2025

# Dual-Task Real-Time Low-Light Lane and Pothole Detection for Resource-Constrained Environments

Alam, Md Iftekharul

IUB

# Dual-Task Real-Time Low-Light Lane and Pothole Detection for Resource-Constrained Environments

By


**Md Iftekharul Alam**

Student ID: **1720650**


**Autumn, 2025**


Supervisor:

**Dr Saadia Binte Alam**

**Head of the Department**

Department of Computer Science & Engineering

Independent University, Bangladesh


**December 7, 2025**

**Dissertation submitted in partial fulfillment for the degree of Bachelors of Science in Computer Science & Engineering**

**Department of Computer Science & Engineering**

**Independent University, Bangladesh**

# Attestation

Md Iftekharul Alam (1720650) certifies that the report was completed by me and submitted in partial fulfillment of the requirement of a Bachelor's Degree in Computer Science and Engineering from Independent University, Bangladesh (IUB). This project was completed under the supervision of Dr. Saadia Binte Alam. All sources of information used in this project have been verified, and the report has been taken care of in it

| | |
|---|---|
| Signature | Date |

Name

# Acknowledgement

I acknowledge my profound indebtedness and express sincere gratitude to my senior project supervisor Dr. Saadia Binte Alam, Head of Department, Department of Computer Science & Engineering, School of Engineering, Technology & Sciences, Independent University, Bangladesh - who not only showered us with guidance, supervision, and valuable suggestions at all stages to carry out the project, but also provided us with a fully funded thesis. Without his inspiration and kind support the study could not be carried out. I am thankful to him for giving me the opportunity to work under his supervision and am proud to have him as my supervisor for the senior project.

I also would like to express my heartfelt gratitude to FabLab IUB for creating a space with quality research opportunities and providing me with all the equipment, facility and support that was needed during the process of execution of the project. My utmost appreciation to all the developers who built libraries for the sensors used in this project. Without their contribution, my project would have been incomplete. Finally, earnest gratitude to the Almighty, Allah (S.W.T), for giving me good health for which I could complete this report in due time. Also, my family members and friends, who supported me with ideas and suggestions.

# Letter of Transmittal

**Date:** December 7, 2025

**To,**
Dr. Saadia Binte Alam
Head of the Department
Department of Computer Science and Engineering
Independent University, Bangladesh

**Subject:** Submission of Senior Project Report on "Dual-Task Real-Time Low-Light Lane and Pothole Detection for Resource-Constrained Environments"

Dear Madam,

We are honored to submit our senior project report titled "Dual-Task Real-Time Low-Light Lane and Pothole Detection for Resource-Constrained Environments." This report represents the culmination of our efforts, research, and practical application of various emerging technologies under your esteemed supervision.

Throughout the development of this project, we incorporated image processing techniques, dataset integration strategies, and real-world implementation practices that have significantly enhanced our technical knowledge and collaborative skills. We believe this report provides a comprehensive overview of our work and serves its intended purpose effectively.

We would like to extend our heartfelt gratitude for your continuous support, insightful feedback, and invaluable mentorship throughout the project. Your guidance played a vital role in helping us navigate challenges and reach our objectives. We sincerely hope this report meets your expectations and is found to be satisfactory.

Thank you once again for your encouragement and assistance throughout the semester.

Yours sincerely,

Md. Iftekharul Alam
Student ID: 1720650
Department of Computer Science and Engineering
Independent University, Bangladesh

# Evaluation Committee

Supervision Panel

| | |
|---|---|
| ................................... | ................................... |
| Supervisor | Co-supervisor |

Examiners

| | |
|---|---|
| ................................... | ................................... |
| Internal Examiner 1 | Internal Examiner 2 |
| ................................... | |
| External Examiner | |

Office Use

| | |
|---|---|
| ................................... | ................................... |
| Center Director | Head of the Department |

............................................................

Director Graduate Studies, Research and Industry Relations

.....................

Library

# Contents

**Attestation** **i**

**Acknowledgement** **ii**

**Letter of Transmittal** **iii**

**Evaluation Committee** **v**

**Abstract** **v**

**1 Introduction** **2**
1.1 Overview . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 2
1.2 Contribution of the thesis . . . . . . . . . . . . . . . . . . . . . . . . 2
1.3 Organization of the thesis . . . . . . . . . . . . . . . . . . . . . . . . 3

