```
}

.navbar-brand:hover {
    color: #0f172a;
    text-shadow: 0 2px 8px rgba(0, 0, 0, 0.2);
}

/* Main container */
.container {
    position: relative;
    z-index: 1;
    margin-top: 80px;
}

/* Hero card */
.hero-card {
    background: rgba(255, 255, 255, 0.95);
    border-radius: 30px;
    padding: 50px;
    box-shadow: 0 30px 80px rgba(0, 0, 0, 0.2);
    backdrop-filter: blur(10px);
    border: 1px solid rgba(255, 255, 255, 0.3);
    animation: slideUp 0.8s ease-out;
    transition: transform 0.3s ease, box-shadow 0.3s ease;
}

.hero-card:hover {
    transform: translateY(-10px);
    box-shadow: 0 40px 100px rgba(0, 0, 0, 0.3);
}

@keyframes slideUp {
    from {
        opacity: 0;
```

```
        transform: translateY(50px);
    }
    to {
        opacity: 1;
        transform: translateY(0);
    }
}

h2 {
    background: linear-gradient(135deg, #ff6ec4, #7873f5, #ff9a56);
    -webkit-background-clip: text;
    -webkit-text-fill-color: transparent;
    background-clip: text;
    font-weight: 800;
    font-size: 2.5rem;
    margin-bottom: 40px;
    animation: glow 2s ease-in-out infinite;
}

@keyframes glow {
    0%, 100% { filter: brightness(1); }
    50% { filter: brightness(1.2); }
}

/* File upload area */
.upload-area {
    border: 3px dashed #ff6ec4;
    border-radius: 20px;
    padding: 60px 30px;
    background: linear-gradient(135deg, rgba(255, 110, 196, 0.1), rgba(120, 115, 245, 0.1));
    transition: all 0.4s cubic-bezier(0.68, -0.55, 0.265, 1.55);
    cursor: pointer;
    position: relative;
    overflow: hidden;
```

```
}

.upload-area::before {
    content: '';
    position: absolute;
    top: -50%;
    left: -50%;
    width: 200%;
    height: 200%;
    background: linear-gradient(45deg, transparent, rgba(255, 110, 196, 0.3), transparent);
    animation: shine 3s infinite;
}

@keyframes shine {
    0% { transform: rotate(0deg); }
    100% { transform: rotate(360deg); }
}

.upload-area:hover {
    border-color: #7873f5;
    background: linear-gradient(135deg, rgba(120, 115, 245, 0.2), rgba(255, 110, 196, 0.2));
    transform: scale(1.02);
    box-shadow: 0 15px 40px rgba(255, 110, 196, 0.3);
}

.upload-area.dragover {
    border-color: #7873f5;
    background: linear-gradient(135deg, rgba(120, 115, 245, 0.3), rgba(255, 110, 196, 0.3));
    transform: scale(1.05);
}

.upload-icon {
    font-size: 4rem;
    color: #ff6ec4;
```

```
    margin-bottom: 20px;
    animation: bounce 2s infinite;
}

@keyframes bounce {
    0%, 100% { transform: translateY(0); }
    50% { transform: translateY(-20px); }
}

.form-control {
    display: none;
}

/* Button styling */
.btn-predict {
    background: linear-gradient(135deg, #ff6ec4, #7873f5);
    border: none;
    padding: 15px 50px;
    font-size: 1.2rem;
    font-weight: 600;
    border-radius: 50px;
    color: white;
    box-shadow: 0 10px 30px rgba(255, 110, 196, 0.4);
    transition: all 0.3s ease;
    margin-top: 30px;
    position:  relative;
    overflow: hidden;
}

.btn-predict::before {
    content: '';
    position: absolute;
    top: 50%;
    left: 50%;
```

```
    width: 0;
    height: 0;
    border-radius: 50%;
    background: rgba(255, 255, 255, 0.3);
    transform: translate(-50%, -50%);
    transition: width 0.6s, height 0.6s;
}

.btn-predict:hover::before {
    width: 300px;
    height: 300px;
}

.btn-predict:hover {
    transform: translateY(-5px);
    box-shadow: 0 15px 40px rgba(255, 110, 196, 0.6);
}

.btn-predict:active {
    transform: translateY(-2px);
}

/* Loading animation */
.loader {
    display: none;
    margin-top: 40px;
}

.loading-spinner {
    width: 80px;
    height: 80px;
    border: 6px solid rgba(255, 110, 196, 0.2);
    border-top: 6px solid #ff6ec4;
    border-radius: 50%;
```

```
    animation: spin 1s linear infinite;
    margin: 0 auto;
}

@keyframes spin {
    0% { transform: rotate(0deg); }
    100% { transform: rotate(360deg); }
}

.loading-text {
    margin-top: 20px;
    color: #ff6ec4;
    font-size: 1.2rem;
    font-weight: 600;
    animation: pulse 1.5s ease-in-out infinite;
}

@keyframes pulse {
    0%, 100% { opacity: 1; }
    50% { opacity: 0.5; }
}

/* Result display */
.result-card {
    background: linear-gradient(135deg, rgba(255, 255, 255, 0.95), rgba(255, 255, 255, 0.9));
    border-radius: 25px;
    padding: 40px;
    margin-top: 40px;
    box-shadow: 0 20px 60px rgba(0, 0, 0, 0.15);
    animation: fadeIn 0.6s ease-out;
}

@keyframes fadeIn {
    from {
```

```
        opacity: 0;
        transform: scale(0.9);
    }
    to {
        opacity: 1;
        transform: scale(1);
    }
}

.result-image {
    border-radius: 20px;
    box-shadow: 0 15px 40px rgba(0, 0, 0, 0.2);
    transition: transform 0.3s ease;
    max-width: 400px;
    animation: zoomIn 0.6s ease-out;
}

@keyframes zoomIn {
    from {
        opacity: 0;
        transform: scale(0.5);
    }
    to {
        opacity: 1;
        transform: scale(1);
    }
}

.result-image:hover {
    transform: scale(1.05) rotate(2deg);
}

.prediction-badge {
    background: linear-gradient(135deg, #10b981, #34d399);
```

```
    color: white;
    padding: 15px 40px;
    border-radius: 50px;
    font-size: 1.5rem;
    font-weight: 700;
    display: inline-block;
    margin-top: 20px;
    box-shadow: 0 10px 30px rgba(16, 185, 129, 0.4);
    animation: popIn 0.6s cubic-bezier(0.68, -0.55, 0.265, 1.55);
}

@keyframes popIn {
    0% {
        opacity: 0;
        transform: scale(0);
    }
    50% {
        transform: scale(1.1);
    }
    100% {
        opacity: 1;
        transform: scale(1);
    }
}

.alert {
    border-radius: 15px;
    animation: shake 0.5s ease-in-out;
}

@keyframes shake {
    0%, 100% { transform: translateX(0); }
    25% { transform: translateX(-10px); }
    75% { transform: translateX(10px); }
```

```
}

/* File preview */
.file-preview {
    margin-top: 20px;
    display: none;
}

.preview-image {
    max-width: 200px;
    border-radius: 15px;
    box-shadow: 0 10px 30px rgba(0, 0, 0, 0.2);
    animation: fadeIn 0.4s ease-out;
}

.file-name {
    margin-top: 15px;
    color: #5b21b6;
    font-weight: 600;
    font-size: 1.1rem;
}

/* Footer */
footer {
    background: rgba(255, 255, 255, 0.1);
    backdrop-filter: blur(20px);
    border-top: 1px solid rgba(255, 255, 255, 0.2);
    color: #1e293b;
    text-align: center;
    padding: 30px 20px;
    margin-top: 80px;
    font-weight: 600;
    position: relative;
    z-index: 1;
```

```
        box-shadow: 0 -8px 32px rgba(0, 0, 0, 0.1);
      }

      footer:hover {
        background: rgba(255, 255, 255, 0.15);
      }
    </style>
</head>
<body>
    <!-- Animated Background Particles -->
    <div class="particles" id="particles"></div>

    <!-- Navigation Bar -->
    <nav class="navbar navbar-dark">
      <div class="container-fluid">
        <a class="navbar-brand" href="#">E Road Damage Classifier</a>
      </div>
    </nav>

    <div class="container text-center">
      <div class="hero-card">
        <h2>Upload an Image for Prediction</h2>

        <!-- Image Upload Form -->
        <form id="upload-form" action="/predict" method="post" enctype="multipart/form-data">
          <div class="upload-area" id="upload-area">
            <div class="upload-icon">
            <h4 style="color: #5b21b6; font-weight: 600;">Drop your image here</h4>
            <p style="color: #6b7280; margin-top: 10px;">or click to browse</p>
                <input class="form-control" type="file" name="file" id="file" accept="image/*"
required>
          </div>
```

```
            <div class="file-preview" id="file-preview">
                <img id="preview-image" class="preview-image" src="" alt="Preview">
                <div class="file-name" id="file-name"></div>
            </div>

            <button type="submit" class="btn btn-predict" id="predict-btn">
                <span>🔍 Analyze Image</span>
            </button>
        </form>

        <!-- Loading Indicator -->
        <div class="loader" id="loading">
            <div class="loading-spinner"></div>
            <p class="loading-text">Analyzing road damage...</p>
        </div>

        <!-- Display Uploaded Image and Prediction -->
        {% if image_path %}
            <div class="result-card">
                    <h4 style="color: #5b21b6; font-weight: 700; margin-bottom: 30px;">Analysis Complete!</h4>
                <img src="{{ image_path }}" class="img-fluid result-image">
                <div class="prediction-badge">{{ prediction }}</div>
            </div>
        {% endif %}

        {% if error %}
            <div class="alert alert-danger mt-4">
                <strong>⚠ Error:</strong> {{ error }}
            </div>
        {% endif %}
    </div>
</div>
```

```
<footer>
    © 2025 Road Damage Detection | Built with £     by Shashidhar & Team
</footer>

<script>
    // Create animated particles
    const particlesContainer = document.getElementById('particles');
    const colors = ['#ff6ec4', '#7873f5', '#ffd93d', '#6bcf7f'];

    for (let i = 0; i < 20; i++) {
        const particle = document.createElement('div');
        particle.className = 'particle';
        particle.style.width = Math.random() * 100 + 50 + 'px';
        particle.style.height = particle.style.width;
        particle.style.left = Math.random() * 100 + '%';
        particle.style.top = Math.random() * 100 + '%';
        particle.style.background = colors[Math.floor(Math.random() * colors.length)];
        particle.style.animationDelay = Math.random() * 5 + 's';
        particle.style.animationDuration = Math.random() * 10 + 10 + 's';
        particlesContainer.appendChild(particle);
    }

    // File upload interactions
    const uploadArea = document.getElementById('upload-area');
    const fileInput = document.getElementById('file');
    const filePreview = document.getElementById('file-preview');
    const previewImage = document.getElementById('preview-image');
    const fileName = document.getElementById('file-name');

    uploadArea.addEventListener('click', () => fileInput.click());

    uploadArea.addEventListener('dragover', (e) => {
```

```
    e.preventDefault();
    uploadArea.classList.add('dragover');
});

uploadArea.addEventListener('dragleave', () => {
    uploadArea.classList.remove('dragover');
});

uploadArea.addEventListener('drop', (e) => {
    e.preventDefault();
    uploadArea.classList.remove('dragover');
    const files = e.dataTransfer.files;
    if (files.length > 0) {
        fileInput.files = files;
        handleFileSelect(files[0]);
    }
});

fileInput.addEventListener('change', (e) => {
    if (e.target.files.length > 0) {
        handleFileSelect(e.target.files[0]);
    }
});

function handleFileSelect(file) {
    const reader = new FileReader();
    reader.onload = (e) => {
        previewImage.src = e.target.result;
        fileName.textContent = file.name;
        filePreview.style.display = 'block';
    };
    reader.readAsDataURL(file);
}
```

```
        // Form submission
        document.getElementById('upload-form').onsubmit = function() {
            document.getElementById('loading').style.display = 'block';
            document.getElementById('predict-btn').disabled = true;
        };
    </script>
</body>
</html>
```

