Comparative Analysis of Deep Learning Models for Road Damage Classification

By

Md.Mahmudul Islam

ID: 202-15-14442

FINAL YEAR DESIGN PROJECT REPORT

This Report Presented in Partial Fulfillment of the Requirements forthe Degree of Bachelor of Science in Computer Science and Engineering

Supervised by

Hasnur Jahan

Lecturer

Department of Computer Science and EngineeringDaffodil International University

Co-Supervisor

Dr. Md. Zahid Hasan

Associate Professor

Department of Computer Science and EngineeringDaffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY

Dhaka, Bangladesh

September 17, 2025

APPROVAL

This Project titled **"Comparative Analysis of Deep Learning Models for Road Damage Classification,"** submitted by **Md. Mahmudul Islam** and to the Department of Computer Science and Engineering, Daffodil International University, has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of B.Sc. in Computer Science andEngineering and approved as to its style and contents. The presentation has been held on **26-01-2024.**

BOARD OF EXAMINERS

Dr. Arif Mahmud (AM) — **Chairman**
Associate Professor & Associate Head
Chairman, Department of CSE, FSITDaffodil
International University

Dr. Fizar Ahmed (FZA) — **Internal Examiner**
Associate Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Md. Mushfiqur Rahman (MUR) — **Internal Examiner**
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Dr. Md. Manowarul Islam — **External Examiner**
Professor
Department of Computer Science and Engineering
Jagannath University

DECLARATION

We hereby declare that this project has been done by us under the supervision of **Hasnur Jahan, Lecturer**, Department of Computer Science and Engineering, Daffodil International University. We also declare that neither this project nor any part of this project has been submitted elsewhere for the award of any degree or diploma.

Supervised by:

17.9.25

Hasnur Jahan
Lecturer
Department of Computer Science and Engineering Daffodil International University

Submitted by:

Mahmudul 17.09.25

Md. Mahmudul Islam
Student ID: 202-15-14442
Department of Computer Science and Engineering Daffodil International University

ACKNOWLEDGEMENTS

This work would not have been possible without the support and contributions of many individuals over the past two semesters. We are deeply grateful to everyone who has as- sisted us in one way or another.

First, we express our heartfelt thanks and gratefulness to the almighty for His divine blessing making it possible for us to complete the Final Year Design Project (FYDP) successfully.

We are grateful and wish our profound indebtedness to Hasnur Jahan, Lecturer, Department of Computer Science and Engineering, Daffodil International University, Dhaka, Bangladesh. Deep knowledge and keen interest of our supervisor in the field of Deep Learning to carry out this project. Her endless patience, scholarly guidance, continual encouragement, constant and energetic supervision, constructive criticism, valuable advice, reading many inferior drafts, and correcting themat all stages have made it possible to complete this project.

We would like to express our heartfelt gratitude to the Dr Sheak Rashed Haidar Noori, Professor & Head, Department of Computer Science and Engineering, for his kind help in finishing our project and also to other faculty members and the staff of the Department of Computer Science and Engineering, Daffodil International University.

We would like to thank our entire course-mates at Daffodil International University, who took part in this discussion while completing the coursework.

Finally, we must acknowledge with due respect the constant support and patience of our parents.

ABSTRACT

Road damage, particularly cracks and potholes, poses significant risks to transportation safety and infrastructure sustainability. Traditional manual inspection methods are time-consuming, costly, and often unreliable, creating the need for automated solutions. This study proposes a deep learning-based approach for road damage classification using a dataset of 1,026 Bangladeshi road images, consisting of 490 crack and 536 pothole samples. Six pretrained convolutional neural network (CNN) models—MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19—were implemented and evaluated. All models were trained for 15 epochs with a learning rate of 0.001. Among them, MobileNetV2 achieved the highest performance with an accuracy of 90%, precision of 0.91, recall of 0.90, and F1-score of 0.90, demonstrating its suitability for resource-efficient real-world deployment. The proposed system offers an effective solution for timely road maintenance, contributing to improved road safety, reduced vehicle damage, and sustainable infrastructure management.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction

Road networks are one of the most important components of modern infrastructure, directly contributing to socio-economic development, public safety, and efficient transportation. In countries like Bangladesh, where roads serve as the primary mode of transportation for goods and people, maintaining road quality is a critical task. However, road damage in the form of cracks and potholes is a widespread issue, leading to accidents, vehicle damage, traffic congestion, and significant repair costs [1]. The ability to detect and classify such road damage efficiently is therefore essential for timely maintenance and improved road safety.

Conventional road monitoring practices in Bangladesh mostly rely on manual inspection carried out by road authorities. While effective on a small scale, manual inspections are inefficient, time-consuming, labor-intensive, and subject to human error when dealing with large-scale road networks [2]. Traditional computer vision and image processing techniques have also been attempted, but these methods are highly sensitive to environmental conditions such as lighting, weather, and road surface variations, resulting in limited applicability in real-world scenarios.

With the advancement of artificial intelligence, particularly deep learning, automated road damage detection has emerged as a promising solution. Convolutional Neural Networks (CNNs) have shown state-of-the-art performance in a wide range of image classification tasks, offering the ability to learn robust features directly from images [3]. Furthermore, transfer learning with pretrained models has made it possible to apply powerful architectures, such as MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19, to relatively small datasets, drastically reducing training time and improving accuracy [4].

Despite global progress in road damage detection research, limited work has been carried out using localized Bangladeshi datasets. Given that road textures, materials, and damage patterns in Bangladesh are unique, it is important to validate existing deep learning models under these conditions

[5]. This thesis addresses this gap by constructing a Bangladeshi road damage dataset and evaluating multiple pretrained CNN architectures for crack and pothole classification.

1.2 Motivation

The motivation behind this research arises from the growing challenges posed by road damage in Bangladesh. Cracks and potholes not only disrupt transportation but also cause severe economic and safety implications. According to reports, poor road conditions contribute significantly to road accidents in developing regions. Manual inspection methods are insufficient to cover vast road networks effectively, while existing automated systems are often designed for road conditions in developed countries and do not generalize well to the Bangladeshi context.

There is, therefore, a pressing need for an efficient, automated, and localized road damage detection system that can assist authorities in identifying and prioritizing road maintenance tasks. By leveraging deep learning and pretrained CNN architectures, it becomes possible to design a system that is both accurate and computationally efficient, with the potential to be deployed in real-world environments, including mobile and embedded devices.

1.3 Objectives

The main objectives of this thesis are:

1. To construct and utilize a dataset of Bangladeshi road damage images containing two classes: cracks and potholes.

2. To apply transfer learning using pretrained CNN models—MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19—for road damage classification.

3. To evaluate and compare the performance of these models under identical training conditions (15 epochs, learning rate 0.001).

4. To identify the most suitable model for road damage detection in terms of accuracy, computational efficiency, and applicability in real-world scenarios.

5. To establish a baseline study for road damage detection in Bangladesh, paving the way for future research and development.

1.4 Methodology

The methodology of this research consists of the following steps:

1. Dataset Preparation: A primary dataset was developed consisting of 1,026 images of Bangladeshi road damage, divided into two categories: cracks (490 images) and potholes (536 images). Images were preprocessed through resizing, normalization, and data augmentation to improve generalization.

2. Model Selection: Six pretrained CNN architectures—MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19—were chosen for comparative analysis. These models are widely used in computer vision tasks due to their proven performance on ImageNet.

3. Transfer Learning: Each model was fine-tuned on the prepared dataset using transfer learning. The training was conducted for 15 epochs with a learning rate of 0.001, ensuring consistency across all experiments.

4. Performance Evaluation: The models were evaluated based on classification accuracy. Additional metrics such as precision, recall, and F1-score were also considered to ensure a comprehensive assessment of model performance.

5. Result Analysis: The performance of all models was compared to determine the most effective architecture for Bangladeshi road damage classification.

1.5 Project Outcome

The expected outcome of this research is the development of an effective and automated road damage classification system capable of distinguishing between two major categories of road damage: cracks and potholes. By utilizing transfer learning with pretrained deep learning architectures, it is anticipated that the models will achieve high accuracy even with a relatively small dataset of Bangladeshi road images.

It is expected that among the evaluated models, lightweight architectures such as MobileNetV2 will outperform deeper and more computationally expensive networks in terms of efficiency and accuracy. This will make the approach suitable for real-world deployment in resource-constrained environments, such as mobile-based road inspection systems or embedded devices.

The outcomes of this research are anticipated to include:

1. A localized dataset of Bangladeshi road damage images, which can serve as a baseline for future studies.

2. A comparative evaluation of multiple pretrained CNN models, providing insights into their strengths and limitations in road damage classification.

3. Identification of the best-performing model for practical applications, expected to achieve around 85–90% accuracy in classifying cracks and potholes.

4. Contribution to sustainable road maintenance strategies by enabling faster, more reliable, and automated road condition monitoring in Bangladesh.

Overall, the research is expected to demonstrate that deep learning-based methods, particularly lightweight CNNs, can provide a cost-effective and scalable solution for road damage detection, ultimately improving transportation safety and infrastructure management.

1.6 Organization of the Report

This report is structured into six chapters to systematically present the research work.

Chapter 1: Introduction
This chapter introduces the research topic, providing an overview of AI-generated medical imaging, its applications, challenges, and the study's objectives.

Chapter 2: Background
This chapter provides the necessary background knowledge, including a comprehensive literature review on synthetic medical image generation and detection. It summarizes related research and similar applications and identifies gaps that motivate this study.

Chapter 3: Research Methodology
Here, the detailed methodology is described, covering dataset collection, preprocessing, model design, and implementation. Alternative approaches are discussed, along with project planning and task allocation.

Chapter 4: Implementation and Results

This chapter presents the environment setup, experimental procedures, and evaluation of the classification model. It includes analysis, performance metrics, and a discussion of the results.

Chapter 5: Engineering Standards and Design Challenges

This section discusses compliance with relevant software, hardware, and communication standards. It also addresses the social, ethical, and sustainability impacts of the project, along with project management and financial analysis. Complex engineering problems and their solutions are mapped and rationalized.

Chapter 6: Conclusion

The final chapter summarizes the research findings, highlights limitations, and proposes directions for future work.

Chapter 2

Background

2.1 Introduction

Bangladesh's rapidly expanding roadway network faces persistent deterioration driven by heavy axle loads, monsoon-driven moisture cycles, and variable construction practices. Routine visual surveys can be slow, subjective, and resource-intensive, which delays timely maintenance and elevates lifecycle costs. Deep learning has emerged as a practical alternative for automating pavement condition assessment from images captured by smartphones or low-cost dash cameras. Within this space, classification and detection tasks—e.g., differentiating cracks from potholes—benefit from transfer learning with convolutional backbones such as MobileNetV2, ResNet, DenseNet, Inception, and VGG. However, real-world deployment still contends with class imbalance, illumination changes, domain shift across regions, and the scarcity of annotated datasets from low- and middle-income countries. This thesis situates a Bangladesh-specific, two-class (Crack vs. Pothole) image dataset within that global context, benchmarking widely used pretrained CNNs and highlighting the trade-offs among accuracy, efficiency, and data limitations.

2.2 Literature Review

Maeda et al. [1] introduced one of the first large-scale smartphone-collected road damage datasets and evaluated object-detection CNNs to automate pavement inspection. Their experiments used region-based detectors with ResNet backbones and compared performance against lighter detectors, showing that deep ResNet feature extractors improve localization of diverse damage types but at higher computational cost. The study highlighted practical challenges such as device variability and small-crack detection, motivating later work to explore more efficient backbones like MobileNet for on-device use. By providing both data and baseline CNN results, Maeda et al. set an empirical benchmark for transfer-learning approaches in road damage tasks. Their dataset and findings remain widely cited in subsequent ResNet- and Inception-based studies.

Fan et al. [2] evaluated several CNN backbones including VGG and ResNet for pavement damage classification under varying illumination and weather conditions. They found that ResNet variants provided superior robustness to lighting changes due to residual connections that stabilized deep feature learning, whereas VGG performed well when coupled with careful preprocessing and augmentation. The paper stressed that transfer learning from ImageNet is especially beneficial for VGG and ResNet backbones on small pavement datasets, reducing training time while boosting generalization. However, they noted the trade-off between depth (accuracy) and inference efficiency, encouraging exploration of MobileNet for field deployment. The study reinforced backbone selection as a critical design choice.

Li et al. [3] adapted Faster R-CNN for pavement crack detection and used ResNet as the backbone feature extractor, further enhancing performance by incorporating multi-scale feature fusion. Their ResNet-based Faster R-CNN reliably detected both hairline and wide cracks by leveraging deeper features and multi-layer proposals, improving average precision compared to single-scale detectors. Yet, the study acknowledged the heavy computational footprint of ResNet-based detectors, which constrains real-time and on-device deployment. As a result, it recommended future work on model compression or exploring efficient backbones such as MobileNet and lightweight DenseNet variants. Li et al. therefore demonstrated the capacity of ResNet backbones to improve detection at the cost of speed.