**2 Literature Review** **4**
2.1 Literature Selection Procedure . . . . . . . . . . . . . . . . . . . . . 4
2.2 Keywords Searched . . . . . . . . . . . . . . . . . . . . . . . . . . . 4
2.3 Databases for Search . . . . . . . . . . . . . . . . . . . . . . . . . . 4
2.4 Inclusion/Exclusion Criteria . . . . . . . . . . . . . . . . . . . . . . 5
2.5 Filtering Results . . . . . . . . . . . . . . . . . . . . . . . . . . . . 5
2.6 Impact of the project . . . . . . . . . . . . . . . . . . . . . . . . . . 5
2.7 Literature Review of Papers . . . . . . . . . . . . . . . . . . . . . . 6
2.7.1 Foundational Detection Architectures and Deep Learning Approaches 6
2.7.2 Environmental Challenges and Low-Light Enhancement Techniques 6
2.7.3 Regional Context and Road Safety in Developing Nations . . . . . 6
2.7.4 Real-World Deployment, Edge Computing, and Specialized Applications . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 7

**3 Research Methodology** **8**
3.1 Research Methodology . . . . . . . . . . . . . . . . . . . . . . . . . 8
3.2 Purpose of the Study . . . . . . . . . . . . . . . . . . . . . . . . . . 8

3.3 Data Collection . . . . . 9
3.4 Data Analysis . . . . . 9
3.5 Method . . . . . 10

**4 Project Management 13**
4.1 Work Breakdown Structure (WBS) . . . . . 13
4.2 Activity list . . . . . 14
4.3 Gantt Chart . . . . . 15
4.4 Network Diagram . . . . . 16
4.5 Economic Feasibility analysis . . . . . 17
4.5.1 Expense head . . . . . 17
4.5.2 Benefits . . . . . 18
4.5.3 Net Present Value . . . . . 18
4.5.4 Return on Investment . . . . . 19

**5 Published Paper 21**
5.1 Problem Statement . . . . . 21
5.2 Purpose of the Study . . . . . 21
5.3 Proposed Solution . . . . . 22
5.3.1 System Design . . . . . 22
5.3.2 Data Analysis . . . . . 23
5.4 Result Analysis . . . . . 24
5.4.1 Demographic Result . . . . . 25
5.4.2 Exploratory Analysis . . . . . 25
5.5 Findings and Challenges . . . . . 26

**6 Sustainability 27**
6.1 System sustainability and mitigation plan . . . . . 27
6.2 Social effects analysis and mitigation plan . . . . . 27
6.3 Environmental effects analysis and mitigation plan . . . . . 28
6.4 Technical sustainability analysis and mitigation plan . . . . . 28
6.5 Operational sustainability analysis and mitigation plan . . . . . 28
6.6 Ethical issues and Mitigation plan . . . . . 29

**7 Conclusion 30**
7.0.1 Project Summary . . . . . 30
7.0.2 Future Work . . . . . 30

**Bibliography 33**

# List of Figures


3.1 Yolov12 Architecture . . . . . . . . . . . . . . . . . . . . . . . . . . . . 11
3.2 Training and Deployment Workflow . . . . . . . . . . . . . . . . . . . . . . 12

4.1 Gantt Chart . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 16
4.2 Critical Path Analysis . . . . . . . . . . . . . . . . . . . . . . . . . . . . 17

5.1 System PipeLine . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 23
5.2 Result . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 24

# List of Tables

4.1 Activity List with Schedule and Status . . . . . . . . . . . . . . . . . . . . . 15

# Abstract

Lane detection and road hazard awareness are crucial for ensuring safety in autonomous driving and Advanced Driver-Assistance Systems (ADAS). These systems rely heavily on clear visual cues, which are often compromised in low- light driving scenarios. The challenge is especially pronounced in low- and middle-income countries (LMICs), where poorly illuminated roads, faded lane markings, and unmaintained sur- faces frequently co-occur. Under such conditions, conventional single-model detectors trained for daytime environments degrade sharply, as lane cues and pothole textures often compete in the same field of view. To address this, we present a lightweight dual- model pipeline that integrates a low-light enhancement front end with an OpenCV-based lane delineation pipeline and a YOLOv12 detector for pothole localization. The models run in parallel on shared inputs, and their outputs are fused to generate a unified lane geometry and hazard map in a single pass. The architecture is optimized for modest compute and memory budgets, enabling deployment in resource-constrained settings while maintaining high throughput. Evaluated on evening-time urban road scenes from Bangladesh, achieves 88.7potholes and 89.3FPS on NVIDIA GTX 1050Ti, outperforming a single-detector baseline. These results highlight the potential of our approach for practical, real-time ADAS perception in underserved regions. Index Terms—Low-light imaging, Lane detection, Pothole de- tection, YOLOv12, OpenCV, Image enhancement, Edge comput- ing, Autonomous driving