**__init__.py :**

```
from __future__import annotations

__version__= "25.2"

def main(args: list[str] | None = None) -> int:
    """This is an internal API only meant for use by pip's own console scripts.

    For additional details, see https://github.com/pypa/pip/issues/7498.
    """
    from pip._internal.utils.entrypoints import _wrapper

    return _wrapper(args)
```

**__main__.py :**

```
import os
import sys

# Remove '' and current working directory from the first entry
# of sys.path, if present to avoid using current directory
# in pip commands check, freeze, install, list and show,
```

```
# when invoked as python -m pip <command>
if sys.path[0] in ("", os.getcwd()):
    sys.path.pop(0)

# If we are running from a wheel, add the wheel to sys.path
# This allows the usage python pip-*.whl/pip install pip-*.whl
if __package__ == "":
    # __file__ is pip-*.whl/pip/__main__.py
    # first dirname call strips of '/__main__.py', second strips off '/pip'
    # Resulting path is the name of the wheel itself
    # Add that to sys.path so we can import pip
    path = os.path.dirname(os.path.dirname(__file__))
    sys.path.insert(0, path)

if __name__ == "__main__":
    from pip._internal.cli.main import main as _main

    sys.exit(_main())
```

**__pip-runner__.py :**

```
"""Execute exactly this copy of pip, within a different environment.

This file is named as it is, to ensure that this module can't be imported via
an import statement.
"""

# /!\ This version compatibility check section must be Python 2 compatible. /!\

import sys

# Copied from pyproject.toml
PYTHON_REQUIRES = (3, 9)
```

```
def version_str(version):  # type: ignore
    return ".".join(str(v) for v in version)

if sys.version_info[:2] < PYTHON_REQUIRES:
    raise SystemExit(
        "This version of pip does not support python {} (requires >={}).".format(
            version_str(sys.version_info[:2]), version_str(PYTHON_REQUIRES)
        )
    )

# From here on, we can use Python 3 features, but the syntax must remain
# Python 2 compatible.

import runpy  # noqa: E402
from importlib.machinery import PathFinder  # noqa: E402
from os.path import dirname  # noqa: E402

PIP_SOURCES_ROOT = dirname(dirname(__file__))

class PipImportRedirectingFinder:
    @classmethod
    def find_spec(self, fullname, path=None, target=None):  # type: ignore
        if fullname != "pip":
            return None

        spec = PathFinder.find_spec(fullname, [PIP_SOURCES_ROOT], target)
        assert spec, (PIP_SOURCES_ROOT, fullname)
        return spec

sys.meta_path.insert(0, PipImportRedirectingFinder())

assert __name__ == "__main__", "Cannot run __pip-runner__.py as a non-main module"
runpy.run_module("pip", run_name="__main__", alter_sys=True)
```

**activate.bat :**

```
@echo off

rem This file is UTF-8 encoded, so we need to update the current code page while executing it
for /f "tokens=2 delims=:." %%a in ('"%SystemRoot%\System32\chcp.com"') do (
   set _OLD_CODEPAGE=%%a
)
if defined _OLD_CODEPAGE (
   "%SystemRoot%\System32\chcp.com" 65001 > nul
)

set "VIRTUAL_ENV=C:\Users\shash\Desktop\Automated Road Damage Detection Using UAV Images and Deep Learning Techniques\App\venv"

if not defined PROMPT set PROMPT=$P$G

if defined _OLD_VIRTUAL_PROMPT set PROMPT=%_OLD_VIRTUAL_PROMPT%
if defined _OLD_VIRTUAL_PYTHONHOME set PYTHONHOME=%_OLD_VIRTUAL_PYTHONHOME%

set "_OLD_VIRTUAL_PROMPT=%PROMPT%"
set "PROMPT=(venv) %PROMPT%"

if defined PYTHONHOME set _OLD_VIRTUAL_PYTHONHOME=%PYTHONHOME%
set PYTHONHOME=

if defined _OLD_VIRTUAL_PATH set PATH=%_OLD_VIRTUAL_PATH%
if not defined _OLD_VIRTUAL_PATH set _OLD_VIRTUAL_PATH=%PATH%

set "PATH=%VIRTUAL_ENV%\Scripts;%PATH%"
set "VIRTUAL_ENV_PROMPT=venv"

:END
```

```
if defined _OLD_CODEPAGE (
    "%SystemRoot%\System32\chcp.com" %_OLD_CODEPAGE% > nul
    set _OLD_CODEPAGE=
)
```

**activate.fish :**

```
# This file must be used with "source <venv>/bin/activate.fish" *from fish*
# (https://fishshell.com/). You cannot run it directly.

function deactivate  -d "Exit virtual environment and return to normal shell environment"
    # reset old environment variables
    if test -n "$_OLD_VIRTUAL_PATH"
        set -gx PATH $_OLD_VIRTUAL_PATH
        set -e _OLD_VIRTUAL_PATH
    end
    if test -n "$_OLD_VIRTUAL_PYTHONHOME"
        set -gx PYTHONHOME $_OLD_VIRTUAL_PYTHONHOME
        set -e _OLD_VIRTUAL_PYTHONHOME
    end

    if test -n "$_OLD_FISH_PROMPT_OVERRIDE"
        set -e _OLD_FISH_PROMPT_OVERRIDE
        # prevents error when using nested fish instances (Issue #93858)
        if functions -q _old_fish_prompt
            functions -e fish_prompt
            functions -c _old_fish_prompt fish_prompt
            functions -e _old_fish_prompt
        end
    end

    set -e VIRTUAL_ENV
    set -e VIRTUAL_ENV_PROMPT
    if test "$argv[1]" != "nondestructive"
```

```
        # Self-destruct!
        functions -e deactivate
    end
end

# Unset irrelevant variables.
deactivate nondestructive

set -gx VIRTUAL_ENV 'C:\Users\shash\Desktop\Automated Road Damage Detection Using UAV
Images and Deep Learning Techniques\App\venv'

set -gx _OLD_VIRTUAL_PATH $PATH
set -gx PATH "$VIRTUAL_ENV/"Scripts $PATH
set -gx VIRTUAL_ENV_PROMPT venv

# Unset PYTHONHOME if set.
if set -q PYTHONHOME
    set -gx _OLD_VIRTUAL_PYTHONHOME $PYTHONHOME
    set -e PYTHONHOME
end

if test -z "$VIRTUAL_ENV_DISABLE_PROMPT"
    # fish uses a function instead of an env var to generate the prompt.

    # Save the current fish_prompt function as the function _old_fish_prompt.
    functions -c fish_prompt _old_fish_prompt

    # With the original prompt function renamed, we can override with our own.
    function fish_prompt
        # Save the return status of the last command.
        set -l old_status $status

        # Output the venv prompt; color taken from the blue of the Python logo.
        printf "%s(%s)%s " (set_color 4B8BBE) venv (set_color normal)
```

```
        # Restore the return status of the previous command.
        echo "exit $old_status" | .
        # Output the original/"old" prompt.
        _old_fish_prompt
    end

    set -gx _OLD_FISH_PROMPT_OVERRIDE "$VIRTUAL_ENV"
end
```

**Activate.ps1 :**

```
<#
.Synopsis
Activate a Python virtual environment for the current PowerShell session.

.Description
Pushes the python executable for a virtual environment to the front of the
$Env:PATH environment variable and sets the prompt to signify that you are
in a Python virtual environment. Makes use of the command line switches as
well as the `pyvenv.cfg` file values present in the virtual environment.

.Parameter VenvDir
Path to the directory that contains the virtual environment to activate. The
default value for this is the parent of the directory that the Activate.ps1
script is located within.

.Parameter Prompt
The prompt prefix to display when this virtual environment is activated. By
default, this prompt is the name of the virtual environment folder (VenvDir)
surrounded by parentheses and followed by a single space (ie. '(.venv) ').

.Example
Activate.ps1
```

Activates the Python virtual environment that contains the Activate.ps1 script.