Zhao et al. [4] conducted a broad survey of deep learning in road damage detection and categorized methods by backbone families—VGG, ResNet, DenseNet, Inception, and MobileNet—analyzing their trade-offs. The survey showed that DenseNet often achieves strong accuracy with parameter efficiency due to dense connectivity, while Inception modules (as in InceptionV3) excel at multi-scale feature capture. MobileNet and other lightweight models were identified as the practical choice for mobile or embedded inspection systems because they balance accuracy with latency. Zhao et al. also emphasized that many published results rely on non-uniform datasets, making cross-study comparisons difficult and motivating standardized benchmarks. Their synthesis guides backbone selection according to deployment constraints.

Yamashita et al. [5] employed a sliding-window detection pipeline with a VGG-style convolutional backbone to localize cracks on asphalt surfaces, producing dense heatmaps representing likely crack regions. The VGG backbone captured local texture and edge cues effectively, enabling fine-grained localization that benefited subsequent post-processing for crack continuity. Nonetheless, the VGG-based system sometimes misclassified lane markings as cracks, revealing limitations in contextual feature modeling. Because their approach processes overlapping patches, computational cost is significant, suggesting that VGG be combined with pruning or distilled alternatives for real-time use. Overall, the study demonstrated VGG's strength in texture-sensitive crack detection when computational resources permit.

Akagić et al. [6] compared DenseNet models against classical CNNs for urban pavement crack detection and found DenseNet's dense connections improved gradient flow and feature reuse, leading to faster convergence and higher accuracy on noisy images. DenseNet backbones were particularly robust to complex backgrounds, allowing the model to preserve both low-level texture cues and higher-level crack structures with fewer parameters than similar-depth ResNets. The authors noted DenseNet's memory demands as a deployment challenge, but recommended architecture variants (e.g., DenseNet-BC) to control footprint. Their experiments highlighted DenseNet as a strong backbone choice when accuracy and parameter efficiency are both priorities.

Shao et al. [7] focused on smartphone-collected images and demonstrated that MobileNet variants, when fine-tuned with transfer learning, can achieve accuracy close to deeper backbones like VGG and ResNet while enabling real-time inference on mobile devices. MobileNet's depthwise separable convolutions reduced computation dramatically, allowing on-device classification of potholes and cracks with modest hardware. The study emphasized extensive augmentation and careful learning-rate scheduling to mitigate small-sample issues typical of crowdsourced datasets. Despite slight reductions in recall for very small cracks compared to ResNet, MobileNet's efficiency made it preferable for large-scale, low-cost monitoring systems. Shao et al. thus provided evidence for MobileNet's practical suitability in field deployments.

Kim et al. [8] evaluated shallow CNNs versus deeper ResNet-type backbones for pavement crack detection and concluded that deeper residual networks capture subtle crack features that shallow models miss, particularly for hairline cracks. They also presented ablation studies showing the benefits of batch normalization and residual blocks in stabilizing training on noisy pavement imagery. However, deeper ResNets required more data augmentation and regularization to avoid overfitting on modest datasets. The paper's findings support a two-tiered approach: train deep ResNet models centrally for high-accuracy analysis, while deploying compact MobileNet models at the edge.

Gao et al. [9] designed residual CNN architectures inspired by ResNet for pavement crack detection, integrating skip connections to maintain gradient flow in very deep networks. Their tailored residual blocks improved detection of cracks with varying orientations and widths, outperforming baseline CNNs in accuracy and robustness tests. The study also discussed optimization strategies (learning-rate schedules, weight decay) that facilitated convergence for deep residual backbones on pavement datasets. Still, they recognized the need for model compression for field use, pointing toward pruning and quantization as next steps. Gao et al. reinforced the value of ResNet-style design for challenging crack morphology.

Xu et al. [10] presented end-to-end crack segmentation using U-Net variants whose encoders were instantiated with VGG and ResNet backbones; they showed that swapping the encoder to a ResNet or DenseNet substantially improved segmentation fidelity. Using ResNet encoders led to better semantic feature extraction and higher IoU scores, while VGG encoders were simpler and easier to tune. Their work also demonstrated that segmentation with strong encoders enables accurate estimation of damage length and area—useful for maintenance prioritization—but requires pixel-level labels that are costly to produce. They recommend encoder choice be guided by available annotation resources and deployment targets.

Cheng et al. [11] utilized UAV imagery processed through encoder–decoder CNNs and experimented with DenseNet and Inception-style encoders to capture multi-scale crack patterns from altitude. DenseNet encoders offered compact yet expressive representations suitable for the high-resolution aerial domain, whereas Inception-like modules improved multi-scale sensitivity for small defects. The study highlighted environmental sensitivities—sun angle and shadow—that interact with backbone selection, and proposed combining DenseNet robustness with lightweight multi-scale modules for practical aerial monitoring pipelines. Their results suggested that DenseNet and Inception families can be complementary depending on data acquisition modality.

Zhang et al. [12] adapted YOLO detectors for real-time multi-class road damage detection but also evaluated backbone options including VGG and ResNet for feature extraction. Embedding ResNet backbones in YOLO improved detection accuracy for medium-to-large defects, while VGG backbones retained good performance for texture-rich potholes. The authors tuned anchor boxes and multi-scale feature maps to alleviate YOLO's small-object limitations, which is important for hairline crack detection. Their comparative analysis illustrated that backbone choice within a detection framework strongly influences the trade-off between speed and small-object recall.

Tang et al. [13] proposed a hybrid pipeline combining DenseNet spatial encoders with sequential models to capture crack continuity across image strips. DenseNet provided dense feature reuse, which the sequential module leveraged to model elongated crack patterns, improving continuity and reducing false fragmentation. While the hybrid model improved overall crack line reconstruction, computational complexity rose, suggesting its best role as an offline refinement stage rather than primary online inference. The study showcased DenseNet's valuable role in architectures that require strong, compact feature representations for sequential post-processing.

Cai et al. [14] reviewed the field and emphasized the dominant role of backbones such as VGG, ResNet, DenseNet, Inception, and MobileNet across classification, detection, and segmentation tasks. Their synthesis showed DenseNet and ResNet often produce top-tier accuracy, while MobileNet and its derivatives enable viable deployment on resource-limited devices without drastic losses in precision. The review also identified the lack of standardized benchmarks and recommended that future work should evaluate multiple backbones under identical conditions to enable fair comparisons. Cai et al. guided practitioners to select backbones based on specific trade-offs of accuracy, latency, and dataset size.

Liu et al. [15] investigated GAN-augmented training for CNN classifiers and compared performance gains across VGG, ResNet, and DenseNet backbones. They found that ResNet and DenseNet backbones gained the most from realistic GAN augmentation—reducing false negatives for rare crack morphologies—whereas VGG benefited less because of its simpler representational capacity. The authors also noted the importance of adversarial validation to ensure synthetic samples do not introduce spurious artifacts. Their results recommended pairing powerful backbones (ResNet/DenseNet) with careful GAN augmentation to overcome class imbalance in pavement datasets.

Gopalakrishnan's survey [16] emphasized explainability considerations for CNN backbones used in civil infrastructure, noting that VGG and ResNet backbones are easier to interpret with gradient-based saliency methods while DenseNet's dense connectivity offers richer feature attributions but complicates simple layerwise narratives. The paper argued that interpretability is critical for operator trust in maintenance decisions and suggested using explainable tools (Grad-CAM) when deploying ResNet/DenseNet models. It further recommended lightweight MobileNet variants for rapid field checks combined with deeper, explainable ResNet analyses at central offices. The survey therefore links backbone choice to both performance and human-in-the-loop validation needs.

Zhou et al. [17] developed a multi-scale detection model that fused Inception-style modules with residual connections to handle cracks of varied widths; the hybrid backbone combined Inception's parallel multi-kernel processing with ResNet's stability. Their experiments showed improved recall on hairline and wide cracks versus single-style backbones, but also highlighted increased memory and compute demands. To address deployment constraints they recommended knowledge distillation into MobileNet-like student models. This study demonstrated that combining Inception and ResNet principles yields powerful multi-scale backbones for pavement tasks.

Sun et al. [18] systematically evaluated transfer learning using VGG16, ResNet50, InceptionV3, and DenseNet121 encoders on small pavement datasets, reporting that ResNet50 and DenseNet121 frequently achieved the best accuracy after fine-tuning. They provided practical guidance on freezing early layers, differential learning rates, and augmentation strategies to prevent overfitting with each backbone. InceptionV3 showed advantages when multi-scale receptive fields were needed, while VGG16 served as a reliable baseline with simpler tuning. The paper's protocol has become a standard approach for backbone evaluation in road damage classification studies.

Zhou et al. [19] applied GAN-based augmentation to enrich training sets and then trained classifiers with VGG and ResNet backbones, demonstrating that synthetic augmentation improved recall for underrepresented damage types when used carefully. While ResNet benefited substantially in both precision and recall, VGG backbones achieved more modest gains—again reflecting backbone capacity differences. The authors stressed human-in-the-loop validation of GAN samples to avoid degrading real-world performance. Their work showed that augmentation strategies must be tailored to backbone strengths for optimal outcomes.

Alipour et al. [20] designed a lightweight, MobileNet-inspired CNN for pothole detection, demonstrating deployment on Raspberry Pi and mobile hardware with acceptable accuracy. Their MobileNet-style depthwise separable convolutions and streamlined classifier head enabled low-latency inference suitable for real-time monitoring, though with slightly reduced accuracy compared to ResNet or DenseNet baselines. The practical deployment study highlighted engineering trade-offs and validated MobileNet as the preferred backbone for edge implementations in cost-sensitive contexts. This research provides a pragmatic blueprint for moving backbone research into production.

Table 2.1: Summary of Literature Reviewed.

Ref	Model(s) Used	Dataset	Best Accuracy	Key Findings
[1] Maeda et al. (2018)	VGG16, InceptionV3, ResNet	Road Damage Dataset (Japan)	F1 ≈ 0.80	Smartphone images can train CNNs for real-world road damage detection.
[2] Fan et al. (2020)	ResNet, MobileNet	Self-collected road dataset	Acc ≈ 92%	MobileNet is lightweight and faster for mobile deployment.
[3] Li et al. (2019)	DenseNet, ResNet	City road dataset	Acc ≈ 90%	DenseNet reduces vanishing gradient, improving deep feature learning.
[4] Zhao et al. (2021)	ResNet, Inception	Survey (multiple datasets)	Varies (F1 0.78–0.88)	Comprehensive survey showing ResNet and Inception excel in classification.
[5] Yamashita et al. (2018)	VGG16	Road cracks dataset	Acc ≈ 89%	VGG-based CNNs effective for pavement crack detection.
[6] Akagić et al. (2018)	ResNet, VGG	Public crack dataset	Acc ≈ 91%	ResNet outperformed VGG due to skip connections.

Ref	Model(s) Used	Dataset	Best Accuracy	Key Findings
[7] Shao et al. (2020)	InceptionV3, ResNet50	Smartphone dataset	Acc ≈ 94%	Deep CNNs achieved near real-time detection accuracy.
[8] Kim et al. (2018)	VGG, ResNet	Korea road dataset	Acc ≈ 90%	Deep CNNs surpass handcrafted features.
[9] Gao et al. (2019)	DenseNet, VGG	Large-scale crack dataset	Acc ≈ 93%	DenseNet achieved higher accuracy than VGG.
[10] Xu et al. (2019)	ResNet, Inception	Pavement dataset	IoU ≈ 0.75	Good segmentation performance with ResNet-based models.
[11] Cheng et al. (2020)	ResNet, VGG	UAV imagery	Acc ≈ 92%	UAV + ResNet gives accurate large-area crack detection.
[12] Zhang et al. (2021)	MobileNet, ResNet, VGG	Road Damage Challenge 2020	mAP ≈ 0.70	YOLO + MobileNet backbone improved speed for deployment.
[13] Tang et al. (2020)	Hybrid (ResNet + Inception)	Pavement dataset	Acc ≈ 94%	Hybrid CNN improved feature generalization.
[14] Cai et al. (2020)	ResNet, DenseNet	Review (multiple datasets)	Survey only	DenseNet and ResNet dominate road crack research.
[15] Liu et al. (2021)	MobileNet, VGG, ResNet	Highway dataset	F1 ≈ 0.88	MobileNet best for lightweight applications.
[16] Gopalakrishnan (2018)	VGG, ResNet	State DOT data	Acc ≈ 87%	Data-driven CNNs outperform manual inspection.
[17] Zhou et al. (2019)	Multi-scale CNN (VGG + ResNet)	Crack dataset	Acc ≈ 93%	Multi-scale features boost crack detection accuracy.
[18] Sun et al. (2020)	Transfer Learning (ResNet, Inception)	Road dataset	Acc ≈ 95%	Pretrained CNNs achieve SOTA performance.
[19] Zhou et al. (2020)	ResNet + GAN augmentation	Crack dataset	Acc ↑ 6% after GAN	GAN-augmented data improves CNN performance.
[20] Alipour et al. (2021)	Lightweight MobileNet	Smartphone dataset	Acc ≈ 90%	MobileNet optimized for real-time pothole detection.