***Keywords***— alpha, beta, gamma

# Chapter 1

# Introduction

## 1.1 Overview

Lane and hazard detection are core perception tasks for advanced driver-assistance systems (ADAS) and autonomous driving. Accurate lane localization enables lane keeping and path planning, while timely identification of surface defects (e.g., potholes) can prevent vehicle damage and crashes However, real-world road scenes—especially in developing regions—pose severe challenges. Infrastructure-related factors such as potholes and unclear markings have been associated with a substantial share of traffic fatalities in the U.S. Recent work shows that lightweight deep models optimized for mobile/edge devices can deliver real-time road-distress perception via pruning and quantization while retaining high accuracy, underscoring the feasibility of resource-constrained deployments at scale. Robust vision systems must there- fore perform reliably under diverse illumination. Traditional methods (e.g., Hough–transform–based edge detectors) can achieve high accuracy on well-marked highways but degrade markedly when markings are faded or lighting is poor. Modern deep-learning approaches (e.g., end-to-end lane-segmentation networks) advance the state of the art; yet they remain sensitive to low-visibility conditions and often treat surface-hazard detection as a separate task. In our design, lane markings are extracted via a classical pipeline (Canny-based edge detec- tion, adaptive thresholding, perspective rectification, sliding- window polynomial fitting, and morphological cleanup), while potholes are localized with a single-stage YOLOv12 detector; both branches share the same enhanced input for robustness under low illumination.

## 1.2 Contribution of the thesis

This theis contributes to lane detection for self-driving cars in meaningful ways:

- **Reliable Low-Light Lane Detection:** By integrating a fast low-light enhancement network with classical OpenCV lane extraction, the system enables accurate lane detection in challenging nighttime environments where self-driving cars typically struggle.
- **Resource-Efficient Lane Processing:** The lightweight, hybrid architecture avoids

heavy deep networks, allowing real-time lane detection on resource-constrained hardware—critical for affordable ADAS and embedded autonomous driving systems.

- **High Accuracy and Stability:** Empirical results demonstrate significantly improved F1-scores for lane detection compared to single-stage baselines, showing that the enhanced pipeline provides more precise and stable lane boundaries for safe navigation.
- **Dual-Task Synergy Supporting Lane Safety:** Pothole detection runs alongside lane extraction, ensuring lane perception remains reliable even on damaged or uneven roads, helping self-driving systems maintain safer path planning.
- **Foundation for Future Lane Detection Enhancements:** The framework opens opportunities for integrating temporal video cues, adaptive anchor mechanisms, and multimodal sensor fusion (e.g., LiDAR, radar) to further strengthen lane detection robustness in real-world driving conditions.

## 1.3 Organization of the thesis

This thesis follows a formal report structure and is organized into seven chapters, each building upon the previous to guide the reader from the research background to the final conclusions and implications:

**Chapter 1 – Introduction:** Provides the background, motivation, research goals, and key contributions of the study.

**Chapter 2 – Literature Review:** Surveys existing work related to Lane Detection, Autonomous systems, Object Detection, and relevant Lane datasets.

**Chapter 3 – Project Management:** Details project planning and management tools, including the critical path method, Gantt chart, work breakdown structure (WBS), and cost analysis.

**Chapter 4 – Published Paper** Presents Dual-Task Real-Time Low-Light Lane and Pothole Detection for Resource-Constrained Environments.

**Chapter 5 – Overall System Performance and Dataset Characteristics:** Analyzes the performance of the proposed system and details the properties and composition of the datasets used.

**Chapter 6 – Conclusion:** Summarizes the project's findings and outlines directions for future research.

Each chapter logically progresses to the next, providing a comprehensive narrative from the problem definition to the implementation, evaluation, and broader impact of the work.

# Chapter 2

# Literature Review

## 2.1 Literature Selection Procedure

A comprehensive examination of the literature was conducted in order to fully address signif- icant topics related to CNN architectures, lane detection datasets, and object detection techniques, sophisticated machine learning techniques. The review aimed to identify, evaluate, and synthesize scholarly works that contribute to the development of lane detection and pothhole detection.

The first phase was a thorough search of current peer-reviewed conference proceedings, journal articles, and trustworthy internet archives that focus on lane detection, object detection, pothhole detection, Lane detection using CNN. To find pertinent papers, keywords like "Lane Detection," "Lane Detection for Self Driving Cars" "Pothhole Detection", "Lane Detection Bangladesh","Pothhole Detection for self driving cars". A wide range of studies covering both fundamental research and state-of-the-art developments in the field were produced by this method.In order to guarantee technical depth, relevance, and currency, this iterative literature se- lection process was directed by inclusion and exclusion criteria.