.Example
Activate.ps1 -Verbose
Activates the Python virtual environment that contains the Activate.ps1 script,
and shows extra information about the activation as it executes.

.Example
Activate.ps1 -VenvDir C:\Users\MyUser\Common\.venv
Activates the Python virtual environment located in the specified location.

.Example
Activate.ps1 -Prompt "MyPython"
Activates the Python virtual environment that contains the Activate.ps1 script,
and prefixes the current prompt with the specified string (surrounded in
parentheses) while the virtual environment is active.

.Notes
On Windows, it may be required to enable this Activate.ps1 script by setting the
execution policy for the user. You can do this by issuing the following PowerShell
command:

PS C:\> Set-ExecutionPolicy -ExecutionPolicy RemoteSigned -Scope CurrentUser

For more information on Execution Policies:
https://go.microsoft.com/fwlink/?LinkID=135170

#>
Param(
    [Parameter(Mandatory = $false)]
    [String]
    $VenvDir,
    [Parameter(Mandatory = $false)]
    [String]

```
    $Prompt
)

<# Function declarations --------------------------------- #>

<#
.Synopsis
Remove all shell session elements added by the Activate script, including the
addition of the virtual environment's Python executable from the beginning of
the PATH variable.

.Parameter NonDestructive
If present, do not remove this function from the global namespace for the
session.

#>
function global:deactivate ([switch]$NonDestructive) {
    # Revert to original values

    # The prior prompt:
    if (Test-Path -Path Function:_OLD_VIRTUAL_PROMPT) {
        Copy-Item -Path Function:_OLD_VIRTUAL_PROMPT -Destination Function:prompt
        Remove-Item -Path Function:_OLD_VIRTUAL_PROMPT
    }

    # The prior PYTHONHOME:
    if (Test-Path -Path Env:_OLD_VIRTUAL_PYTHONHOME) {
        Copy-Item -Path Env:_OLD_VIRTUAL_PYTHONHOME -Destination Env:PYTHONHOME
        Remove-Item -Path Env:_OLD_VIRTUAL_PYTHONHOME
    }

    # The prior PATH:
    if (Test-Path -Path Env:_OLD_VIRTUAL_PATH) {
        Copy-Item -Path Env:_OLD_VIRTUAL_PATH -Destination Env:PATH
```

```
    Remove-Item -Path Env:_OLD_VIRTUAL_PATH
    }

    # Just remove the VIRTUAL_ENV altogether:
    if (Test-Path -Path Env:VIRTUAL_ENV) {
        Remove-Item -Path env:VIRTUAL_ENV
    }

    # Just remove VIRTUAL_ENV_PROMPT altogether.
    if (Test-Path -Path Env:VIRTUAL_ENV_PROMPT) {
        Remove-Item -Path env:VIRTUAL_ENV_PROMPT
    }

    # Just remove the _PYTHON_VENV_PROMPT_PREFIX altogether:
    if (Get-Variable -Name "_PYTHON_VENV_PROMPT_PREFIX" -ErrorAction SilentlyContinue)
{
        Remove-Variable -Name _PYTHON_VENV_PROMPT_PREFIX -Scope Global -Force
    }

    # Leave deactivate function in the global namespace if requested:
    if (-not $NonDestructive) {
        Remove-Item -Path function:deactivate
    }
}

<#
.Description
Get-PyVenvConfig parses the values from the pyvenv.cfg file located in the
given folder, and returns them in a map.

For each line in the pyvenv.cfg file, if that line can be parsed into exactly
two strings separated by `=` (with any amount of whitespace surrounding the =)
then it is considered a `key = value` line. The left hand string is the key,
the right hand is the value.
```

```
If the value starts with a `'` or a `"` then the first and last character is
stripped from the value before being captured.

.Parameter ConfigDir
Path to the directory that contains the `pyvenv.cfg` file.
#>
function Get-PyVenvConfig(
    [String]
    $ConfigDir
) {
    Write-Verbose "Given ConfigDir=$ConfigDir, obtain values in pyvenv.cfg"

    # Ensure the file exists, and issue a warning if it doesn't (but still allow the function to continue).
    $pyvenvConfigPath = Join-Path -Resolve -Path $ConfigDir -ChildPath 'pyvenv.cfg' -ErrorAction
Continue

    # An empty map will be returned if no config file is found.
    $pyvenvConfig = @{ }

    if ($pyvenvConfigPath) {

        Write-Verbose "File exists, parse `key = value` lines"
        $pyvenvConfigContent = Get-Content -Path $pyvenvConfigPath

        $pyvenvConfigContent | ForEach-Object {
            $keyval = $PSItem -split "\s*=\s*", 2
            if ($keyval[0] -and $keyval[1]) {
                $val = $keyval[1]

                # Remove extraneous quotations around a string value.
                if ("'""".Contains($val.Substring(0, 1))) {
                    $val = $val.Substring(1, $val.Length - 2)
                }
```

```
            $pyvenvConfig[$keyval[0]] = $val
            Write-Verbose "Adding Key: '$($keyval[0])'='$val'"
        }
    }
  }
  return $pyvenvConfig
}

<# Begin Activate script --------------------------------- #>

# Determine the containing directory of this script
$VenvExecPath = Split-Path -Parent $MyInvocation.MyCommand.Definition
$VenvExecDir = Get-Item -Path $VenvExecPath

Write-Verbose "Activation script is located in path: '$VenvExecPath'"
Write-Verbose "VenvExecDir Fullname: '$($VenvExecDir.FullName)"
Write-Verbose "VenvExecDir Name: '$($VenvExecDir.Name)"

# Set values required in priority: CmdLine, ConfigFile, Default
# First, get the location of the virtual environment, it might not be
# VenvExecDir if specified on the command line.
if ($VenvDir) {
  Write-Verbose "VenvDir given as parameter, using '$VenvDir' to determine values"
}
else {
  Write-Verbose "VenvDir not given as a parameter, using parent directory name as VenvDir."
  $VenvDir = $VenvExecDir.Parent.FullName.TrimEnd("\\/")
  Write-Verbose "VenvDir=$VenvDir"
}

# Next, read the `pyvenv.cfg` file to determine any required value such
# as `prompt`.
$pyvenvCfg = Get-PyVenvConfig -ConfigDir $VenvDir
```

```
# Next, set the prompt from the command line, or the config file, or
# just use the name of the virtual environment folder.
if ($Prompt) {
    Write-Verbose "Prompt specified as argument, using '$Prompt'"
}
else {
    Write-Verbose "Prompt not specified as argument to script, checking pyvenv.cfg value"
    if ($pyvenvCfg -and $pyvenvCfg['prompt']) {
        Write-Verbose "  Setting based on value in pyvenv.cfg='$($pyvenvCfg['prompt'])'"
        $Prompt = $pyvenvCfg['prompt'];
    }
    else {
        Write-Verbose "  Setting prompt based on parent's directory's name. (Is the directory name passed to venv module when creating the virtual environment)"
        Write-Verbose "  Got leaf-name of $VenvDir='$(Split-Path -Path $venvDir -Leaf)'"
        $Prompt = Split-Path -Path $venvDir -Leaf
    }
}

Write-Verbose "Prompt = '$Prompt'"
Write-Verbose "VenvDir='$VenvDir'"

# Deactivate any currently active virtual environment, but leave the
# deactivate function in place.
deactivate -nondestructive

# Now set the environment variable VIRTUAL_ENV, used by many tools to determine
# that there is an activated venv.
$env:VIRTUAL_ENV = $VenvDir

$env:VIRTUAL_ENV_PROMPT = $Prompt

if (-not $Env:VIRTUAL_ENV_DISABLE_PROMPT) {
```

```
    Write-Verbose "Setting prompt to '$Prompt'"

    # Set the prompt to include the env name
    # Make sure _OLD_VIRTUAL_PROMPT is global
    function global:_OLD_VIRTUAL_PROMPT { "" }
    Copy-Item -Path function:prompt -Destination function:_OLD_VIRTUAL_PROMPT
      New-Variable -Name _PYTHON_VENV_PROMPT_PREFIX -Description "Python virtual
environment prompt prefix" -Scope Global -Option ReadOnly -Visibility Public -Value $Prompt

    function global:prompt {
      Write-Host -NoNewline -ForegroundColor Green "($_PYTHON_VENV_PROMPT_PREFIX)
"
      _OLD_VIRTUAL_PROMPT
    }
}

# Clear PYTHONHOME
if (Test-Path -Path Env:PYTHONHOME) {
    Copy-Item -Path Env:PYTHONHOME -Destination Env:_OLD_VIRTUAL_PYTHONHOME
    Remove-Item -Path Env:PYTHONHOME
}

# Add the venv to the PATH
Copy-Item -Path Env:PATH -Destination Env:_OLD_VIRTUAL_PATH
$Env:PATH = "$VenvExecDir$([System.IO.Path]::PathSeparator)$Env:PATH"
```

**deactivate.bat :**

```
@echo off

if defined _OLD_VIRTUAL_PROMPT (
    set "PROMPT=%_OLD_VIRTUAL_PROMPT%"
)
```

```
set _OLD_VIRTUAL_PROMPT=

if defined _OLD_VIRTUAL_PYTHONHOME (
    set "PYTHONHOME=%_OLD_VIRTUAL_PYTHONHOME%"
    set _OLD_VIRTUAL_PYTHONHOME=
)

if defined _OLD_VIRTUAL_PATH (
    set "PATH=%_OLD_VIRTUAL_PATH%"
)

set _OLD_VIRTUAL_PATH=

set VIRTUAL_ENV=
set VIRTUAL_ENV_PROMPT=

:END
```