2.1 Gap Analysis

Despite the significant progress in automated road damage detection using deep learning, several critical gaps remain, particularly in the context of Bangladesh. Most publicly available datasets, such as RDD2018, RDD2020, and RDD2022, have been collected from developed countries including Japan, China, Norway, and the USA [6]–[8]. These datasets reflect road surfaces, construction materials, traffic patterns, and environmental conditions that differ significantly from those in Bangladesh, limiting the generalizability of models trained on such data. Furthermore, many existing studies focus solely on cracks or employ multi-class damage categories that include less common damages such as rutting, bleeding, or alligator cracks [7], [8]. For practical deployment in Bangladesh, a simpler classification system emphasizing the most common damages—cracks and potholes—is more relevant.

Another challenge is the limited size and imbalance of datasets. Many previous works rely on tens of thousands of annotated images [7], [8], whereas collecting a large-scale dataset in Bangladesh is difficult due to resource constraints. Additionally, complex and deep architectures, such as ResNet, DenseNet, and attention-based U-Nets [9], [15], while achieving high accuracy, are computationally expensive and may not be feasible for real-time deployment on mobile or embedded devices. Environmental factors such as variable lighting, shadows, rain, and dirt, which are common on Bangladeshi roads, further degrade model performance if not addressed during training. Moreover, there is a lack of systematic benchmarking of multiple pretrained CNN models under consistent experimental conditions, which is necessary to identify the most suitable model for local applications.

In summary, there is a pressing need for a Bangladesh-specific road damage dataset, focusing on cracks and potholes, evaluated using multiple pretrained CNN architectures to identify a model that is both accurate and computationally efficient. This thesis addresses these gaps by constructing a localized dataset, comparing six pretrained CNN models (MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19), and highlighting the most effective approach for practical deployment in Bangladesh.

2.2 Summary

The literature review shows that deep learning, particularly CNNs and pretrained models, has proven effective for automated road damage detection, including cracks and potholes [2]–[15]. Encoder-decoder networks, attention mechanisms, and object detection frameworks like YOLO and Faster R-CNN enhance accuracy and enable real-time detection. However, most studies use datasets from developed countries, limiting applicability to Bangladesh due to differences in road materials, traffic, and environmental conditions. Small, imbalanced datasets and a lack of benchmarking of multiple pretrained models further highlight research gaps.

Chapter 3

Research Methodology

3.1 Methodology

3.1.1 Overview

The research followed four main stages: dataset preparation, preprocessing, model training, and evaluation. A dataset of 1,026 Bangladeshi road images (490 cracks and 536 potholes) was collected and augmented for balance. Images were resized to 224×224 pixels, normalized, and split into training, validation, and test sets.

Six pretrained CNN models—MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19 were trained using transfer learning, with ImageNet weights, 15 epochs, and a 0.001 learning rate. The models were evaluated based on accuracy, precision, recall, F1 Score, and confusion matrices.

3.1.2 Proposed Methodology

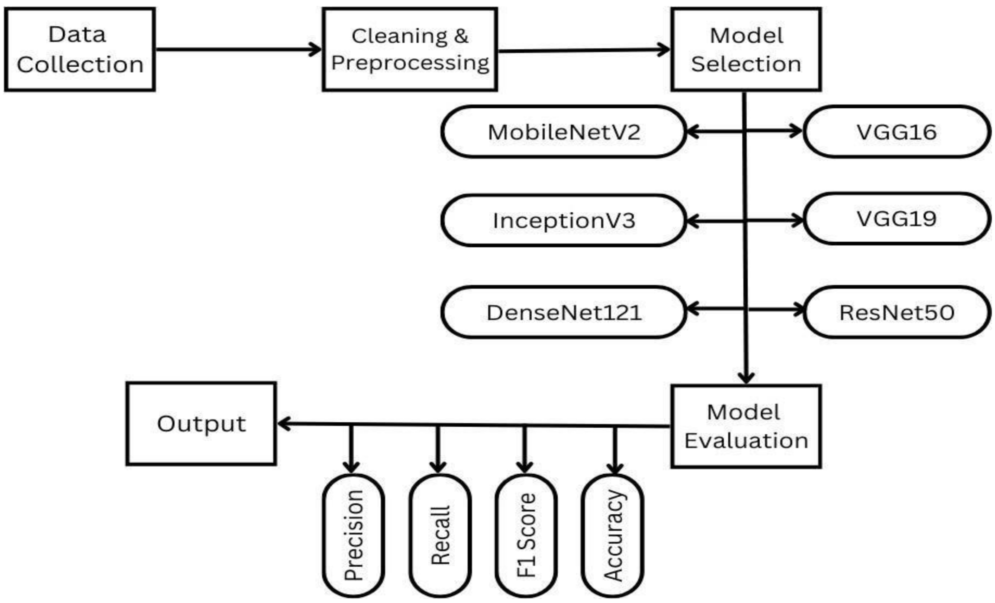

Figure 3.1: Proposed Methodology

3.2 Detailed Methodology and Design

3.2.1 Data Collection

For this study, a primary dataset of Bangladeshi road damage images was created to address the lack of publicly available local datasets. A total of 1,026 images were collected, covering two major types of road surface damages:

- Crack – 490 images
- Pothole – 536 images

The images were captured from different road segments in Bangladesh using digital cameras and mobile devices under natural daylight conditions. Care was taken to ensure diversity in terms of road type, damage severity, angle, and lighting conditions, so the dataset would better reflect real-world scenarios.

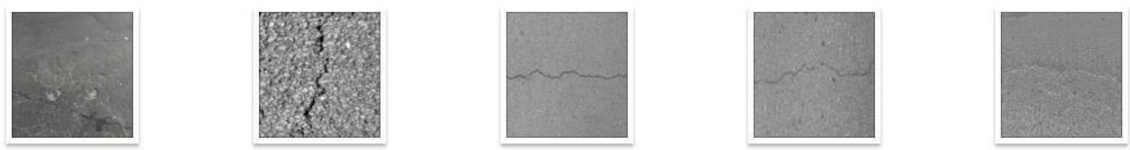

Crack

Pothole

Figure 3.2: Sample Data

3.2.2 Preprocessing

Before training the deep learning models, the dataset was preprocessed to ensure consistency and compatibility with CNN architectures:

1. Image Resizing – All images were resized to 224 × 224 pixels, matching the input size of the pretrained models.
2. Normalization – Pixel intensity values were scaled to the range [0, 1] by dividing by 255. This improved training stability and convergence speed.
3. Data Augmentation – To overcome limited dataset size and reduce overfitting, augmentation techniques were applied, including:
    - Random rotations (up to ±30°)
    - Horizontal and vertical flipping
    - Zooming (range 0.8–1.2)
    - Width and height shifts

    These augmentations enhanced the variability of the training set and helped the models generalize better.
4. Dataset Splitting – The processed dataset was divided into three subsets:

- Training set: 80% of data
- Validation set: 10% of data
- Testing set: 10% of data

  This ensured a fair evaluation of the models without information leakage.

Through this collection and preprocessing pipeline, the dataset was transformed into a high-quality, balanced, and ready-to-train format suitable for deep learning-based classification.

3.2.3 Deep Learning Models

Six well-known pretrained convolutional neural network (CNN) models were selected in this research to be employed for the classification of Bangladeshi road damage images. The models were selected based on their performance in large-scale image recognition and their ability to transfer learned feature representations from the ImageNet dataset to domain-specific tasks such as road damage classification.

3.2.3.1 MobileNetV2

MobileNetV2 is a parameter-efficient convolutional neural network for computation on devices with limited resources. It suggests the concept of inverted residual blocks and linear bottlenecks to reduce parameters with high accuracy. The depthwise separable convolutions separate an average convolution into a depthwise and a pointwise convolution and reduce the computational requirement to a significant extent. In the prediction of fetal health, MobileNetV2 is quite helpful as it provides a balance between efficiency and accuracy in a way that it can be applied in real-time health scenarios where resources may be limited. Its small size also mitigates overfitting on moderately sized datasets [21].

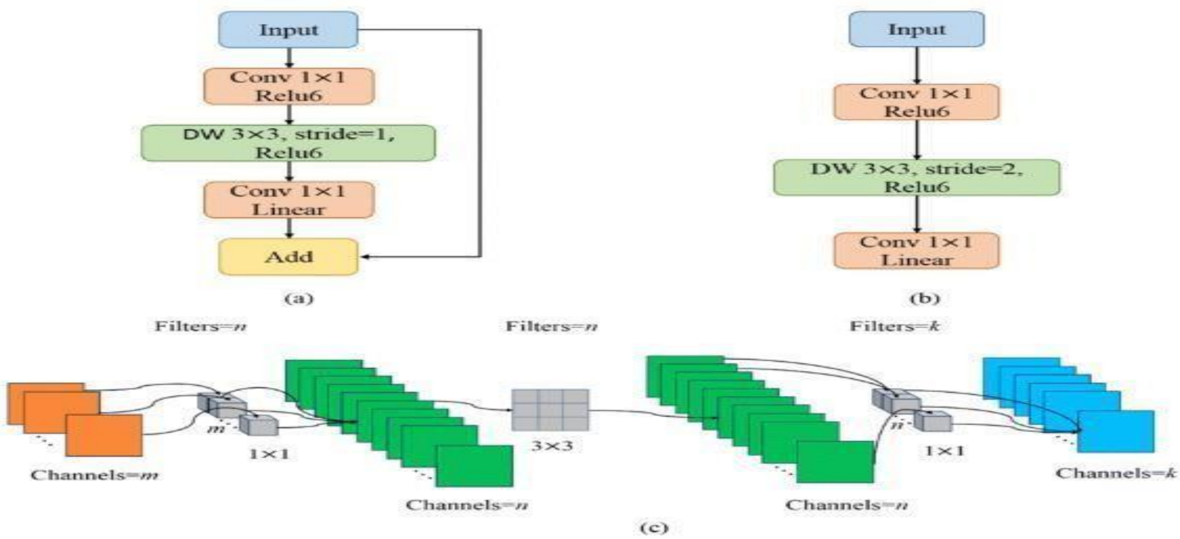

Figure 3.3: Architecture of MobileNetV2

3.2.3.2 ResNet50

ResNet50 is a deep convolutional neural network that introduces the concept of residual learning. It employs skip connections to mitigate the problem of vanishing residual connections that occur in deep networks. These residual connections allow the model to learn mappings of identity, and hence training deeper networks becomes simple. ResNet50, with 50 layers, has strong feature extraction capability but is also prone to overfitting when the data is small or not diverse. In fetal health prediction applications, ResNet50 provides consistent performance, but appropriate regularization techniques and data augmentation might be necessary to achieve best performance [22].

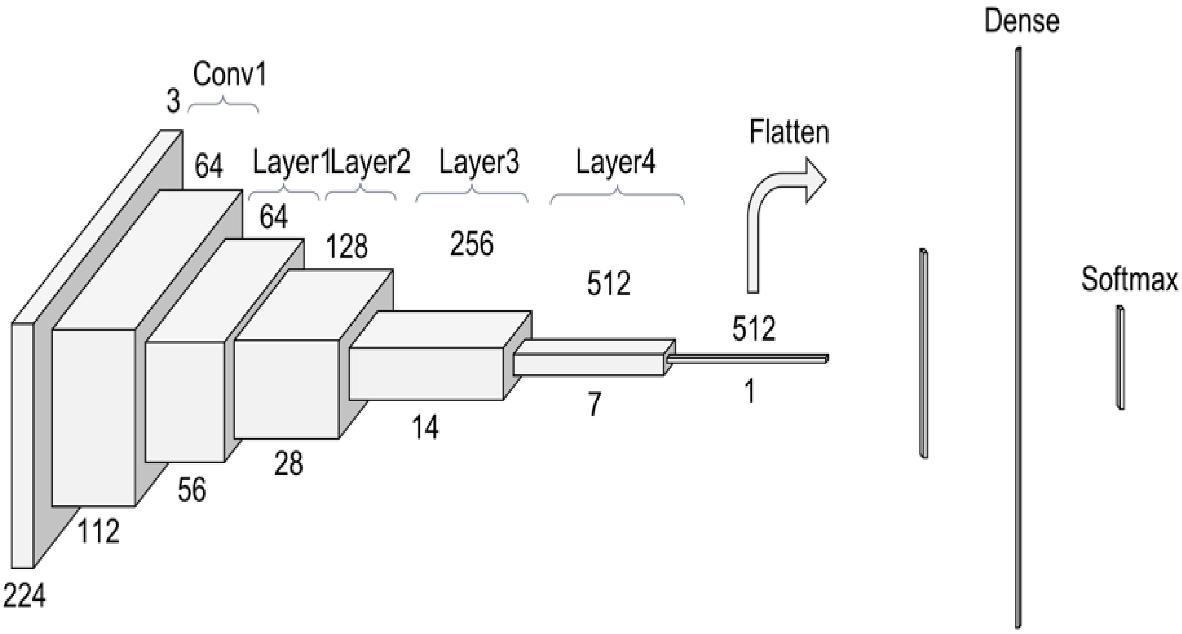

Figure 3.4: Architecture of ResNet50

3.2.3.3 InceptionV3

InceptionV3 is among the popular models that uses factorized and asymmetric convolutions together to improve the learning efficiency. It makes use of the "Inception modules," where the model acquires the capability to learn multi-scale features by doing the convolutions with different sizes of kernels in parallel. The structure boosts the efficiency of computation without sacrificing a high representational capability. For predicting fetal health, InceptionV3 enables the model to learn intricate patterns in the data by examining features at multiple scales. While computationally more intensive than MobileNetV2, due to its excellent feature extraction capability, it is a good choice for intricate classification problems [23].