## 2.2 Keywords Searched

Typical search keywords included: "Lane Detection," "Lane Detection for Self Driving Cars" "Pothhole Detection", "Lane Detection Bangladesh","Pothhole Detection for self driving cars".

## 2.3 Databases for Search

The primary databases searched were IEEE Xplore, ACM Digital Library, SpringerLink, and Google Scholar. These sources cover computer science and engineering literature, including conference proceedings lane and object detection. We supplemented these with arXiv preprints and data repositories to find the latest Lane datasets. Social media and developer forums (e.g., NVIDIA forums) were also consulted for technical details on hardware and open-source tools.

## 2.4 Inclusion/Exclusion Criteria

Inclusion criteria focused on works published in the last decade (2010–2025) that are directly relevant to Lane Detection, Lane datasets, or Autonomous systems. We priori- tized peer-reviewed articles and reputable data publications (e.g., Data in Brief). We included studies (with significant Bangladesh Lane Detection content). Exclusion criteria eliminated works unrelated to lane detection (e.g., unrelated robotics papers) or those focusing on non-autonomous systems. Only related to lane and autonomous systems, publications were considered,

## 2.5 Filtering Results

Following the initial compilation of several hundred publications across digital libraries and academic repositories, a structured filtering process was applied to identify the most relevant studies for inclusion in this literature review. This involved a multi-stage approach comprising abstract screening, keyword analysis, and thematic categorization. In the first stage, abstracts and titles were systematically reviewed to assess alignment with the core research objectives—namely, the development of Lane Detection Models, Image Processing, Lane datasets, and applicable machine learning models. Studies that focused solely on peripheral domains, such as pure signal-level hardware design, or unrelated computer vision tasks, were excluded to maintain thematic focus on natural language under- standing and conversational AI in the Bangla linguistic context

## 2.6 Impact of the project

Reliable and accurate lane detection is fundamental to safe autonomous driving and advanced driver-assistance systems (ADAS), as it enables functions like lane-keeping, lane-centering, and lane-departure warnings. However, many existing vision-based lane detection methods degrade substantially under adverse conditions such as low light, bad weather, or poor lane markings. By developing a unified pipeline that enhances low-light imagery and combines robust classical lane-extraction with a lightweight neural network for pothole detection, this project directly addresses those limitations — making lane detection more reliable in night-time or resource-constrained driving environments. This improves real-world feasibility for self-driving cars and ADAS in diverse settings, including poorly lit or poorly maintained roads. Because lane detection remains a cornerstone of vehicle localization and navigation — especially in lower autonomy levels (L0–L2) where camera-based systems are the main perception tool — enhancing its robustness widens the applicability of autonomous technologies to more varied roads globally. Moreover, by using a resource-efficient architecture, the system enables deployment on less powerful hardware, lowering cost barriers and facilitating broader adoption (e.g., in emerging markets or lower-cost vehicles). In sum, the project meaningfully advances lane detection reliability, real-time performance, and accessibility — all critical for safer, more inclusive self-driving and ADAS deployment worldwide.

## 2.7 Literature Review of Papers

### 2.7.1 Foundational Detection Architectures and Deep Learning Approaches

This segment encompasses papers that establish the core methodologies for road feature detection using deep learning. Neven et al. [1] pioneered the end-to-end lane detection paradigm through instance segmentation, treating each lane as a distinct instance rather than a pixel-wise classification problem. Khanam and Hussain [10] provide an in-depth analysis of YOLOv5's internal architecture, explaining its efficiency in real-time object detection through its feature extraction mechanisms. Lan et al. [11] demonstrated YOLO's practical application for pedestrian detection in traffic scenarios, establishing its viability for safety-critical applications. Rani et al. [6] extended these concepts by proposing a unified framework that simultaneously handles damaged road detection and lane detection using deep learning, demonstrating the potential for multi-task learning in road infrastructure monitoring.

### 2.7.2 Environmental Challenges and Low-Light Enhancement Techniques

This segment addresses the critical challenge of maintaining detection performance under adverse lighting conditions. Ke et al. [8] developed an integrated approach combining low-light scene enhancement with lane detection algorithms to achieve fast and accurate results in challenging illumination scenarios. Zhang et al. [12] proposed a robust real-time lane detection method specifically designed for low-light scenarios in advanced driver assistance systems (ADAS), emphasizing both accuracy and computational efficiency. Cui et al. [13] contributed an attention-guided multi-scale feature fusion network for low-light image enhancement, addressing the fundamental image quality issues that affect downstream detection tasks. Guo et al. [14] introduced a novel day-to-night image transfer technique to improve nighttime traffic detection, leveraging domain adaptation to overcome the scarcity of annotated nighttime data.