**pyvenv.cfg :**

```
home = C:\Users\shash\AppData\Local\Programs\Python\Python313
include-system-site-packages = false
version = 3.13.7
executable = C:\Users\shash\AppData\Local\Programs\Python\Python313\python.exe
command = C:\Users\shash\AppData\Local\Programs\Python\Python313\python.exe -m venv C:\Users\shash\Desktop\Automated Road Damage Detection Using UAV Images and Deep Learning Techniques\App\venv
```

**app.py :**

```
from flask import Flask, render_template, request
import torch
import torchvision.transforms as transforms
import torchvision.models as models
from PIL import Image
import os
import torch.nn as nn

app = Flask(__name__)
UPLOAD_FOLDER = "static/uploads"
os.makedirs(UPLOAD_FOLDER, exist_ok=True)
app.config['UPLOAD_FOLDER'] = UPLOAD_FOLDER

# Define class names
classes = ["good", "poor", "satisfactory", "very_poor"]

def load_model(model_path, device):
    model = models.resnet101(weights=None)
    num_features = model.fc.in_features
    model.fc = nn.Sequential(
        nn.Linear(num_features, 512),
        nn.ReLU(),
        nn.Dropout(0.5),
        nn.Linear(512, len(classes))
    )
    state_dict = torch.load(model_path, map_location=device)
    if "module." in list(state_dict.keys())[0]:
        from collections import OrderedDict
        new_state_dict = OrderedDict()
        for k, v in state_dict.items():
            name = k.replace("module.", "")
            new_state_dict[name] = v
```

```
        state_dict = new_state_dict
    model.load_state_dict(state_dict)
    model.to(device)
    model.eval()
    return model

# Define Image Transformation
transform = transforms.Compose([
    transforms.Resize((224, 224)),
    transforms.ToTensor(),
    transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])
])

def predict_image(model, image_path, device):
    image = Image.open(image_path).convert("RGB")
    image = transform(image).unsqueeze(0).to(device)
    with torch.no_grad():
        output = model(image)
        _, predicted_class = torch.max(output, 1)
        predicted_label = classes[predicted_class.item()]
    return predicted_label

# Load model
model_path = "model/best_resnet101_model.pth"
device = torch.device("cuda" if torch.cuda.is_available() else "cpu")
if os.path.exists(model_path):
    model = load_model(model_path, device)
else:
    model = None

@app.route('/')
def index():
    return render_template('index.html')
```

```
@app.route('/predict', methods=['POST'])
def predict():
    if 'file' not in request.files:
        return render_template('index.html', error="No file uploaded")
    file = request.files['file']
    if file.filename == '':
        return render_template('index.html', error="No selected file")
    file_path = os.path.join(app.config['UPLOAD_FOLDER'], file.filename)
    file.save(file_path)

    if model:
        prediction = predict_image(model, file_path, device)
        return render_template('index.html', image_path=file_path, prediction=prediction)
    else:
        return render_template('index.html', error="Model file not found")

if __name__ == '__main__':
    app.run(debug=True)
```

## 6.2. Implementation

**Front-End Implementation**

The front-end of the automated road damage detection system is designed to provide a simple, responsive, and user-friendly interface for interacting with the system. It is developed using HTML, CSS, and JavaScript, ensuring compatibility across different devices and browsers. The interface allows users to upload road images captured manually or through UAV systems for analysis.

The user interface includes modules such as image upload, preview display, and result visualization. When a user uploads an image, it is immediately displayed on the screen for confirmation. JavaScript is used to handle client-side validation, ensuring that only valid image formats are accepted before sending data to the backend.

The result display module presents the predicted road condition category (Good, Satisfactory, Poor, or Very Poor) along with confidence scores. The design focuses on clarity and simplicity so that users can easily interpret the results without technical knowledge. Proper styling and structured layouts enhance usability and overall user experience.

**Back-End Implementation (Flask)**

The backend of the system is implemented using Flask, a lightweight and flexible Python web framework. Flask is chosen for its simplicity and ability to integrate seamlessly with machine learning models. The backend handles image processing, model inference, and communication between the front-end and the deep learning model.

The system follows a modular structure where different functionalities such as image handling, preprocessing, and prediction are organized into separate components. The uploaded image is received through Flask routes and passed to the preprocessing pipeline, where operations like resizing, normalization, and formatting are applied.

After preprocessing, the image is fed into the trained ResNet101 model, which performs classification based on learned features. The prediction result is then sent back to the front-end in a structured format (JSON response), where it is displayed to the user. Flask ensures smooth handling of requests and efficient communication between components.

**Model Integration and Processing Workflow**

The deep learning model (ResNet101) is integrated into the Flask backend as a core processing component. When an image is uploaded, it passes through a preprocessing pipeline that includes resizing to the required input size, normalization of pixel values, and conversion into tensor format suitable for the model.

The processed image is then passed to the ResNet101 model, which extracts features through multiple convolutional and residual layers. The model outputs a probability distribution across the four road condition classes. The class with the highest probability is selected as the final prediction.

The prediction result, along with confidence scores, is returned to the front-end in JSON format. This

structured communication ensures efficient data transfer and quick response times. The workflow is optimized to support near real-time predictions, making the system suitable for practical deployment.

**Deployment and Reliability**

The system is designed to be deployed on a local server or cloud environment using Flask. It supports scalable deployment, allowing integration with cloud platforms for handling large datasets and real-time processing. The application ensures reliability through proper error handling, input validation, and consistent response mechanisms.

Static files such as HTML, CSS, and JavaScript are served efficiently, while API endpoints are tested to ensure stable performance under different usage conditions. The system is capable of handling multiple requests and maintaining consistent performance, making it reliable for real-world applications.

# 7. TESTING

## 7.1 Types of System Testing

System testing was performed to ensure that our automated road damage detection system works correctly, reliably, and efficiently under real-world conditions. The testing process focused on validating all modules of the system, including the front-end interface, Flask backend, preprocessing pipeline, and the ResNet101 model. The objective was to verify system behavior, identify defects, and ensure that all functional requirements of the project were met.

### 7.1.1 Unit Testing

In our project, unit testing was carried out to validate individual components independently.

- The image upload module was tested to ensure that different image formats (JPG, PNG) were accepted and processed correctly.
- The preprocessing module was tested to verify proper resizing, normalization, and transformation of input images before feeding them into the model.
- The model prediction module was tested using sample images to confirm that the ResNet101 model correctly identifies road damage.

By performing unit testing, we were able to detect errors at an early stage, such as incorrect input handling and preprocessing mismatches. This ensured that each component functioned correctly before integration.

### 7.1.2 Integration Testing

Integration testing was conducted to verify that all modules of the system worked together seamlessly.

- The interaction between the front-end (HTML, CSS, JavaScript) and the Flask backend was tested to ensure smooth data transfer.
- The connection between the preprocessing pipeline and the ResNet101 model was validated to ensure correct input-output flow.
- The system was tested end-to-end by uploading images and verifying whether the prediction results were correctly displayed on the interface.

This testing helped ensure proper communication between components and eliminated issues like data mismatch and API response errors.

### 7.1.3 System Testing

System testing was performed on the complete system to evaluate its performance in real-world scenarios.

- The entire workflow, from image upload to final prediction output, was tested.
- The system was tested with multiple real-time road images to verify accuracy and consistency.
- Performance testing ensured that the system responded within acceptable time limits.

The results showed that the system was able to detect road damage accurately and provide reliable outputs, confirming that the project meets its intended objectives.

## 7.2 Testing Strategies

A structured testing strategy was followed to ensure the reliability, accuracy, and stability of the automated road damage detection system. The testing process was carried out in a systematic manner, starting from individual component validation and progressing toward complete system verification. This step-by-step approach helped in identifying issues early and ensured smooth integration of all modules.

Both manual and functional validation approaches were used during testing. Different test scenarios were designed to simulate real-world usage, such as uploading various types of road images under different conditions. This ensured that the system could handle diverse inputs and produce consistent outputs. Special attention was given to validating the interaction between the front-end and Flask backend, as well as the performance of the ResNet101 model during prediction. The system was tested with different image qualities, sizes, and road conditions to evaluate its robustness and generalization capability.

The testing strategy also focused on ensuring smooth user interaction, proper error handling, and consistent response generation. Input validation mechanisms were verified to prevent invalid or unsupported data from affecting system performance. This ensured that the system remains stable and user-friendly under all operating conditions.

Overall, the testing strategy ensured that the system performs efficiently, delivers accurate predictions, and maintains reliability across different scenarios, making it suitable for real-world deployment.

### 7.2.1 Test Strategy and Approach

A structured testing strategy was followed to ensure the system performs accurately and reliably. Testing was conducted in a phased manner, starting from unit testing, followed by integration testing, and finally system testing. This approach helped in identifying issues at early stages and ensured smooth interaction between all modules.

Both manual testing and scenario-based validation were used to simulate real-world conditions. The system was tested using different road images with variations in lighting, texture, and damage levels to ensure robustness. Special focus was given to verifying the interaction between the front-end (HTML, CSS, JavaScript) and the Flask backend, along with the performance of the ResNet101 model.

### 7.2.2 Test Objectives

The main objectives of testing were to ensure that all components of the system function correctly and meet the required performance standards. The testing process aimed to verify that the system produces accurate predictions, handles user inputs efficiently, and maintains stability during operation.

It also focused on validating proper data flow from image upload to final classification, ensuring that preprocessing and model prediction work seamlessly. Additionally, the objective was to confirm that the system behaves consistently under different input conditions.

### 7.2.3 Features Tested

The major features tested in the system include image upload functionality, preprocessing pipeline, model prediction using ResNet101, and result display on the user interface. Each feature was validated to ensure correct operation and expected output.

The interaction between front-end and backend was also tested to ensure smooth data transmission. Input validation mechanisms were checked to confirm that invalid or unsupported image formats are handled properly. Overall system usability and response time were also evaluated.

### 7.2.4 Integration Testing Strategy

Integration testing focused on verifying the interaction between different modules of the system. This includes communication between the front-end interface and Flask backend, as well as the connection

between preprocessing steps and the ResNet101 model.