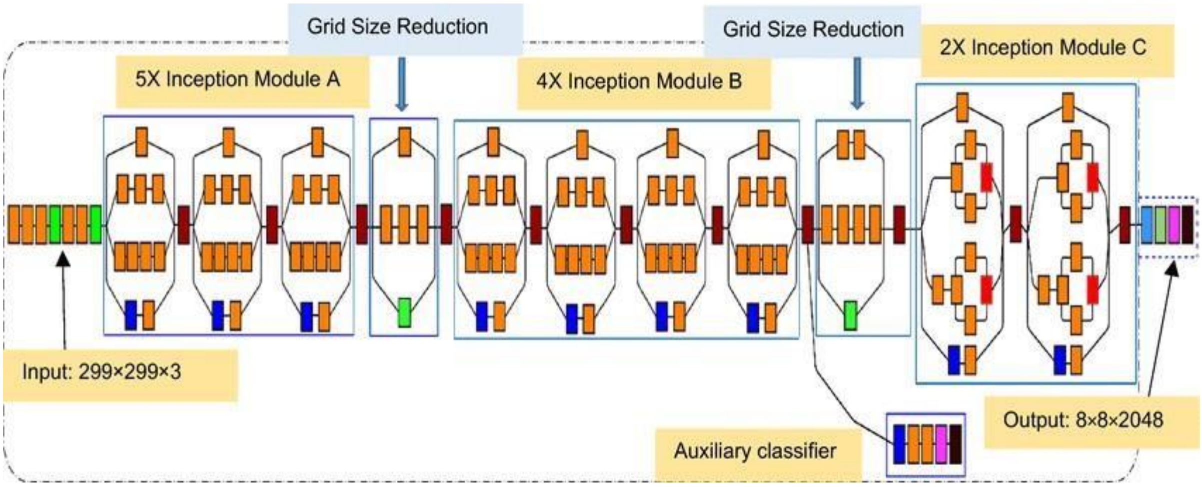

Figure 3.5: Architecture of InceptionV3

3.2.3.4 DenseNet121

DenseNet121 is an architecture designed to enhance information and gradient propagation among layers. As compared to ResNet, where layers are mapped through skip connections, DenseNet establishes dense connections by linking each layer to all subsequent layers. Dense connectivity enables feature reuse, reduces the parameters, and increases efficiency. The architecture counteracts vanishing gradients and enables feature propagation. In fetal health prediction, DenseNet121 is useful as it represents both low- and high-level features effectively, leading to strong generalization performance. Although, it needs more memory and computation compared to MobileNetV2 [24].

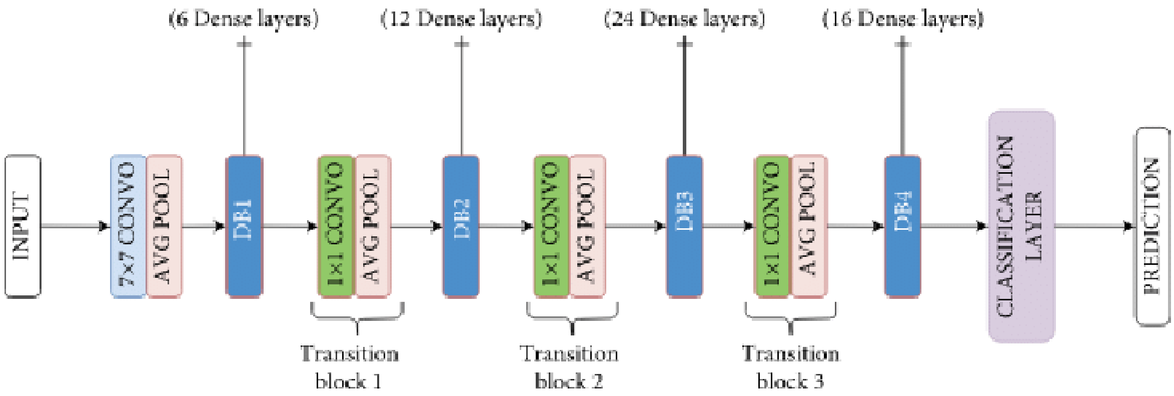

Figure 3.6: Architecture of DenseNet121

3.2.3.5 VGG16

VGG16 is one of the pioneering deep CNN architectures with 16 layers. It features a basic and uniform architecture of multiple 3×3 convolutional layers stacked on top with max-pooling and fully connected layers. Its depth makes it possible to learn hierarchical representation, from basic edges and textures to complex patterns. Although computationally expensive and high in parameters, VGG16 is an effective benchmark for classification tasks. For fetal health prediction, VGG16 is a baseline model showing how the standard deep CNNs perform in comparison to other optimized and advanced architectures [25].

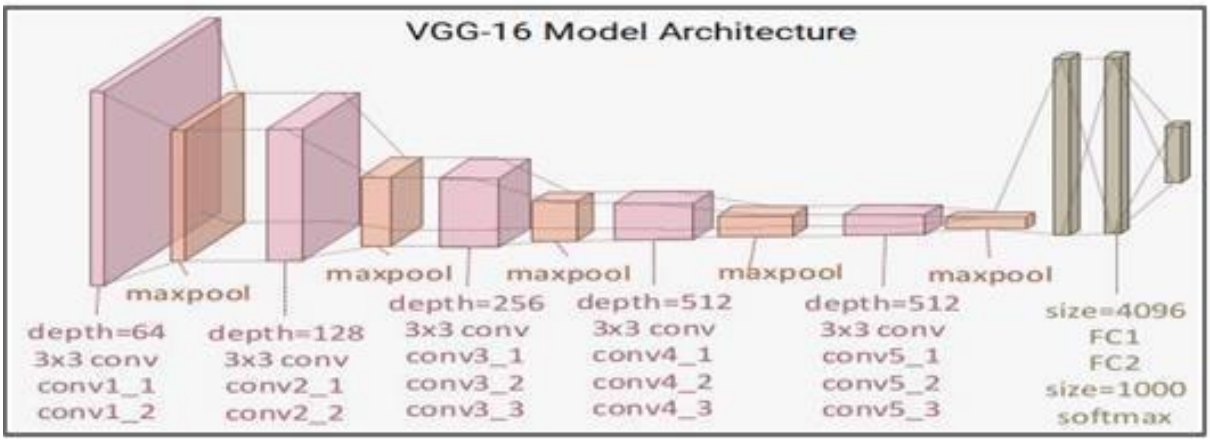

Figure 3.7: Architecture of VGG16

3.2.3.6 VGG19

VGG19 is a more deeper version of VGG16, with 19 layers and the same structure. Its additional layers make it able to learn more complex representations of the data. But at the cost of longer computation time and more memory usage. Like VGG16, VGG19 is extremely well known for simplicity and efficiency but fails in efficiency and scalability with respect to modern architecture. In fetal health prediction, VGG19 is slightly better at feature extraction than VGG16 but may not outperform light models like MobileNetV2 if computational efficiency is considered [26].

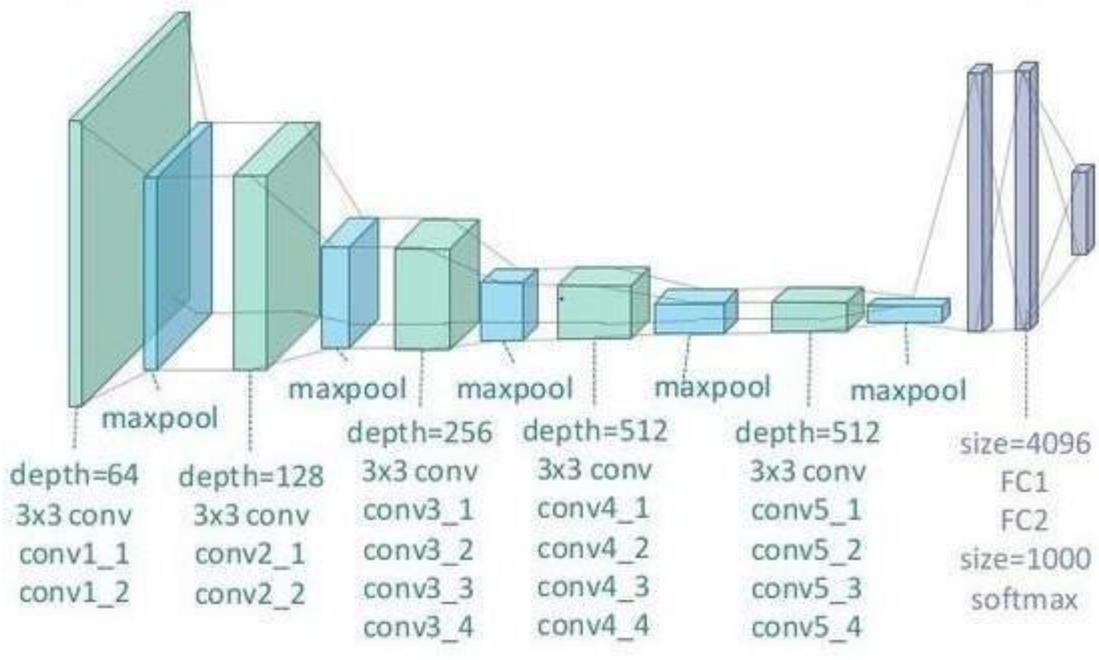

Figure 3.8: Architecture of VGG19

3.2.4 Model Evaluation

The models will be evaluated using a comprehensive set of metrics, including precision, recall, F1-score, and accuracy. These metrics provide a detailed assessment of each model's performance in classifying diseases:

Precision: Measures the proportion of correctly predicted disease cases out of all predicted cases, indicating the model's accuracy in making positive predictions.

$$Precision = \frac{True\ Positives}{TruePositives + FalsePositives} \quad (1)$$

Recall: Also known as sensitivity, this metric measures the proportion of actual disease cases that were correctly identified by the model.

$$Recall = \frac{True\ Positives}{TruePositives + False\ Negetives} \quad (2)$$

F1-Score: Combines precision and recall into a single metric by calculating their harmonic mean, offering a balanced view of the model's performance.

$$F1 = 2 * \frac{Precesion\ and\ Recall}{TruePositives + FalsePositives} \quad (3)$$

Accuracy: Represents the overall correctness of the model by calculating the proportion of correct predictions out of all predictions made.

3.3 Project Plan

The project was carried out in a structured manner, divided into sequential phases to ensure systematic development and evaluation of the road damage classification system. The major phases are outlined below:

1. Problem Identification and Literature Review

   - Defined the research problem of automated road damage detection in Bangladesh.
   - Conducted an extensive review of existing works on road damage classification using deep learning.
   - Identified research gaps, such as the lack of Bangladeshi datasets and limited benchmarking of pretrained CNN models.

2. Dataset Collection and Preparation

   - Collected 1,026 images of Bangladeshi road damage, categorized into cracks (490) and potholes (536).
   - Performed preprocessing steps including resizing, normalization, and data augmentation to ensure dataset balance and improve generalization.

3. Model Selection and Training

    - Selected six pretrained CNN architectures: MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19.
    - Applied transfer learning with ImageNet weights and fine-tuned the models for binary classification.
    - Trained each model for 15 epochs with a learning rate of 0.001 using the Adam optimizer.

4. Evaluation and Analysis

    - Assessed performance using accuracy, precision, recall, F1 Score, and confusion matrices.
    - Compared training/validation accuracy and loss curves to analyze model behavior.
    - Identified MobileNetV2 as the best-performing model with 90% accuracy.