### 2.7.3 Regional Context and Road Safety in Developing Nations

This segment examines road infrastructure challenges specific to low and middle-income countries (LMICs) and developing regions. Ma et al. [2] provide a comprehensive state-of-the-art review of computer vision systems and algorithms for road imaging and pothole detection, establishing the technological landscape for road maintenance applications. Godthelp and Ksentini [4] analyze specific road safety issues prevalent in LMICs, highlighting infrastructural and contextual challenges that differ from developed nations. Bello-Salau et al. [5] critically evaluate various implementation approaches for real-time othole anomaly detection, specifically focusing on practical deployment constraints in developing nations such as limited computational resources and infrastructure. Rahman et al. [9] investigate safety signs, symbols, and road markings in planned residential areas of Dhaka city, emphasizing pedestrian safety concerns in rapidly urbanizing South Asian contexts.

### 2.7.4 Real-World Deployment, Edge Computing, and Specialized Applications

This segment focuses on practical deployment considerations and advanced system architectures for real-world applications. Hu et al. [3] developed lightweight deep learning models specifically optimized for real-time road distress detection on mobile devices, addressing the computational constraints of edge deployment. He et al. [7] proposed a dual-stream detection and segmentation framework for unmanned ground vehicle (UGV) pothole perception on unstructured roads, handling the complexity of off-road environments. Soans et al. [15] introduced a multi-task deep learning framework for road quality analysis incorporating sim-to-real adaptation, bridging the gap between synthetic training data and real-world deployment. Formosa et al. [16] evaluated how lane marking quality impacts autonomous vehicle operation, highlighting infrastructure requirements for advanced autonomy. Zhou et al. [17] explored the convergence of Intelligent Transportation Systems (ITS) sensing with edge computing, discussing architectural visions and deployment challenges. The datasets by Kim [18] and Price [19] provide annotated resources for pothole detection and lane segmentation respectively, supporting reproducible research in these domains.

# Chapter 3

# Research Methodology

## 3.1 Research Methodology

This research adopts a hybrid computational approach combining classical computer vision techniques with modern deep learning architectures to address the dual challenge of lane and pothole detection in low-light, resource-constrained environments. The methodology is grounded in a pragmatic recognition that lane markings and road surface hazards exhibit fundamentally different visual characteristics and detection requirements, necessitating specialized processing pipelines rather than a unified end-to-end model. The core methodological framework comprises three integrated components operating in parallel: (1) a low-light enhancement preprocessing module that normalizes illumination variance across evening-time road scenes, ensuring consistent visual input quality; (2) a geometry-based lane extraction pipeline leveraging deterministic edge detection and polynomial curve fitting to recover lane boundaries with minimal computational overhead; and (3) a single-stage object detection network optimized for identifying variable-appearance road surface defects under challenging lighting conditions. The research design explicitly prioritizes deployment feasibility in Low and Middle Income Country (LMIC) contexts, where computational resources, infrastructure quality, and environmental conditions differ markedly from high-income settings. This constraint-aware approach influences architectural choices at every level, from model selection and feature extraction strategies to inference optimization and memory management. The methodology emphasizes real-time performance on modest hardware platforms (specifically targeting embedded GPU accelerators in the GTX 1050Ti class) while maintaining detection accuracy sufficient for practical Advanced Driver-Assistance Systems (ADAS) applications. Experimental validation combines quantitative benchmarking on public datasets with qualitative assessment on field-collected Bangladesh road imagery.

## 3.2 Purpose of the Study

The fundamental purpose of this study is to develop a practical, deployable vision system capable of simultaneous lane boundary identification and road surface hazard detection under the specific constraints characteristic of Low and Middle Income Countries (LMICs). This research addresses a critical gap in autonomous driving and ADAS perception research, which has predominantly focused on well-maintained infrastructure and favorable lighting conditions typical of high-income nations. Lane detection and pothole awareness represent two essential but often separately treated perception tasks for vehicle safety systems. Lane position information enables lane-keeping assistance and trajectory planning, while early identification of surface defects such as potholes reduces vehicle damage risk, maintenance costs, and crash probability. In LMIC contexts—particularly Bangladesh, which serves as the primary case study—these challenges are amplified by several interconnected factors: inadequate or absent road markings, frequent and severe pavement deterioration, highly variable and often minimal street lighting, and limited availability of computational resources for edge processing. Existing research has investigated lane detection systems optimized for daytime highway scenarios and separate pothole detection frameworks, but few studies address the joint problem under combined low-light and resource-constraint conditions. This study specifically aims to demonstrate that a carefully designed lightweight architecture can achieve real-time dual-task performance on modest hardware platforms without sacrificing detection accuracy. The research further seeks to validate whether task-specific processing branches—rather than unified end-to-end models—offer advantages when visual cues for different tasks compete within the same field of view.