The strategy ensured that data flows correctly through each stage without loss or corruption. It is also validated that the prediction results generated by the model are accurately returned and displayed to the user. This helped in confirming that all integrated components work together seamlessly.

### 7.2.5 Acceptance Criteria

The system is considered acceptable when it successfully processes input images and provides accurate road condition classification results. The application must run without errors and handle different types of inputs efficiently.

The user interface should be responsive and easy to use, and the system should maintain consistent performance during repeated usage. All modules must function correctly together, ensuring reliable end-to-end operation.

### 7.2.6 Overall Test Results

All planned test cases were executed successfully, and the system demonstrated stable performance across different scenarios. The model produced accurate predictions for various types of road damage images, confirming its effectiveness.

The integration between front-end, backend, and model components worked as expected, with no major issues observed. The system showed consistent behavior and reliable performance during testing.

## 7.3. Sample Test Cases

| S.No. | Test Case | Expected Result | Result |
|---|---|---|---|
| 1 | Upload Valid Road Image | Image is successfully uploaded and displayed | Pass |
| 2 | Upload Invalid File Format | System rejects file and shows error message | Pass |
| 3 | Image Preprocessing | Image is resized and normalized correctly | Pass |
| 4 | Model Prediction (Clear Damage Image) | Correct damage class is predicted | Pass |
| 5 | Model Prediction (Good Road Image) | Classified as Good/Satisfactory | Pass |
| 6 | Backend Processing (Flask API) | Image is processed and response returned | Pass |
| 7 | Extremely Blurred Image | System predicts with reduced confidence | Fail |
| 8 | Night-Time / Poor Lighting Image | Correct classification under low visibility | Fail |
| 9 | Mixed Damage Types in One Image | Detect dominant damage class | Partial Pass |
| 10 | Non-Road Image Input (e.g., buildings) | System rejects or misclassifies | Fail |

**Table no 7.3.** Test cases

**Test Case 1:**

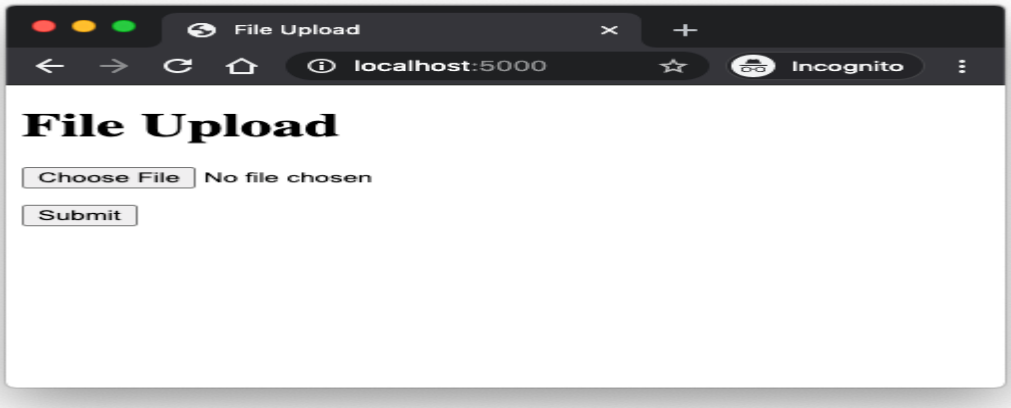


Fig 7.3.1 User Login

**Description:** The user uploads a valid road image through the interface. The system accepts the image and displays a preview before processing.

**Test Case 2:**

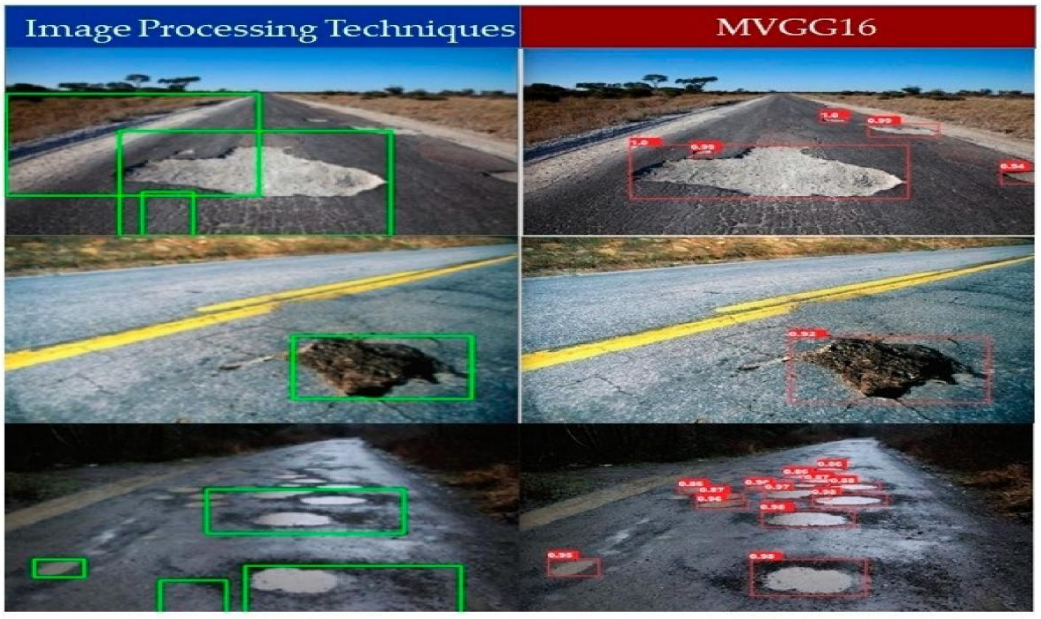


**Fig 7.3.2** Image Processing

**Description:** The system processes a damaged road image and predicts the correct class such as Poor or Very Poor with confidence score.

**Test Case 3:**

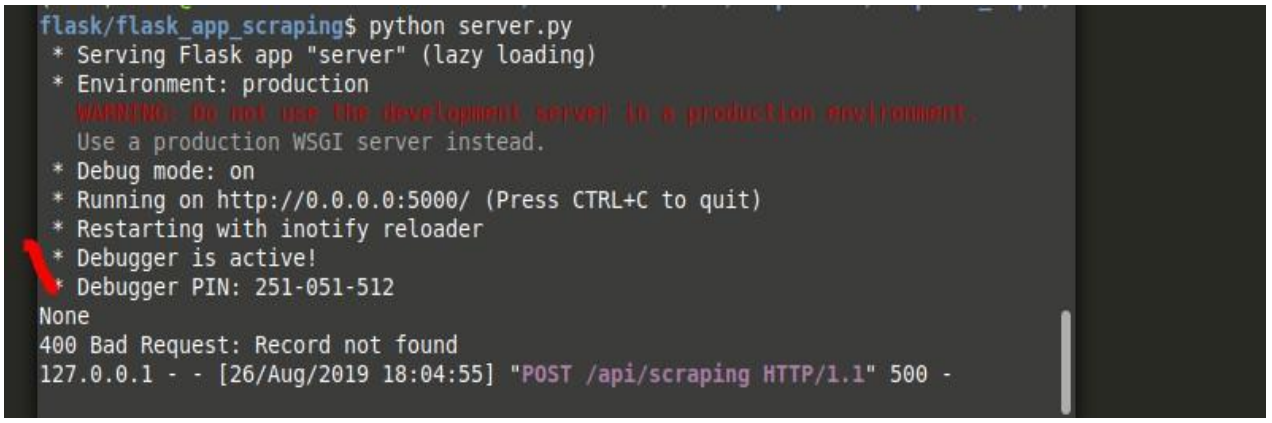


**Fig 7.3.3** Redirection to WebApp

**Description:** The Flask backend receives the image, processes it using the model, and returns the prediction result successfully.

# 8. OUTPUT SCREENS

## Evaluation Metrics:

Test Accuracy: 0.9515

| Attributes/Metrics | Precision | Recall | F1-score | Support |
|---|---|---|---|---|
| Good | 0.99 | 1.00 | 0.99 | 84 |
| Poor | 0.89 | 0.87 | 0.88 | 39 |
| Satisfactory | 0.90 | 0.92 | 0.91 | 51 |
| Very_poor | 1.00 | 0.97 | 0.98 | 32 |
| Accuracy | | | 0.95 | 206 |
| Macro_avg | 0.95 | 0.94 | 0.94 | 206 |
| Weighted_avg | 0.95 | 0.95 | 0.95 | 206 |