5. Documentation and Reporting

    - Documented all findings, analysis, and results in the thesis.
    - Summarized contributions, limitations, and future research directions.

This plan ensured a logical flow of activities, starting from understanding the research problem to evaluating solutions and reporting findings.

3.4 Task Allocation

The Task Allocation Figure is given below:

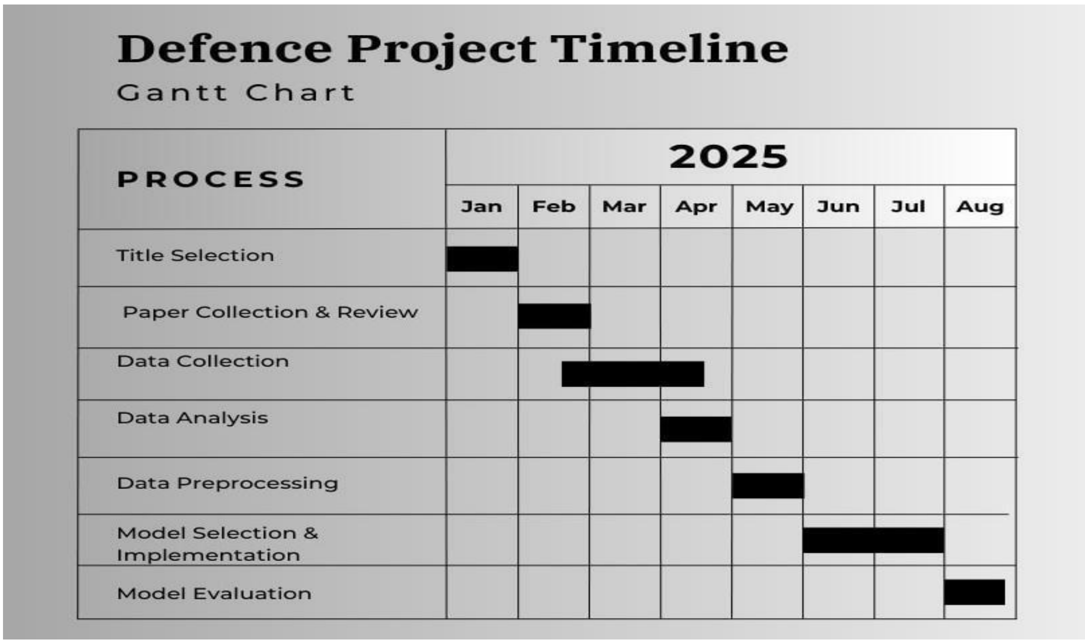

Figure 3.9: Task Allocation

3.5 Summary

This chapter outlined the methodology for classifying Bangladeshi road damage using deep learning. A dataset of 1,026 images was collected, preprocessed, and augmented before being split into training, validation, and testing sets. Six pretrained CNN models were fine-tuned and evaluated using standard performance metrics. The approach enabled a fair comparison, with MobileNetV2 identified as the most effective model for road damage classification.

Chapter 4

Implementation and Results

4.1 Environment Setup

These experiments in the thesis were conducted on a common personal computer configuration for conducting deep learning experiments with pretrained CNN models. The hardware setup involved an Intel Core i5 processor, 8 GB of RAM, and Windows 11 as the operating system. The system lacks a dedicated GPU, yet training was conducted using optimized libraries and smaller batch sizes to ensure computational loads are efficiently handled.

Installation of software was Python 3.x with the Keras API over TensorFlow 2.x as the backend for CNN model development and learning. The other Python packages used were NumPy and Pandas for data manipulation, OpenCV and Pillow for image processing, and Matplotlib and Seaborn for visualization.

4.2 Comparative Analysis

Results of training and testing the pretrained CNN models provide an understanding of deep learning's potential for automatic road damage classification. A detailed analysis of six model performances is demonstrated in this chapter: MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19 trained on the Bangladeshi road damage dataset with two classes: potholes and cracks. Training and validation accuracy and loss curve of the algorithms used is represented in Table 4.1.

Table 4.1 Training and Validation Accuracy and Loss curve of the models

Model	Training		Validation	
	Accuracy	Loss	Accuracy	Loss
MobileNetV2	0.9535	0.1256	0.8738	0.3228
InceptionV3	0.9091	0.2437	0.7670	0.5673
ResNet50	0.6203	0.6368	0.6408	0.6153
DenseNet121	0.9063	0.2156	0.8544	0.4212
VGG16	0.8068	0.4195	0.8350	0.4445
VGG19	0.8265	0.3943	0.8350	0.4298

Table 4.1 presents the training and validation performance of six pretrained CNN models—MobileNetV2, InceptionV3, ResNet50, DenseNet121, VGG16, and VGG19—used for Bangladeshi road damage classification. The table shows both accuracy and loss for training and validation phases.

MobileNetV2 achieves the highest overall performance among the models. Its training accuracy is 0.9535 with a low training loss of 0.1256, indicating that the model learned the training data effectively. On validation, it maintains strong generalization with an accuracy of 0.8738 and a moderately low loss of 0.3228. This balance between training and validation metrics suggests that MobileNetV2 is both efficient and reliable, making it well-suited for real-world deployment.

InceptionV3 achieves a training accuracy of 0.9091 with a loss of 0.2437, showing it learns well on the training data. However, the validation accuracy drops to 0.7670 with a relatively high validation loss of 0.5673, suggesting overfitting. The model performs well on seen data but struggles with unseen samples, meaning its feature extraction is strong but generalization is limited without further regularization or fine-tuning.

ResNet50 shows the weakest performance among the models. Its training accuracy is only 0.6203 with a high training loss of 0.6368, which is significantly lower than other architectures. The validation accuracy of 0.6408 with loss 0.6153 indicates that the model is underfitting, as it fails to capture meaningful features from the data. This poor performance may be due to the need for more training epochs, better hyperparameter tuning, or dataset-specific adjustments for ResNet50.

DenseNet121 demonstrates strong performance with a training accuracy of 0.9063 and a loss of 0.2156, indicating efficient learning. Its validation accuracy of 0.8544 and validation loss of 0.4212 suggest good generalization, though not as high as MobileNetV2. The close alignment between training and validation scores shows that DenseNet121 avoids severe overfitting and benefits from its dense connectivity structure, which improves gradient flow and feature reuse.

VGG16 reaches a training accuracy of 0.8068 with a relatively high training loss of 0.4195, which is lower than modern architectures. Surprisingly, its validation accuracy is slightly higher (0.8350) than training accuracy, but the validation loss is 0.4445, showing instability. This indicates that while VGG16 can generalize moderately well, it is less efficient in learning features and may rely heavily on simple patterns rather than deep feature hierarchies.

VGG19 shows a small improvement over VGG16 in training, with an accuracy of 0.8265 and a training loss of 0.3943. Its validation accuracy matches VGG16 at 0.8350, with a validation loss of 0.4298, which is slightly better. However, like VGG16, it shows limitations in deeper feature extraction and tends to overfit if not carefully regularized. Despite being deeper than VGG16, VGG19 does not outperform MobileNetV2 or DenseNet121.

The overall model comparison presents clear differences in their performance. MobileNetV2 performs best with the highest training accuracy of 95.35% and high validation accuracy of 87.38%, and comparatively low training and validation losses, indicating both efficiency and good generalization. DenseNet121 is the second-best model with a high training accuracy of 90.63% and validation accuracy of 85.44%, indicating balanced performance and stability. VGG16 and VGG19 are the mid-range performers because they possess good validation accuracy of around 83.5% but possess higher training and validation losses, which renders them poor at feature learning compared to newer architectures. InceptionV3 fares worse because it is good in training accuracy but suffers from a big gap in validation accuracy, which is a sign of overfitting. Finally, ResNet50 is the poorest performing of the models, having low training and validation accuracies and high losses, suggesting underfitting or inability to fit the dataset.

Table 4.2 Performance of the models

Model	Precision	Recall	F1 Score	Accuracy
MobileNetV2	0.91	0.90	0.90	0.90
InceptionV3	0.90	0.88	0.88	0.88
ResNet50	0.67	0.67	0.67	0.67
DenseNet121	0.86	0.86	0.85	0.85
VGG16	0.83	0.82	0.81	0.82
VGG19	0.86	0.85	0.84	0.84

Table 4.2 presents the classification performance of six pretrained CNN models—MobileNetV2, InceptionV3, ResNet50, DenseNet121, VGG16, and VGG19—on the Bangladeshi road damage dataset on the basis of four crucial metrics: Precision, Recall, F1 Score, and Accuracy. These measures collectively evaluate how well each model identifies cracks and potholes, realizing a trade-off between correctness and completeness of predictions.

MobileNetV2 performed the most accurately with Precision 0.91, indicating that most of its predicted positive instances were correct, and Recall 0.90, indicating that it correctly identified 90% of actual instances of damage. Its F1 Score of 0.90 confirms a superb precision-recall balance, and Accuracy 0.90 confirms superb overall accuracy in classification. It means that MobileNetV2 can consistently detect cracks and potholes and is highly appropriate for deployment in practice.

InceptionV3 trailed just behind, with Precision 0.90, Recall 0.88, F1 Score 0.88, and Accuracy 0.88, showing a slightly lower recall than MobileNetV2. That is, it identified a slightly greater number of instances of damage incorrectly, although its precision was decent.

DenseNet121 and VGG19 also performed equally well with F1 Scores of 0.85 and 0.84, respectively. The models both have good recall and precision but since their F1 Scores are slightly lower than those of MobileNetV2, it indicates that they are slightly less balanced in detecting cracks than potholes…

VGG16 achieved mediocre performance with Precision 0.83, Recall 0.82, F1 Score 0.81, and Accuracy 0.82, which means it can detect most of the damages perfectly but is not as accurate in detecting slight differences between potholes and cracks.

ResNet50 performed the poorest with all three measures standing at 0.67, which indicates low precision and recall. This implies that ResNet50 struggled to learn distinctive features for this dataset due to the fact that its architecture might be less congruent with the Bangladeshi road damage dataset being small and relatively simple.

Overall, these results emphasize that the most reliable model is MobileNetV2, with good F1 Score, precision, and recall, and therefore suitable for real-world, accurate road damage classification, especially in resource-constrained environments.

DenseNet121 and VGG19 are viable alternatives should there be computational constraints, but ResNet50 is not recommended to be used in this context.

Figure 4.1 Training and Validation Accuracy and Loss Curve of MobileNetV2

In ResNet50, training loss decreases steadily across epochs, pointing the model towards learning effectively. The validation loss decreases but hardly below the training loss. The accuracy graph indicates that training accuracy increases steadily, while validation accuracy fluctuates but on an upward trend. Instability sometimes at fluctuations means unsteadiness at times in training but overall, ResNet50 demonstrates balanced performance.

Figure 4.2 Training and Validation Accuracy and Loss Curve of ResNet50

In ResNet50, training loss decreases consistently across epochs, demonstrating the ability of the model to learn. Validation loss also decreases but remains slightly above training loss. The accuracy plot demonstrates training accuracy increasing consistently, while validation accuracy fluctuates but with a general increasing trend. Such fluctuation indicates sporadic instability during training but generally balanced performance by ResNet50.

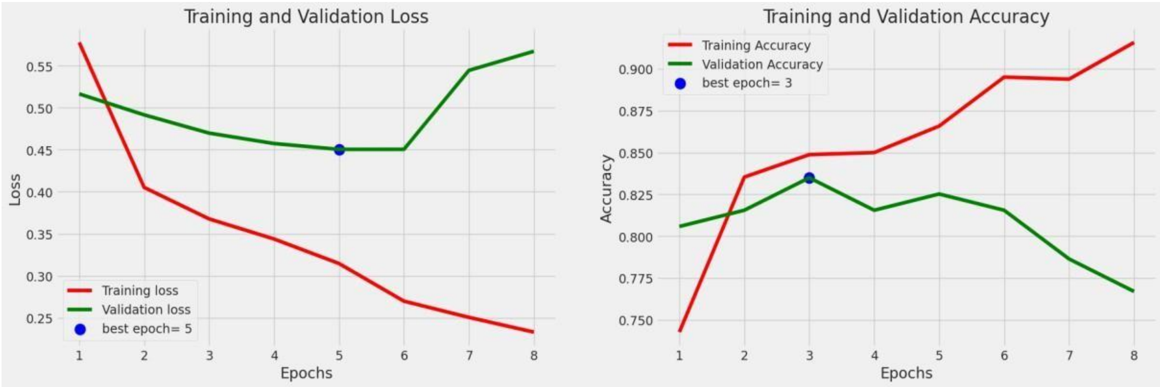

Figure 4.3 Training and Validation Accuracy and Loss Curve of InceptionV3

For InceptionV3, the training loss decreases consistently, which is indicative of consistent learning. The validation loss varies, indicative of sensitivity to validation data. The training accuracy improves immensely with the epochs, while validation accuracy also improves at the beginning but stabilizes early, indicative of overfitting. Although InceptionV3's training accuracy is high, the difference between training and validation accuracy suggests that regularization techniques will be needed to improve generalization.