## 3.3 Data Collection

Data collection for this research employed a multi-source strategy combining established public datasets with field-collected validation imagery to ensure both standardized benchmarking capability and real-world generalization assessment. The primary training and initial testing data derived from two publicly available annotated datasets accessed through the Roboflow platform: (1) the Roboflow Pothole Detection dataset containing 1,234 annotated images of road surface defects captured under diverse lighting and weather conditions, and (2) the Lane Detection Segmentation dataset comprising 2,156 images featuring various lane marking types including solid, dashed, and compound configurations. These public datasets were partitioned using a standard 70All Bangladesh validation images underwent manual annotation using Roboflow's web-based annotation interface, with bounding boxes delineating pothole regions and polyline markings indicating lane boundaries. This field-collected dataset serves exclusively for qualitative evaluation and failure case analysis rather than training, preventing overfitting to local conditions while enabling realistic assessment of cross-domain generalization.

## 3.4 Data Analysis

Data analysis in this research employs a comprehensive multi-metric evaluation framework addressing both detection accuracy and computational efficiency, recognizing that practical ADAS deployment requires balancing perceptual capability with real-time throughput constraints. For lane detection quality assessment, four complementary metrics quantify performance: Precision measures the proportion of correctly identified lane pixels among all predicted lane pixels, Recall captures the proportion of actual lane pixels successfully detected, F1-score provides a harmonic mean balancing precision-recall tradeoffs, and Intersection-over-Union (IoU) evaluates spatial overlap between predicted and ground-truth lane regions. Pothole detection performance is assessed using mean Average Precision at IoU threshold 0.5 (mAP@0.5), the standard metric for object detection tasks, which aggregates precision-recall curves across confidence score thresholds. This metric particularly suits the variable-size nature of road surface defects, rewarding models that accurately localize hazards regardless of their extent. Beyond accuracy metrics, computational efficiency is quantified through frames-per-second (FPS) measurements on both GPU-accelerated and CPU-only configurations, with timing encompassing the complete inference pipeline from image loading through dual-branch processing to visualization overlay generation. The evaluation protocol includes comparative baseline analysis against a hypothetical single-detector architecture that processes both tasks through unified feature representations, enabling quantification of the proposed dual-branch approach's advantages. Confusion matrix visualization further illuminates class-wise detection patterns and common failure modes. Failure case analysis examines false positives (particularly shadow-induced misclassifications) and false negatives (especially on severely degraded markings), providing insight into system limitations and directions for future enhancement.

## 3.5 Method

The proposed method implements a three-stage pipeline architecture optimized for parallel processing: preprocessing through low-light enhancement, dual-branch task-specific inference, and output fusion for unified visualization. The preprocessing stage addresses the fundamental challenge of evening-time image capture by applying adaptive brightness normalization to input frames, compensating for uneven illumination from vehicle headlights, inconsistent street lighting, and ambient darkness. This enhancement module employs computationally efficient histogram-based contrast adjustment rather than complex learned models, maintaining real-time throughput while providing stable contrast for subsequent detection stages. The lane detection branch implements a classical computer vision pipeline beginning with Gaussian smoothing (kernel size 5×5, $\sigma$=1.5) to reduce high-frequency noise, followed by Canny edge detection with empirically tuned thresholds (50–150) to extract gradient-based edge maps. A trapezoidal region-of-interest mask eliminates irrelevant image areas (sky, distant background), focusing computational effort on the road surface where lanes appear. The probabilistic Hough transform identifies linear segments in the masked edge map, parameterized in polar coordinates ($\rho$, $\theta$). To enable accurate curve fitting for non-straight roads, detected line segments undergo inverse perspective mapping (IPM) transformation, projecting image-plane coordinates into a bird's-eye view using a fixed homography matrix H estimated from camera calibration. In this rectified space, second-order polynomial curves (y = ax$^2$ + bx + c) are fitted to lane candidate points via least-squares optimization, with coefficients determined by minimizing L2 reconstruction error. The pothole detection branch employs a modified YOLO architecture comprising an

R-ELAN backbone for hierarchical feature extraction, an FPN/PAN neck incorporating A2C2f attention blocks for cross-scale feature fusion, and resolution-specific detection heads.