**Table no. 8.1:** Classification Report

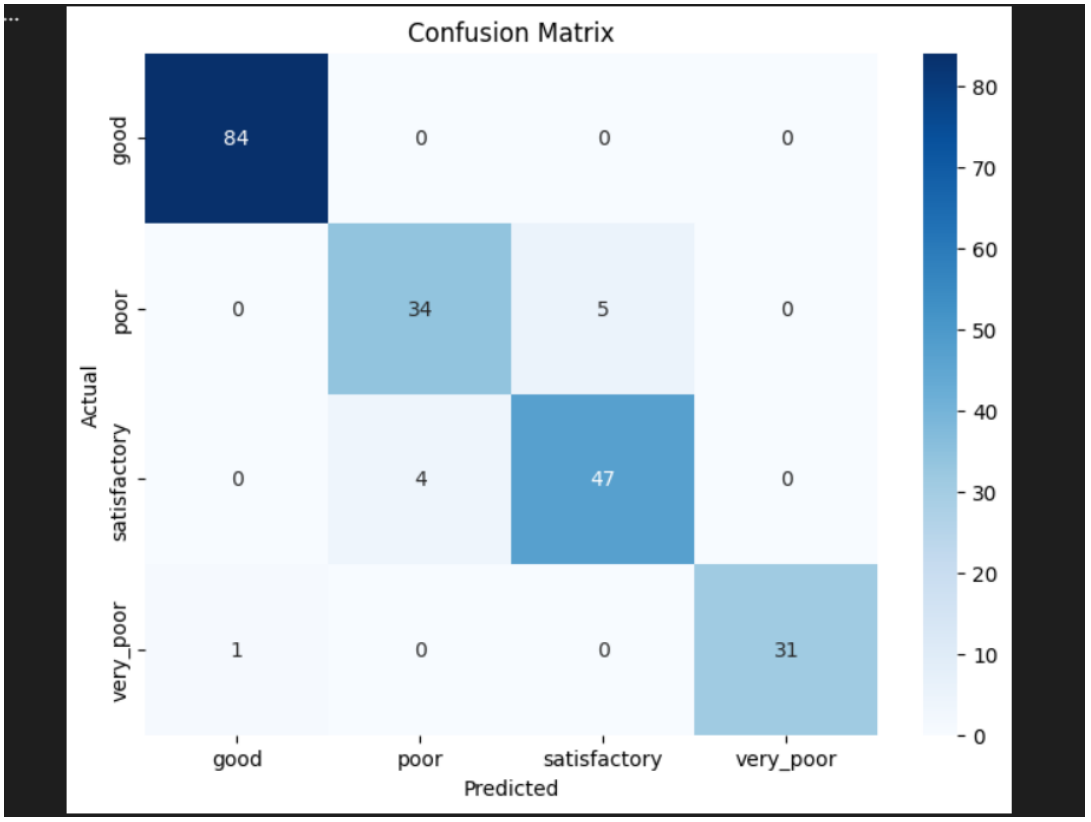


**Fig no. 8.2:** Confusion Matrix

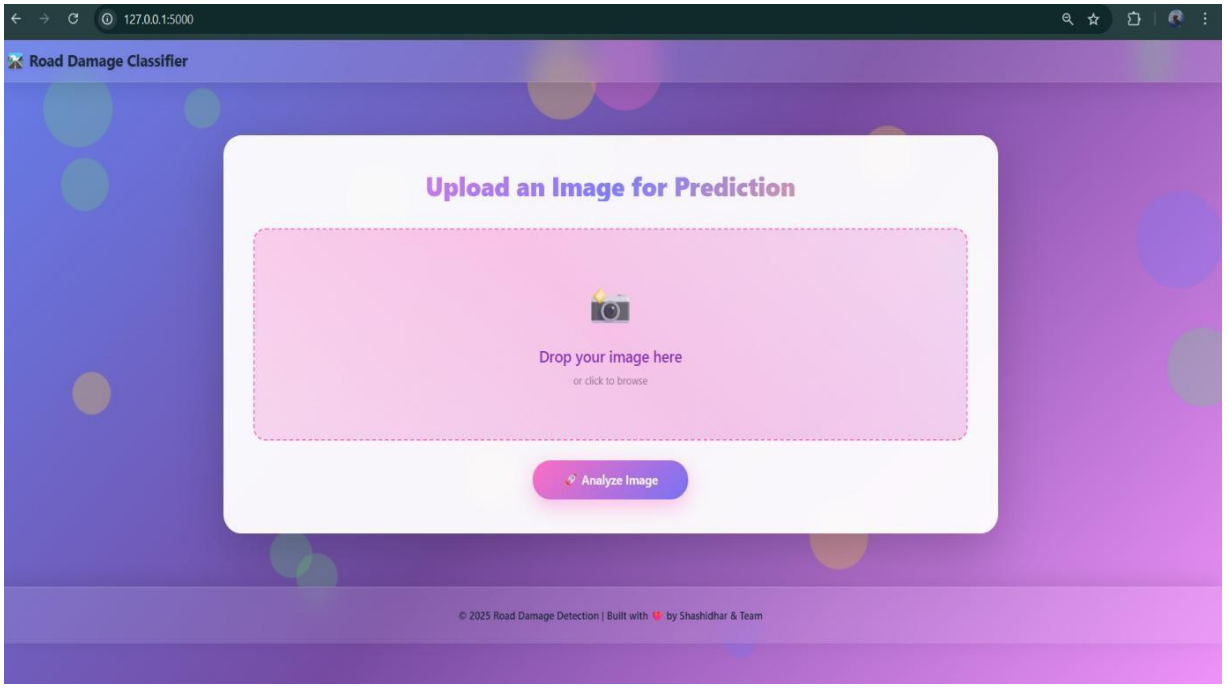


**Fig 8.3 Road Damage Detection Landing Page**

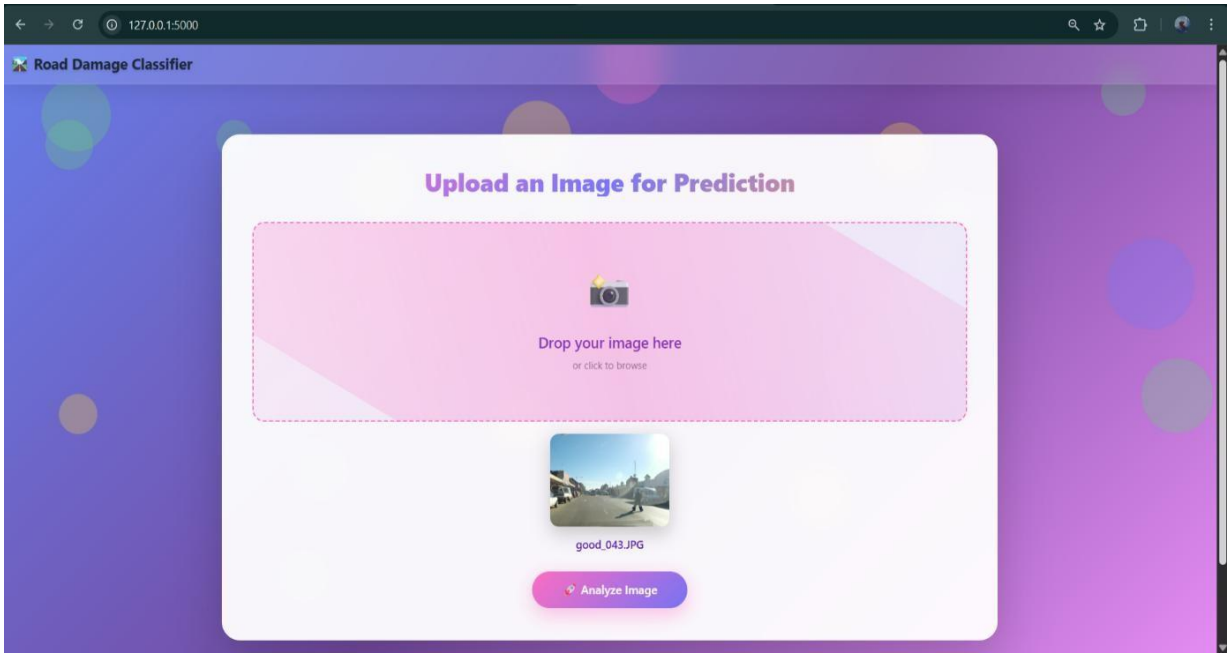


**Fig 8.4 Image-1 Analysis for RDD**

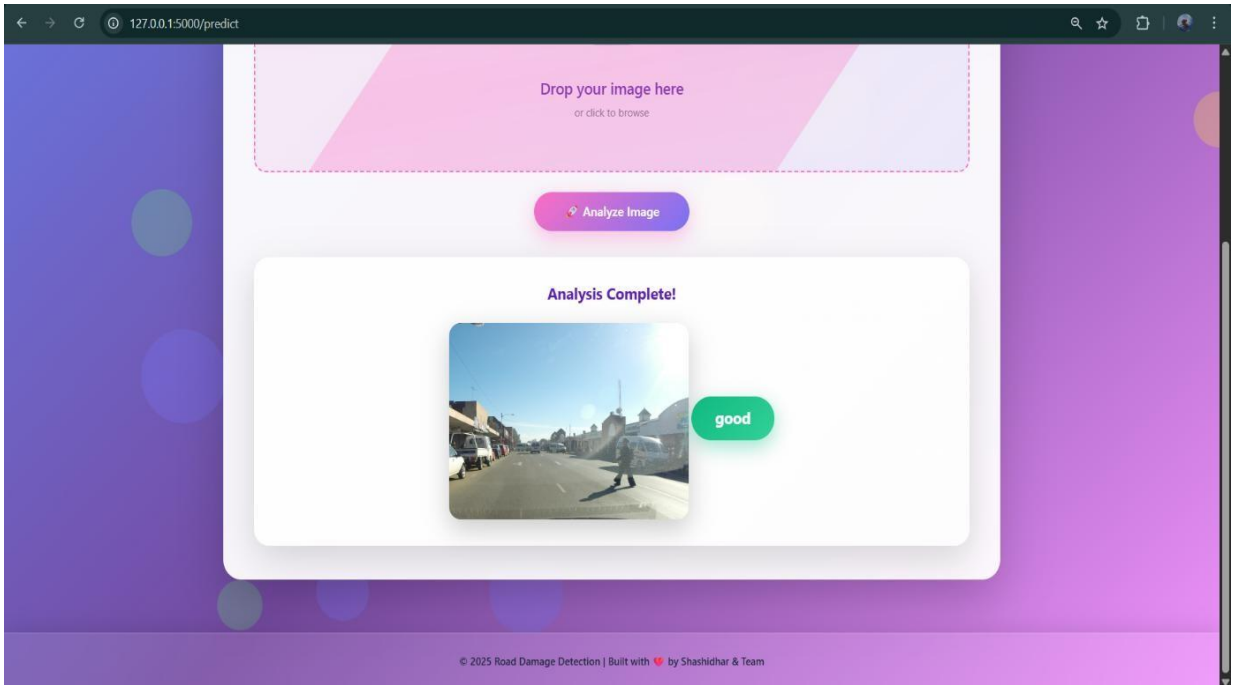


**Fig 8.5 Analysis result of Image-1**

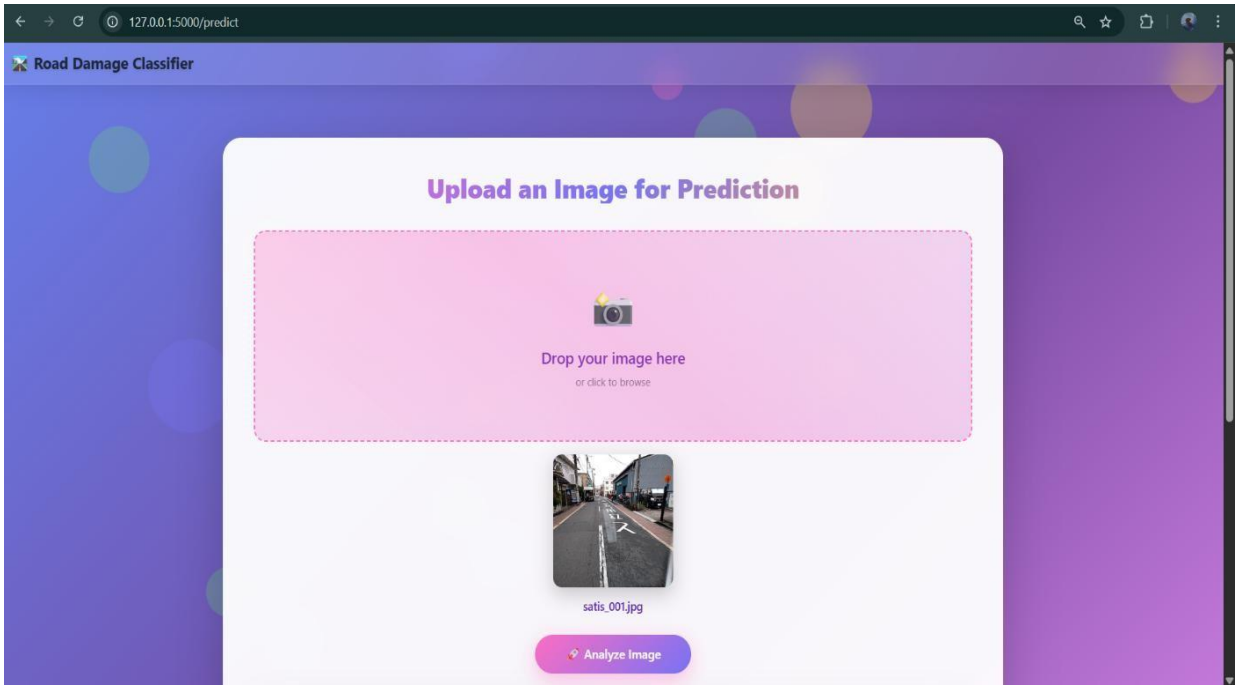


**Fig 8.6 Image-2 Analysis for RDD**

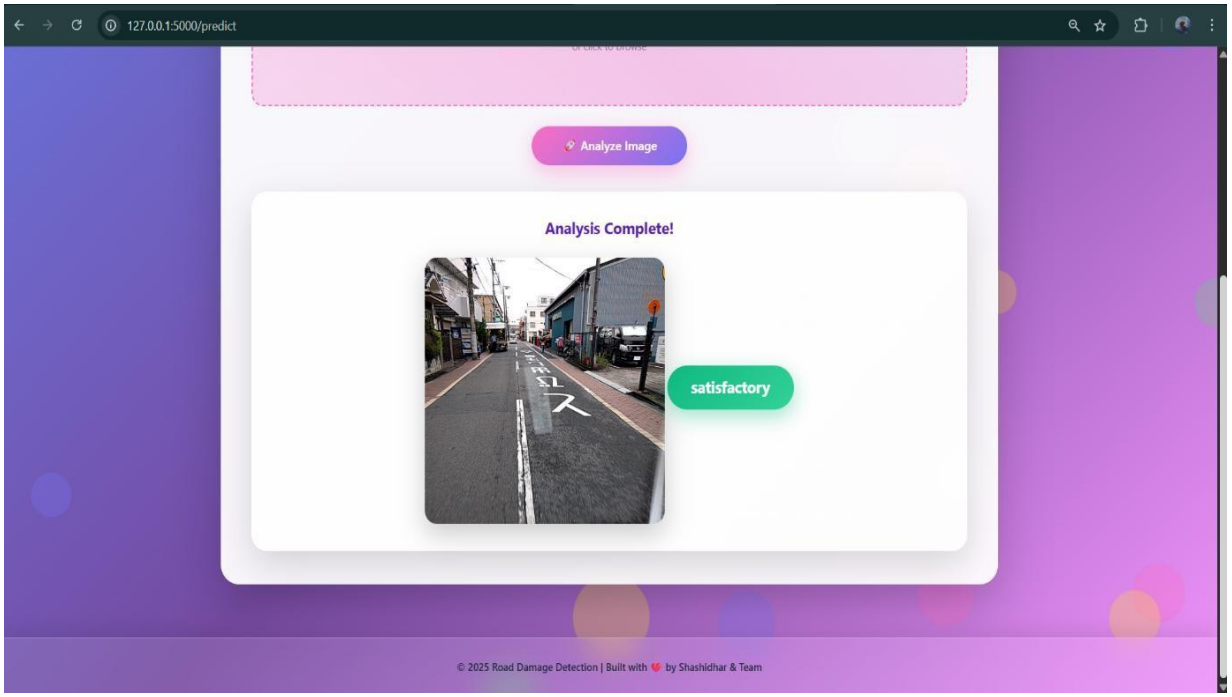


**Fig 8.7 Analysis result of Image-2**

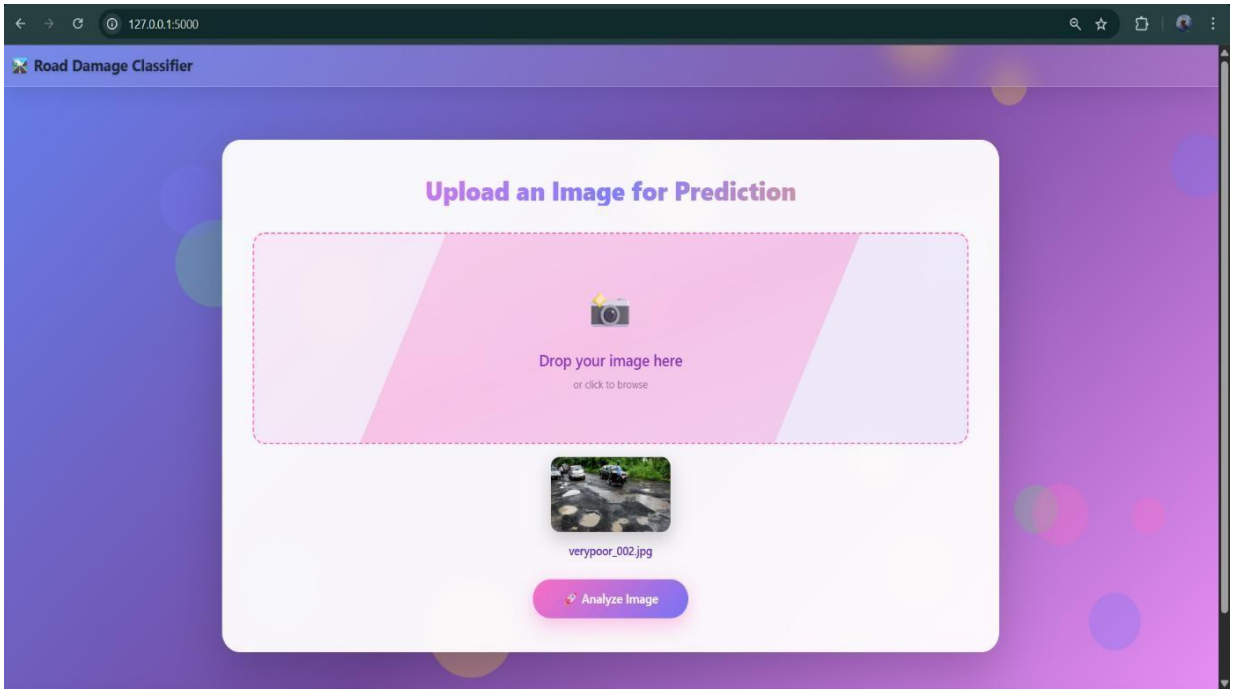


**Fig 8.8 Image-3 Analysis for RDD**

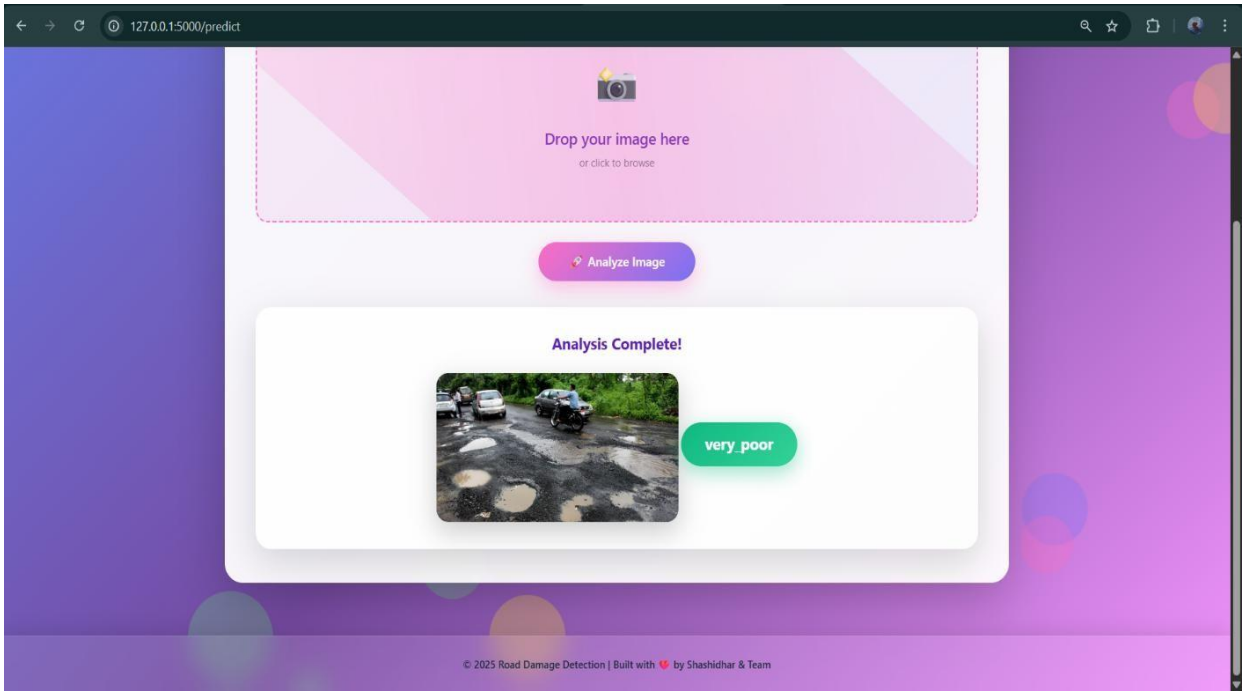


**Fig 8.9 Analysis result for Image-3**

# 9. CONCLUSION

This study aims to bridge the gap between deep learning and real-world infrastructure management by developing an automated road damage classification system. By leveraging CNN with ResNet101, transfer learning, and optimization techniques, we aim to create a system that is accurate, efficient, and deployable for road monitoring applications. The outcomes of this research will help government agencies, smart city planners, and transportation departments maintain road networks proactively and cost-effectively.

The proposed system successfully demonstrates how advanced computer vision methods can analyze diverse road images and classify damage levels with high precision. Through a carefully structured data preparation pipeline—consisting of cleaning, normalization, augmentation, and tensor conversion—the model receives high-quality inputs, enabling it to learn subtle and complex road-surface patterns. ResNet101's deep residual architecture excels at capturing fine cracks, texture variations, pothole edges, and severe surface degradation, leading to robust classification across Good, Satisfactory, Poor, and Very Poor categories.

Furthermore, the integration of transfer learning significantly reduces training time and enhances performance, especially when working with limited datasets. The modular system architecture allows seamless interaction between preprocessing modules, model inference, storage units, and result visualization. This ensures easy maintainability, scalability, and future extensions—such as incorporating segmentation models, drone-based live road inspection, or GIS-based road-mapping tools.

By automating what is traditionally a slow, manual, and error-prone inspection process, the system provides a reliable tool for real-time road-condition monitoring. It supports quicker decision-making, reduces human effort, and helps maintenance teams prioritize repairs based on accurate condition assessments. Overall, this research not only demonstrates the feasibility of deep learning for road-damage detection but also lays the groundwork for intelligent transportation systems that contribute to safer, smarter, and more sustainable infrastructure management