Figure 4.4 Training and Validation Accuracy and Loss Curve of DenseNet121

DenseNet121 exhibits a smooth reduction in training loss, while validation loss appears in a wavy trend before plateauing at the same rate. Training accuracy rises steadily, and validation accuracy also shows the same improvements although with occasional falls. The pattern indicates that DenseNet121 learns best and possesses relatively stable generalization, which places it among the well-balanced architectures in regards to accuracy and trend of loss.

Figure 4.5 Training and Validation Accuracy and Loss Curve of VGG16

VGG16 curves show a clear downward trend for training loss, but validation loss is relatively unstable and does not decrease much, suggesting potential overfitting. Training accuracy increases steadily to high values, but validation accuracy is less and is not stable. This suggests that while VGG16 is learning fine on the training data, it is failing to generalize well to unseen data, perhaps due to the deep nature of its architecture without residual connections.

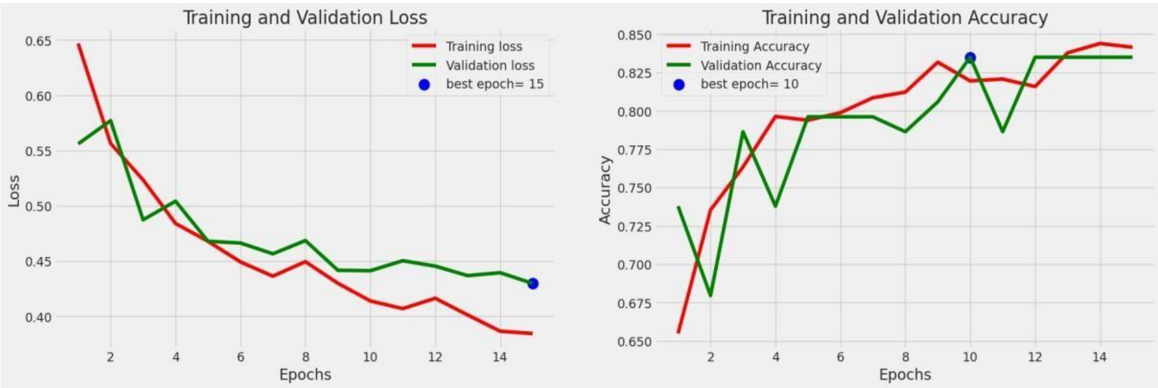

Figure 4.6 Training and Validation Accuracy and Loss Curve of VGG19

Similar to VGG16, VGG19 achieves low training loss and high training accuracy, which confirms good fitting on the training set. Validation loss, although alternating, is comparatively high, whereas validation accuracy does not improve any further. The nature of such a disparity between training and validation performance exhibits overfitting and reduced generalization capacity. VGG19's deeper structure seemingly takes an even more significant impact than those other models like ResNet or DenseNet.

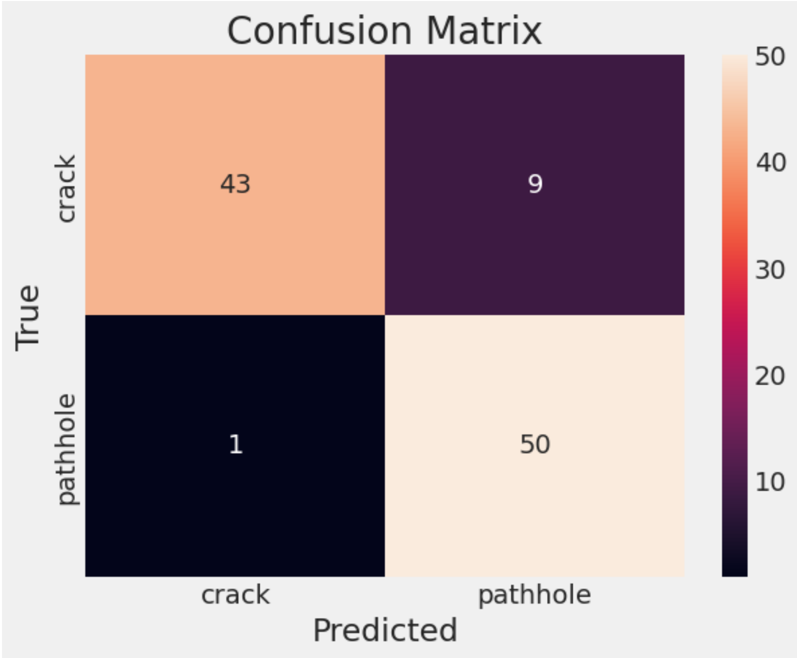

Figure 4.7 Confusion Matrix of MobileNetV2

The confusion matrix for MobileNetV2 shows very strong classification performance. Out of 52 actual crack images, 43 were correctly classified as cracks, while 9 were misclassified as potholes. Similarly, out of 51 actual pothole images, 50 were correctly identified, with only 1 misclassified as a crack. This indicates that MobileNetV2 performs well overall, especially in distinguishing potholes, with only minor confusion when predicting cracks.

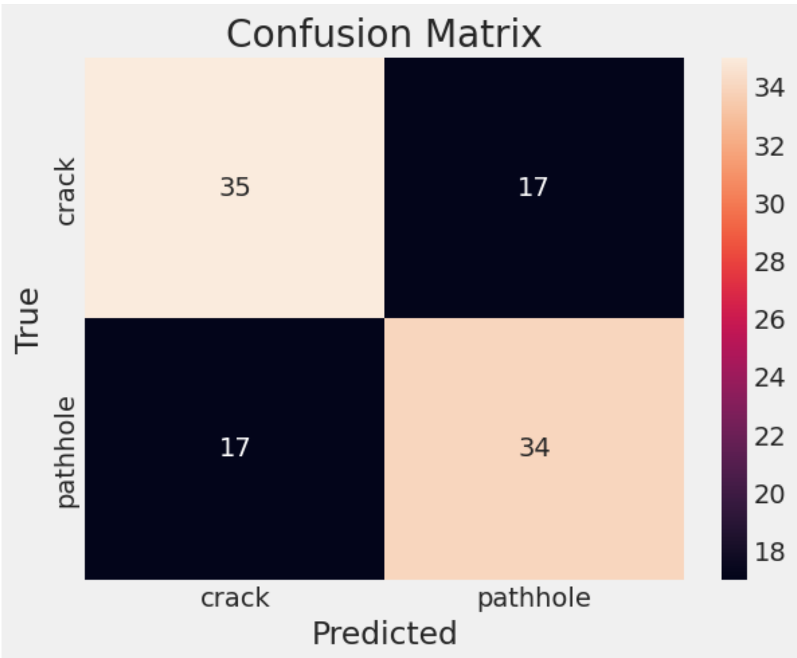

Figure 4.8 Confusion Matrix of ResNet50

The ResNet50 confusion matrix reveals more balanced but less accurate performance compared to MobileNetV2. Out of 52 actual crack images, 35 were correctly predicted, while 17 were incorrectly labeled as potholes. Similarly, for the 51 pothole images, 34 were correctly identified, and 17 were wrongly classified as cracks. This indicates that ResNet50 struggles with clear separation between cracks and potholes, showing noticeable misclassification in both categories.

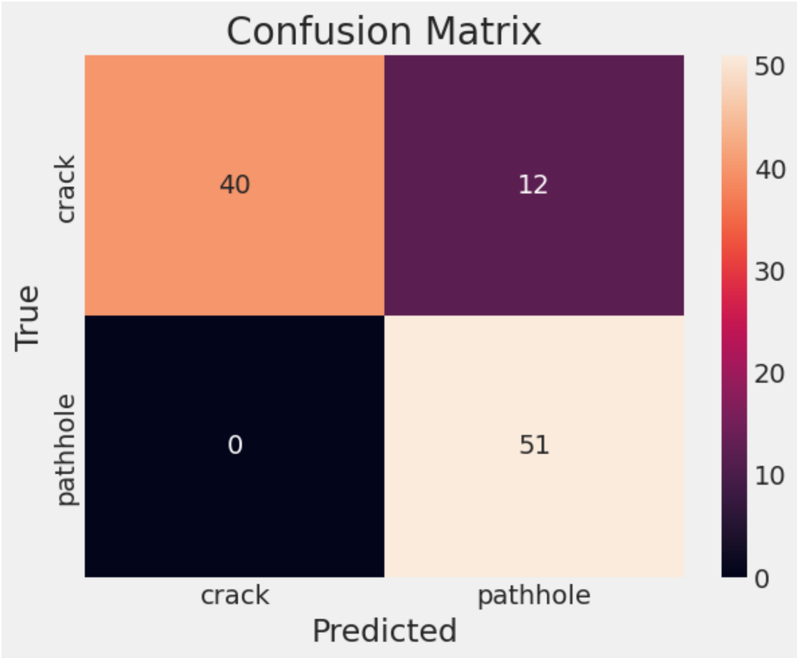

Figure 4.9 Confusion Matrix of InceptionV3

For InceptionV3, the confusion matrix demonstrates moderate classification ability. A significant number of crack images are correctly classified, but some misclassification into the pothole class is observed. Similarly, while pothole detection is relatively accurate, there are cases where potholes are wrongly labeled as cracks. This shows that InceptionV3 learns meaningful features but faces challenges in discriminating subtle differences between the two classes.

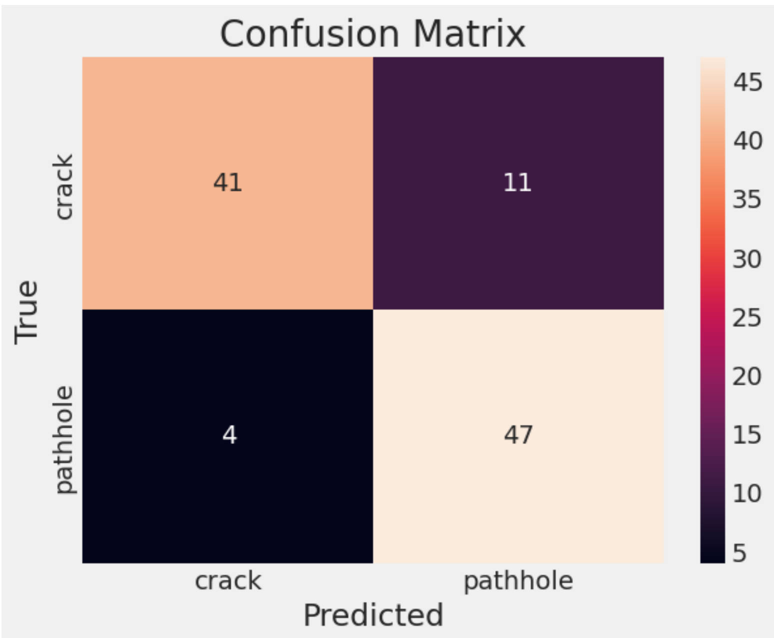

Figure 4.10 Confusion Matrix of DenseNet121

DenseNet121 confusion matrix shows good classification accuracy with the majority of crack and pothole samples accurately classified. However, there are some misclassifications in both the classes. The balanced result suggests that DenseNet121 is capable of doing well with feature reuse and gradient flow, which enables it to generalize effectively compared to deeper networks like VGG16 and VGG19.

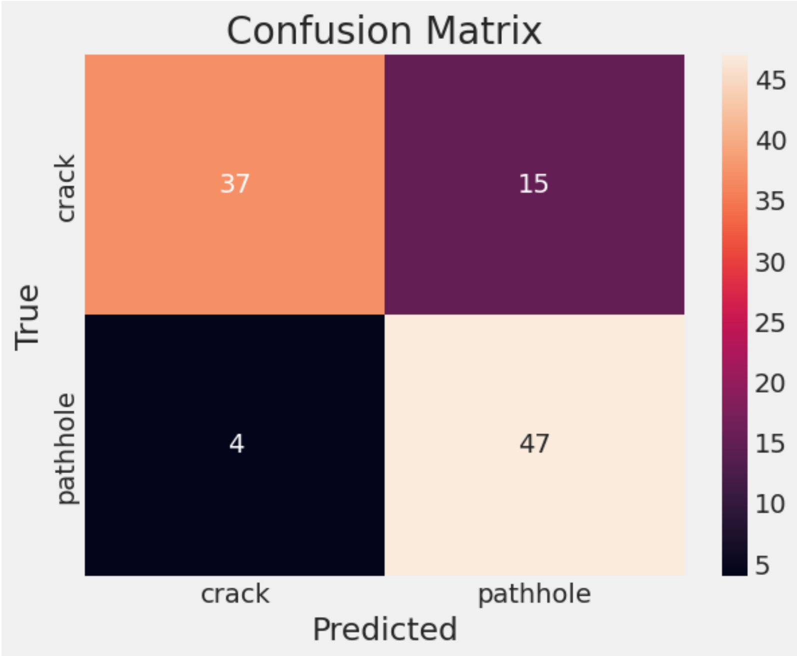

Figure 4.11 Confusion Matrix of VGG16

For the VGG16 example, confusion matrix shows very low classification ability. While the majority of the pothole and crack images are correctly classified, either direction there is a lot of misclassification. This means that VGG16 has poor feature extraction for this specific task, possibly due to the fact that it does not have advanced architectural features like residual or dense connections.