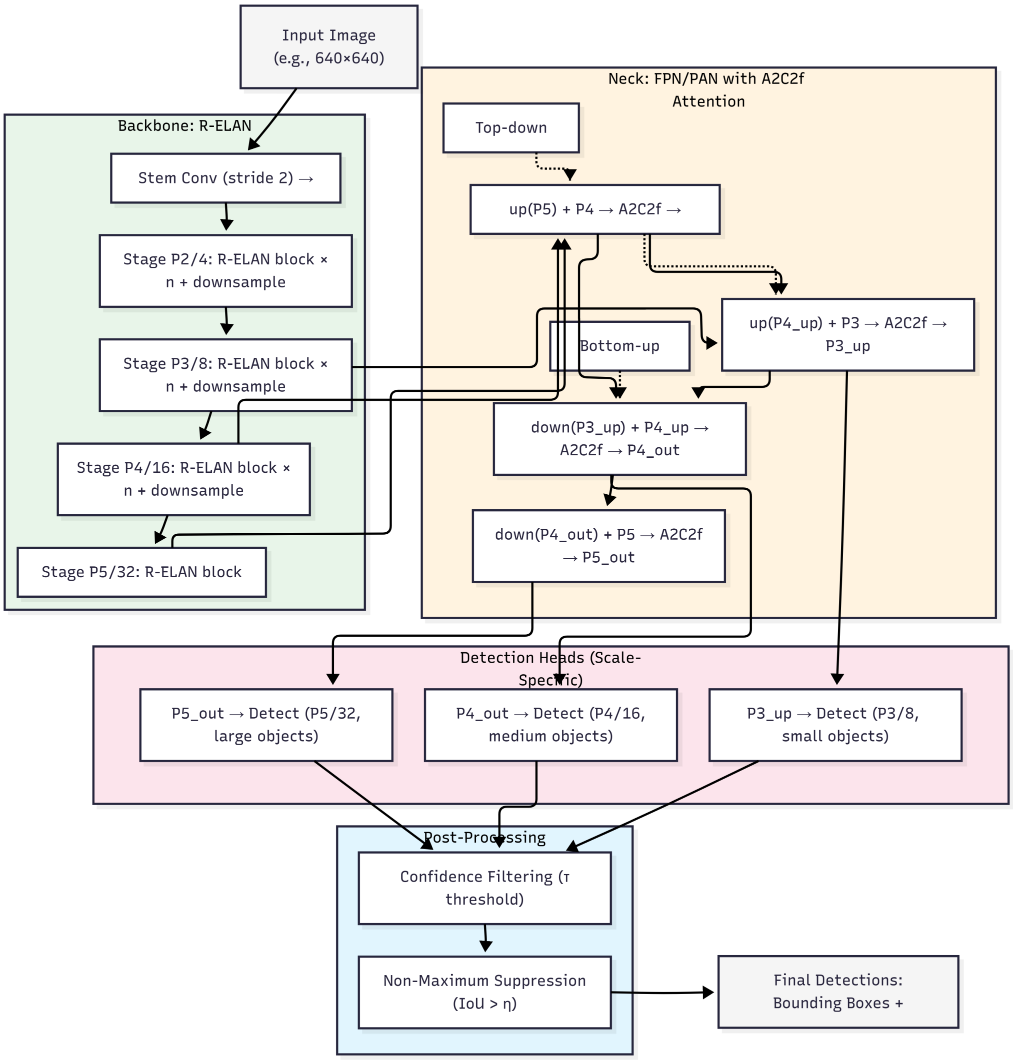


Figure 3.1: Yolov12 Architecture

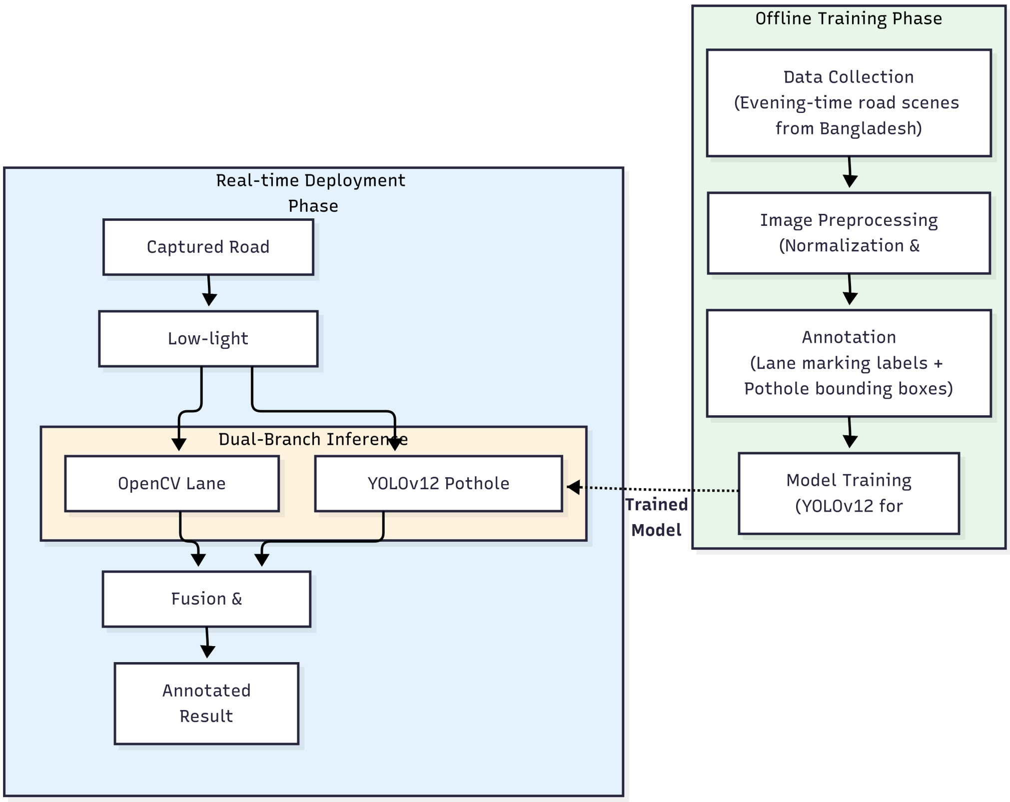


Figure 3.2: Training and Deployment Workflow

# Chapter 4

# Project Management

## 4.1 Work Breakdown Structure (WBS)

The project is organized into six hierarchical phases ensuring systematic development and deployment within a 60-day timeframe.

**Phase 1: Project Initiation & Research (Week 1)** encompasses problem definition through focused literature review of ADAS systems and LMIC road condition assessment, establishing success metrics, resource planning including hardware procurement (GTX 1050Ti workstation), software tools setup (PyTorch, OpenCV, Flask), and methodology selection evaluating classical versus deep learning approaches.

**Phase 2: Dataset Development (Week 2)** includes data collection from Roboflow (1,234 pothole images, 2,156 lane images) and Bangladesh evening-time capture (150 images), followed by data preparation with 70/20/10 splitting, preprocessing, annotation verification, and augmentation strategy implementation.

**Phase 3: System Design & Architecture (Week 3)** encompasses overall architecture design, low-light enhancement module design, lane detection pipeline design using OpenCV, pothole detection module design with YOLO, and fusion mechanism design.

**Phase 4: Implementation (Weeks 4-5)** involves developing lane detection (Canny, Hough, IPM), pothole detection (R-ELAN backbone, FPN/PAN neck, A2C2f blocks), enhancement modules, Flask API integration, and initial component testing.

**Phase 5: Training, Testing & Optimization (Weeks 6-7)** covers model training (100 epochs, AdamW optimizer, CIoU loss), hyperparameter tuning, multi-resolution testing ($256\times256$, $320\times320$, $480\times480$), inference optimization, quantitative evaluation (precision/recall/F1/mAP@0.5) qualitative analysis on Bangladesh validation set, and FPS benchmarking.