# 10. FUTURE ENHANCEMENTS

The current system establishes a strong conceptual and technical foundation for automated road-damage classification using deep learning; however, several enhancements can significantly extend its real-world applicability, accuracy, and scalability. Future improvements may focus on enriching the model's capability by incorporating multi-label and multi-damage classification, enabling the system to identify multiple defect types—such as cracks, potholes, rutting, and surface deformation—within a single image instead of assigning only one class. Another promising direction involves integrating advanced object detection frameworks like YOLOv8 or Faster R-CNN to localize the exact damaged regions. This would allow municipal authorities and smart-city systems to visualize and prioritize zones requiring urgent intervention, enhancing the precision of maintenance workflows.

Beyond classification and detection, deployment-centric advancements can transform the system into a fully operational smart-infrastructure tool. Lightweight model optimization using quantization, pruning, ONNX, and TensorRT can support real-time inference on mobile phones, drones, and edge devices such as Raspberry Pi or NVIDIA Jetson. Drone-based automated road scanning combined with IoT-enabled sensors can enable large-scale, real-time monitoring of highways and urban roads. Additionally, incorporating active learning or federated learning would allow the model to continuously adapt to new road textures, climates, and lighting conditions without requiring centralized datasets. Integrating the system with GIS dashboards, smart-city platforms, and automated maintenance pipelines will ultimately transform this research from a standalone classifier into a comprehensive, intelligent, and scalable road-infrastructure management solution.

# 11. REFERENCES


1. Chen, W., Liu, H., Zhang, Q., “Real-Time Road Damage Detection Using YOLOv8 for Intelligent Transportation Systems,” IEEE Access, vol. 13, pp. 4567–4578, 2025.

2. Patel, R., Kumari, S., “Adaptive Road Surface Damage Detection Using Vision Transformers and UAV Imagery,” IEEE Transactions on Intelligent Transportation Systems, vol. 26, no. 3, pp. 1124–1138, 2025.

3. Rasheed, S., Khan, M., Ali, F., “UAV-Based Road Damage Detection Using YOLO Deep Learning Approach,” International Journal of Innovative Research in Technology, vol. 11, no. 1, pp. 102–110, 2025.

4. Jiang, Y., Wang, X., Li, Z., “Improved YOLOv8 with Attention Mechanism for Road Damage Detection,” Springer Nature, vol. 14, no. 2, pp. 210–225, 2024.

5. Cheng, C., Zhao, L., Sun, Y., “Multiscale Feature Fusion-Based Road Damage Detection Using Improved YOLOv5,” Journal of Remote Sensing, vol. 18, no. 4, pp. 334–348, 2024.

6. Silva, L. A., Pereira, M., Costa, R., “Automated Road Damage Detection Using UAV Images and Deep Learning Techniques,” IEEE Conference on Smart Infrastructure, pp. 89–96, 2023.

7. Zhang, Y., Liu, J., Chen, K., “Road Damage Detection Using YOLOv3 with Multi-Level Attention Mechanism,” Expert Systems with Applications, vol. 202, pp. 117–129, 2022.

8. Arya, D., Maeda, H., Ghosh, S., “RDD2022: A Multi-Country Road Damage Detection Dataset for Deep Learning,” arXiv preprint arXiv:2209.08538, 2022.

9. Arya, D., Maeda, H., Ghosh, S., “Deep Learning-Based Road Damage Detection System for Multiple Countries,” IEEE International Conference on Big Data, pp. 1–8, 2021.