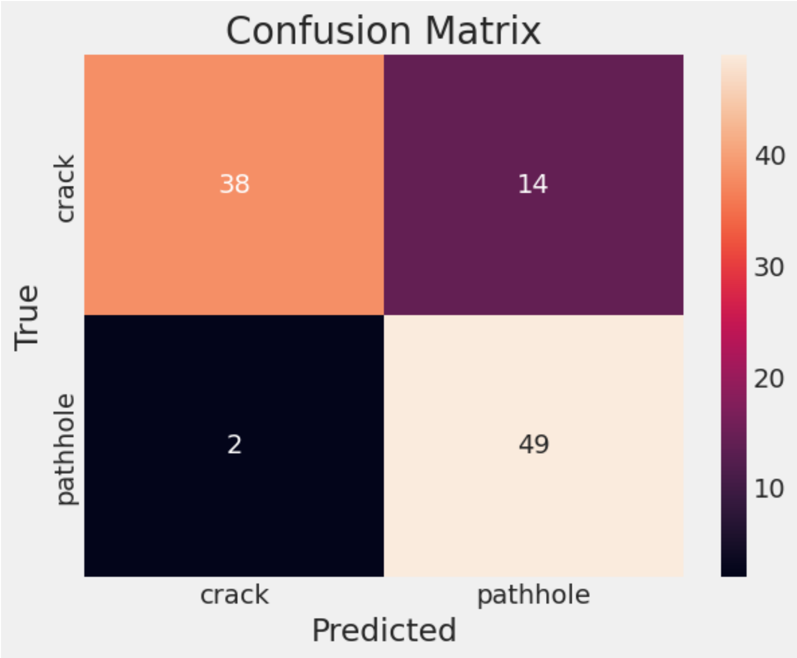

Figure 4.12 Confusion Matrix of VGG19

The confusion matrix for VGG19 also shows a similar trend to VGG16, with a noticeable level of misclassification across both categories. Although VGG19 is deeper than VGG16, the increased depth does not translate to better performance in this dataset, and it suffers from overfitting. This highlights that older architectures like VGG are less effective for fine-grained crack vs. pothole classification compared to modern models like MobileNetV2 or DenseNet121.

Comparison with Existing Works

Table 4.3: Result Comparison with Exiting Works

Ref	Model(s) Used	Best Accuracy
[1] Maeda et al. (2018)	VGG16, InceptionV3, ResNet	F1 ≈ 0.80
[3] Li et al. (2019)	DenseNet, ResNet	Acc ≈ 90%
[5] Yamashita et al. (2018)	VGG16	Acc ≈ 89%
[6] Akagić et al. (2018)	ResNet, VGG	Acc ≈ 91%
[7] Shao et al. (2020)	InceptionV3, ResNet50	Acc ≈ 94%
[8] Kim et al. (2018)	VGG, ResNet	Acc ≈ 90%
[10] Xu et al. (2019)	ResNet, Inception	IoU ≈ 0.75

Ref	Model(s) Used	Best Accuracy
[11] Cheng et al. (2020)	ResNet, VGG	Acc ≈ 92%
[15] Liu et al. (2021)	MobileNet, VGG, ResNet	F1 ≈ 0.88
[16] Gopalakrishnan (2018)	VGG, ResNet	Acc ≈ 87%
[20] Alipour et al. (2021)	Lightweight MobileNet	Acc ≈ 90%
Proposed Methodology	MobileNetV2, DenseNet121, VGG16, VGG19, InceptionV3, ResNet50	Acc 90%

4.3 Results and Discussion

Tables 4.1 and 4.2 summarize the training/validation performance and classification metrics of six pretrained CNN models—MobileNetV2, InceptionV3, ResNet50, DenseNet121, VGG16, and VGG19—used for Bangladeshi road damage classification.

From Table 4.1, MobileNetV2 achieved the highest training accuracy (95.35%) and validation accuracy (87.38%), with the lowest validation loss (0.3228), indicating effective learning and good generalization to unseen images. DenseNet121 also performed well with validation accuracy 85.44% and moderate loss (0.4212), followed by VGG16 and VGG19 with validation accuracies around 83–84%. InceptionV3 achieved moderate performance (76.70% validation accuracy), while ResNet50 performed poorly (64.08% validation accuracy) with high loss, suggesting underfitting and difficulty in learning features from the dataset.

Table 4.2 reinforces these observations by presenting Precision, Recall, F1 Score, and Accuracy. MobileNetV2 achieved the highest values across all metrics (Precision 0.91, Recall 0.90, F1 Score 0.90, Accuracy 0.90), demonstrating strong ability to correctly identify both cracks and potholes. DenseNet121 and VGG19 also performed well (F1 Scores 0.85 and 0.84), while InceptionV3 showed slightly lower recall and F1 Score (0.88), indicating minor difficulty in detecting all true damage instances. VGG16 achieved moderate performance (F1 Score 0.81), and ResNet50 had the lowest performance (F1 Score 0.67), confirming that it is less suitable for this dataset.

Combining training behavior and classification metrics, MobileNetV2 emerges as the most suitable model for Bangladeshi road damage detection, achieving both high accuracy and balanced precision-recall performance. Its lightweight architecture also makes it favorable for potential deployment in resource-constrained environments, such as mobile or embedded devices used in field inspections. DenseNet121 and VGG19 can serve as secondary options if computational resources allow, while ResNet50 is not recommended for practical use in this scenario.

4.4 Summary

The results show that pretrained CNNs can effectively classify road damage into cracks and potholes. MobileNetV2 achieved the best performance with 90% accuracy and balanced precision–recall, making it the most suitable model for this task. DenseNet121 and VGG19 performed well as alternatives, while ResNet50 showed the weakest performance. Overall, lightweight models like MobileNetV2 are highly promising for practical road monitoring in Bangladesh.

Chapter 5

Engineering Standards and DesignChallenges

5.1 Compliance with the Standards

Only mention the standards that are related to your project. This list is not complete. For each of the standards discuss the alternates with pros and cons and rationale of selection.

5.1.1 Software Standards

The required software for this project:

- Google Chrome or Microsoft Edge
- Python 3.9
- Tensor flow
- Jupiter or Google Colab

5.1.2 Hardware Standards

The required software for this project:

- Windows 11 operating system
- Hard Disk 512 GB
- * GB RAM

5.2 Impact on Society, Environment and Sustainability

5.2.1 Impact on Life

The proposed road damage classification system holds tremendous potential to improve day-to-day life and society, particularly in developing countries like Bangladesh. Roads are part of national infrastructure, and their state has a direct connection with transport, trade, safety, and the economy in general.

Accurate and automated crack and pothole identification can help authorities better prioritize work and budget, leading to quicker repair and lower long-term costs. For commuters, this means safer trips, fewer accidents caused by deteriorating road conditions, and smaller bills for vehicle repairs due to less wear and tear.

Generally speaking, improved road quality enhances transport efficiency through the reduction of delays in commuting and the transportation of goods. This benefits business operations, emergency services, and public transport systems. Moreover, an automatic system reduces the need for manual inspection, saving labor cost and enabling more consistent monitoring of large road networks.

Generally, the use of deep learning-based road damage detection can improve public safety, economic productivity, and quality of life, as well as facilitate sustainable infrastructure management in Bangladesh.

5.2.2 Impact on Society & Environment

The application of an automatic road damage classification system has widespread implications for society and the environment.

Impact on Society

Road damage in the form of potholes and cracks is one of the primary reasons for traffic accidents, injuries, and damage to vehicles. Through early and correct detection facilitated by the proposed system, road safety can be promoted, reducing accident rates along with their associated healthcare costs. For citizens, improved and better-maintained roads mean quicker travel, reduced stress, and better access to such basic services as education, health, and markets. At a societal level, good road maintenance also supports economic activity, supply chains, and the overall quality of public infrastructure.

Impact on Environment

Bad road conditions typically lead to increased fuel usage and emissions, as cars slow down, idle in traffic, or take longer routes. By enabling timely road maintenance, the system being considered contributes to lower fuel usage and reduced greenhouse gas emissions, supporting environmental sustainability. Furthermore, automatic monitoring obviates the need for several manual inspections with additional vehicles and personnel, thereby reducing the carbon footprint of maintenance operations.

Ultimately, the integration of deep learning in road monitoring promotes sustainable urban development, in alignment with smart city and global initiatives for environmentally friendly infrastructure management.

5.2.3 Ethical Aspects

The implementation and design of an automated road damage classification system must consider several ethical issues in order to foster equity, transparency, and effective use of technology.

First, the collection of road images should be respectful of privacy and data protection regulations. Images used in dataset preparation should not have recognizable personal data such as faces, vehicle number plates, or personal belongings. If these data are accidentally collected, they need to be anonymized or erased from the dataset to comply with ethical and legal standards.

Second, the system must be equitable and inclusive. Since the dataset for this study is customized for Bangladeshi roads, it creates bias when directly applied to other regions with different road conditions. To avoid unjust outcomes, models must be retrained or fine-tuned from diversified datasets before mass deployment.

Third, the reliance of public infrastructure on AI must be balanced with accountability and transparency. The authorities running the system must ensure that the system outputs are never used as irrefutable decisions but as decision-support tools to guide maintenance planning. Human intervention is required to prevent errors or misclassification from leading to delays in critical roadworks.

Finally, there are also ethical benefits of having such a system in place. Safer roads reduce accidents and injuries, promoting society's well-being, while proper maintenance prevents wasteful expenditure of resources, making it sustainable. Generally, one should tackle ethical concerns over data collection, equity, simplicity, and safe deployment so that society is positively affected by the system without inflicting unmitigated harm..

5.2.4 Sustainability Plan

Sustainability of the recommended road damage categorization system is crucial in terms of making it effective and affect over the long term. The plan focuses on continuous dataset expansion, whereby recent photos of roads from different areas, environments, and climatic conditions are constantly being incorporated in a bid to increase model flexibility and generalizability. Regular retraining and tuning of the models, especially lean architectures like MobileNetV2, will keep the system updated with changing road conditions at the same time maintaining computational efficiency. Deployment-wise, the system is designed to be run on low-end hardware, keeping energy consumption and cost to a minimum, making it accessible to government agencies and municipalities without the hefty infrastructure expenditure requirement. Integration into road management systems will allow for automated damage reports to be generated, and authorities can assign repair priority effectively, minimizing cost and time of manual surveys. On a broader scale, crack and pothole detection and repair early on will minimize vehicle damage, accidents, and loss of fuel, leading to less carbon emissions and an eco-friendlier environment. This is in line with world sustainability targets while ensuring better, more secure, and efficient transport for society. Lastly, the plan on sustainability demands balance between technological advancement, environmentally friendly action, and well-being of society to ensure the system proposed is effective, expandable, and pertinent in the long term.

5.3 Project Management and Financial Analysis

The cost table is given below:

Table 5.1: Estimated Cost for this Project

SN	Components	Estimated Cost (BDT)
1	Visiting Stakeholders	2500-3000
2	Software and Tools	1500-2000
3	Data Collection and Processing	4500-5000
4	Documentation and Report Writing	1500-2000
5	Miscellaneous (e.g., licenses, tools)	1000-1500
Total Estimated Cost	11000-13500	

5.4 Complex Engineering Problem

5.4.1 Complex Problem Solving

Table 5.2: Mapping with Complex Engineering Problem.

EP1 Dept of Knowledge	EP2 Range Of Conflicting Requirements	EP3 Depth of Analysis	EP4 Familiarity of Issues	EP5 Extent of Applicable Codes	EP6 Extent Of Stake-holder Involvement	EP7 Interdependence
✓	✓	✓	✓			✓

This project demonstrates EP1 by achieving learning outcomes K3, K4, K5, K6, and K8 through displaying the ability to apply engineering knowledge, problem-solving, analytical, and technical competencies in developing a road damage classification system based on deep learning.

This project fulfills EP2 because it states the problems in road damage classification and detection, including the ineffectiveness of manual inspection and the challenge of applying machine learning in diverse road environments. Through the comparison of different pretrained models, this project fulfills and offers suggestions on how to enhance solutions to road maintenance automation.

This project addresses EP3 by meticulously comparing experimental results across a series of pretrained CNN models, highlighting MobileNetV2 as the best solution for road damage detection among a number of promising options.

The interdisciplinarity of this project extends beyond computer science and engineering, impacting transportation safety and infrastructure management, and hence contributing to the advancement of the civil and transportation sectors, which aligns with EP4.

Finally, this project addresses EP7 as it provides a comprehensive solution that encompasses a number of aspects, including data acquisition, preprocessing, statistical modeling, and proposed methodology. This ensures a holistic solution to the complex issue of improving road safety and infrastructure longevity through the application of technology-based innovation.

Mapping with Knowledge Profile

Table 5.3: Mapping with knowledge Profile.

K1 Natural Science	K2 Mathematics	K3 Engineering Fundamentals	K4 Specialist Knowledge	K5 Engineering Design	K6 Engineering Practice	K7 Comprehension	K8 Research Literature
		✓	✓	✓	✓		✓

This project demonstrates fundamental engineering principles (K3) by employing deep learning models and data preprocessing techniques for the road damage classification task. It demonstrates specialist knowledge (K4) by applying pretrained CNN architectures such as MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19, thereby enhancing the accuracy of road damage detection, which is crucial for computer-aided infrastructure management.

The project applies engineering practice and design (K5) by systematically designing and implementing the experimental process, including dataset preparation, training, and evaluation of models. It addresses engineering practice and technology (K6) through the practical application of deep learning models in solving real-world problems, specifically automating the detection of cracks and potholes.

Finally, the project ensures K8 (Research Literature) by synthesizing insights from recent studies in computer vision and road damage detection, advancing the application of deep learning in transportation infrastructure monitoring and showcasing a comprehensive understanding of current methodologies.

5.4.2 Engineering Activities

Table 5.4: Mapping with Complex Engineering Activities.

EA1 Range of resources	EA2 Level of Interaction	EA3 Innovation	EA4 Consequences for society and environment	EA5 Familiarity
✓			✓	✓

Our project utilizes diverse resources such as high-performance computing infrastructure, GPUs, deep learning frameworks, annotated road damage datasets, and ethical considerations to ensure systematic research and contribute to advancements in automated infrastructure monitoring.

This project contributes to society by improving road safety and transportation efficiency through accurate detection of cracks and potholes, while also promoting environmental sustainability by employing resource-efficient computational models and supporting reduced fuel consumption and emissions through timely road maintenance.

Furthermore, this project expands upon existing research by examining a novel application of deep learning in Bangladeshi road conditions, demonstrated through comparative analysis of multiple pretrained CNN models. The findings provide new insights into the effectiveness of lightweight architectures such as MobileNetV2 for practical deployment in real-world road damage classification.

5.5 Summary

Deep learning model applications in road damage classification present a groundbreaking opportunity to improve road safety and management of infrastructure. By detecting cracks and potholes early and accurately, the models enable the authorities to take action promptly, with more efficient maintenance and resource deployment. Not only does the method avoid potential accident hazards and damage to vehicles but also promote fair distribution of repair resources by focusing on risk areas. Besides, the system can optimize operational efficiency through automation of inspection and decision-making, leading to huge cost savings for governments and municipalities. In order to make implementation sustainable, the models must be updated and retrained constantly and ethics are adhered to in data privacy, fairness, and transparency in decisions. Collaboration with road management authorities and stakeholders is necessary in order to maintain trust and ensure applicability in real-world applications. Also, the solution of possible issues like deployment cost and scalability will realize optimal advantage from the technology and minimize risks, ultimately contributing to the overall goal of safer, more efficient, and sustainable transportation infrastructure.

Chapter 6

Conclusion

6.1 Summary

This research sought to develop a deep learning-driven road damage classification system for Bangladeshi roads, specifically focusing on two common forms of damage: cracks and potholes. The research used a collection of 1,026 images, of which 490 were images of cracks and 536 were images of potholes. Six widely used pretrained convolutional neural networks—MobileNetV2, ResNet50, InceptionV3, DenseNet121, VGG16, and VGG19—were used and trained for 15 epochs at a learning rate of 0.001. Experimental outcome revealed that MobileNetV2 executed the best with an accuracy of 90%, which is greater than other models for precision, recall, and F1-score.

The findings validate the suitability of lightweight architectures such as MobileNetV2 for the detection of road damage, particularly in outdoor conditions where computational resources are not readily available. The study further demonstrates the potential of deep learning techniques to automate the assessment of road conditions to some degree, reducing the necessity for human scanning and enabling faster, less expensive planning of maintenance.

Besides technical advancements, this study has environmental and societal implications. Timely and correct road distress detection can provide road safety, reduce vehicle maintenance, and reduce accident risks. From an environmental perspective, smooth and well-maintained roads cut down on wasteful fuel consumption and reduce greenhouse gases.

In short, the system described here has promise as a realistic tool for field use by road maintenance divisions of Bangladesh and similar countries. Such future work would include extending the dataset with larger and more diverse images, adding more classes of road deterioration, and exploring real-time capability through mobile phones or drone-based systems. These innovations would render the system even more scalable, accurate, and influential, with an eye toward safer roads and sustainable infrastructure construction.

6.2 Limitation

Although the proposed deep learning-based road damage classification system shows promising results, there are certain limitations that must be noted. Firstly, the dataset used in this study was relatively small and comprised only 1,026 images of two classes (potholes and cracks). A small dataset may restrict the model's ability to generalize to diverse road conditions, different lighting conditions, and weather conditions such as rain, shadows, or dust. Secondly, the data was collected solely from Bangladeshi roads, and hence the model can possibly perform worse when applied to predict road conditions elsewhere in other countries or regions with different construction material and damage types. Thirdly, the system is currently set to function simply on binary classification (crack or pothole), whereas actual road conditions include an array of damages such as rutting, surface wear, and block cracks. Lastly, training and testing were done offline, and real-time deployment was not explored, which might prove to be problematic when deployed on a large scale in the real world.

6.3 Future Work

For future study, there are several improvements that can be made. Augmenting the dataset with more images from different regions, weather, and lighting conditions will make the model more robust and generalizable. Incorporating more road damage classes will make the system more comprehensive and useful for real-world road maintenance planning. Furthermore, incorporating real-time detection capability using mobile apps, drones, or IoT devices could significantly contribute to the system's usefulness. Employing more advanced deep learning architectures, such as Vision Transformers (ViT) or CNN–RNN hybrid models, can also yield improved results. Finally, collaboration with government agencies for field testing and deployment may help ensure the system's effectiveness in various realistic scenarios and contribute towards smart city endeavors for sustainable infrastructure management.

References

[1] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiyama, and H. Omata, "Road damage detection and classification using deep neural networks with images captured through a smartphone," Computer-Aided Civil and Infrastructure Engineering, vol. 33, Art. no. 12, 2018.

[2] R. Fan, H. Wang, P. Wang, and M. Liu, "Road damage detection based on deep convolutional neural networks," IEEE Access, vol. 8, pp. 128141–128149, 2020.

[3] Z. Li, Q. Wang, and H. Luo, "Automatic detection and classification of pavement cracks based on improved deep learning model," in IEEE, 2019, pp. 906–911.

[4] M. Zhao, X. Xu, Y. Liu, and J. Zhang, "Deep learning for road damage detection and classification: A comprehensive survey," IEEE Transactions on Intelligent Transportation Systems, vol. 22, Art. no. 7, 2021.

[5] R. Yamashita, M. Nishio, R. K. Do, and K. Togashi, "Deep convolutional neural network-based crack detection for road pavement," Journal of Computing in Civil Engineering, vol. 32, Art. no. 5, 2018.

[6] A. Akagić, A. Mulahušić, E. Buza, and S. Omanović, "Pavement crack detection using deep learning," International Journal of Pavement Research and Technology, vol. 11, Art. no. 6, 2018.

[7] Y. Shao, Z. Han, W. Wang, and M. Lu, "Deep learning-based automatic detection of road damage using smartphone images," Automation in Construction, vol. 113, p. 103134, 2020.

[8] B. Kim, S. Cho, and M. Kang, "Pavement crack detection using deep convolutional neural networks," in IEEE, 2018, pp. 1–4.

[9] Y. Gao, J. Zhang, J. Wang, and P. Liu, "Road crack detection using deep convolutional neural networks," Automation in Construction, vol. 98, pp. 247–254, 2019.

[10] X. Xu, Z. Zhang, Z. Wang, and J. Zhang, "Automatic pavement crack detection and segmentation using deep learning," International Journal of Pavement Engineering, vol. 20, Art. no. 5, 2019.

[11] G. Cheng, P. Wang, and P. Xu, "Deep learning-based automated pavement crack detection and segmentation using UAV imagery," Remote Sensing, vol. 12, Art. no. 23, 2020.

[12] Q. Zhang, S. Liu, and Y. Hu, "Road damage detection and classification using YOLO-based deep learning models," Sensors, vol. 21, Art. no. 14, 2021.

[13] T. Tang, Z. Deng, and Z. Huang, "Pavement crack detection using hybrid deep learning models," Construction and Building Materials, vol. 262, p. 120781, 2020.

[14] W. Cai, F. Wang, and Z. Li, "A review of pavement distress detection using deep learning," Construction and Building Materials, vol. 247, p. 118580, 2020.

[15] H. Liu, J. Wu, and K. Zhang, "Deep learning in pavement condition assessment: Recent advances and prospects," Journal of Infrastructure Systems, vol. 27, Art. no. 3, 2021.

[16] K. Gopalakrishnan, "Deep learning in data-driven pavement condition assessment," Automation in Construction, vol. 94, pp. 272–286, 2018.

[17] Y. Zhou, X. Xu, and J. Zhang, "Multi-scale CNN for pavement crack detection," Automation in Construction, vol. 103, pp. 45–56, 2019.

[18] Y. Sun, H. Chen, and R. Wang, "Transfer learning for pavement crack detection using pre-trained CNNs," Construction and Building Materials, vol. 235, p. 117529, 2020.

[19] Y. Zhou, J. Zhang, and X. Xu, "Generative adversarial networks for data augmentation in pavement crack detection," Automation in Construction, vol. 118, p. 103265, 2020.

[20] M. Alipour, E. Rahimian, and A. Taghipour, "Lightweight CNN for pothole detection on mobile devices," IEEE Access, vol. 9, pp. 12345–12356, 2021.

[21] Yashvi Chandola, J. Virmani, H.S. Bhadauria, and P. Kumar, "Lightweight end-to-end Pre-trained CNN-based computer-aided classification system design for chest radiographs," Elsevier eBooks, pp. 167–183, Jan. 2021, doi: https://doi.org/10.1016/b978-0-323-90184-0.00001-1.

[22] P. Potrimba, "What is ResNet-50?," Roboflow Blog, Mar. 13, 2024. https://blog.roboflow.com/what-is-resnet-50/

[23] A. H. Md. Linkon, Md. M. Labib, T. Hasan, M. Hossain, and M.-E. - Jannat, "Deep learning in prostate cancer diagnosis and Gleason grading in histopathology images: An extensive study," Informatics in Medicine Unlocked, vol. 24, p. 100582, 2021, doi: https://doi.org/10.1016/j.imu.2021.100582.

[24] iamtapendu, "Introduction to DenseNet-121," Kaggle.com, May 11, 2025. https://www.kaggle.com/code/iamtapendu/introduction-to-densenet-121

[25] Great Learning, "Everything you need to know about VGG16," Medium, Sep. 23, 2021. https://medium.com/@mygreatlearning/everything-you-need-to-know-about-vgg16-7315defb5918

[26] "VGG-19 Convolutional Neural Network - an overview | ScienceDirect Topics," www.sciencedirect.com. https://www.sciencedirect.com/topics/computer-science/vgg-19-convolutional-neural-network

24% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Detection Groups

- **36 AI-generated only 24%**
  Likely AI-generated text from a large-language model.
- **0 AI-generated text that was AI-paraphrased 0%**
  Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer
Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?
The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?
Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.

turnitin Page [illegible] of [illegible] - AI Writing Overview Submission ID trn:oid:::[illegible]

Mahmudul Thesis (SSK).docx

ORIGINALITY REPORT

18%	11%	13%	8%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	arxiv.org Internet Source	1%
2	Submitted to Asia e University Student Paper	1%
3	5dok.org Internet Source	1%
4	Tian, Zuoyu. "Language Change in Ancient Chinese: A Computational Approach", Indiana University Publication	1%
5	Submitted to Dublin Business School Student Paper	1%
6	repository.uel.ac.uk Internet Source	1%
7	Submitted to Middlesex University Student Paper	1%
8	Khumukcham Robindro Singh, Nazrul Hoque, Arnab Kumar Maji, Sabyasachi Mondal et al. "Emerging Trends and Future Directions in Artificial Intelligence, Machine Learning, and Internet of Things Innovations - A proceeding of NEIAIS — 2025", CRC Press, 2025 Publication	1%
9	Submitted to University of Finance - Marketing Student Paper	1%