10. Doshi, K., Yilmaz, Y., “Road Damage Detection Using Deep Ensemble Learning and YOLOv4,” arXiv preprint arXiv:2011.00728, 2020.

11. Maeda, H., Sekimoto, Y., Seto, T., Kashiyama, T., Omata, H., “Road Damage Detection and Classification Using Deep Neural Networks with Smartphone Images,” Computer-Aided Civil and Infrastructure Engineering, vol. 35, no. 6, pp. 537–553, 2020.

12. Li, S., Zhao, X., Zhou, G., “Automatic Pixel-Level Multiple Damage Detection of Concrete Structure Using Fully Convolutional Network,” Computer-Aided Civil and Infrastructure Engineering, 2019.

13. Bang, S., Park, S., Kim, H., Kim, H., “Encoder–Decoder Network for Pixel-Level Road Crack Detection,” IEEE Access, vol. 7, pp. 102244–102254, 2019.

14. Maeda, H., Sekimoto, Y., Seto, T., Kashiyama, T., Omata, H., “Road Damage Detection Using Deep Neural Networks with Images Captured Through a Smartphone,” arXiv preprint arXiv:1801.09454, 2018.

15. Fan, R., Ai, H., “Road Surface Crack Detection Using Deep Convolutional Neural Network,” IEEE International Conference on Image Processing, 2018.

16. Jenkins, M. D., Carr, T. A., Iglesias, M. I., Buggy, T., Morison, G., “Automatic Pavement Crack Detection Using Deep Learning-Based Convolutional Neural Network,” IEEE International Conference on Image Processing, 2018.

17. Zhang, D., Li, Q., Chen, Y., Cao, M., He, L., Zhang, B., “An Efficient and Reliable Crack Detection Method Using Convolutional Neural Networks,” Neurocomputing, vol. 272, pp. 886–895, 2018.

18. Eisenbach, M., Stricker, R., Seichter, D., et al., “How to Get Pavement Distress Detection Ready for Deep Learning? A Systematic Approach,” International Joint Conference on Neural Networks (IJCNN), 2017.

19. Chun, P., Yamane, T., Izumi, S., Sakai, K., “Automatic Detection Method of Cracks from Concrete Surface Imagery Using Deep Learning,” IEEE Transactions on Automation Science and Engineering, 2017.

20. Cha, Y. J., Choi, W., Büyüköztürk, O., “Deep Learning-Based Crack Damage Detection Using Convolutional Neural Networks,” Computer-Aided Civil and Infrastructure Engineering, vol. 32, no. 5, pp. 361–378, 2017.

21. Shi, Y., Cui, L., Qi, Z., Meng, F., Chen, Z., “Automatic Road Crack Detection Using Random Structured Forests,” IEEE Transactions on Intelligent Transportation Systems, vol. 17, no. 12, pp. 3434–3445, 2016.

22. Zhang, L., Yang, F., Zhang, Y. D., Zhu, Y. J., “Road Crack Detection Using Deep Convolutional Neural Network,” IEEE International Conference on Image Processing (ICIP), 2016.

23. Amhaz, R., Chambon, S., Idier, J., Baltazart, V., “Automatic Crack Detection on Two-Dimensional Pavement Images: An Algorithm Based on Minimal Path Selection,” IEEE Transactions on Intelligent Transportation Systems, vol. 17, no. 10, pp. 2718–2729, 2016.

24. Oliveira, H., Correia, P. L., “Automatic Road Crack Detection and Characterization,” IEEE Transactions on Intelligent Transportation Systems, vol. 14, no. 1, pp. 155–168, 2013.

25. Zou, Q., Cao, Y., Li, Q., Mao, Q., Wang, S., “CrackTree: Automatic Crack Detection from Pavement Images,” Pattern Recognition Letters, vol. 33, no. 3, pp. 227–238, 2012.