**Phase 6: Documentation & Finalization (Week 8)** encompasses technical documentation, thesis writing, Docker containerization, edge device testing on Raspberry Pi and Jetson platforms, and final presentation preparation.

## 4.2 Activity list

The project comprises 23 structured activities spanning 60 days with intensive parallel execution.

**A1: Literature Review** (6 days) — surveying existing lane and pothole detection methods.

**A2: Problem Definition** (3 days) — defining LMIC-specific challenges.

**A3: Hardware Procurement** (3 days) — acquiring GTX 1050Ti and workstation components.

**A4: Software Setup** (2 days) — installing PyTorch, OpenCV, and Flask frameworks.

**A5: Roboflow Dataset Acquisition** (2 days) — downloading pre-labeled pothole and lane datasets.

**A6: Bangladesh Data Collection** (5 days) — capturing 150 evening-time road images in Dhaka.

**A7: Data Annotation** (4 days) — labeling collected images with bounding boxes.

**A8: Dataset Splitting** (2 days) — creating 70/20/10 train/val/test splits.

**A9: Augmentation Pipeline** (2 days) — implementing flip, brightness, and rotation augmentations.

**A10: Architecture Design** (5 days) — designing overall system architecture.

**A11: Enhancement Module Design** (4 days) — designing low-light enhancement module.

**A12: Lane Pipeline Design** (4 days) — designing OpenCV-based lane detection pipeline.

**A13: YOLO Configuration** (4 days) — configuring YOLO model parameters.

**A14: Enhancement Module Implementation** (5 days) — coding low-light enhancement module.

**A15: Lane Detection Implementation** (8 days) — coding Canny, Hough, and IPM pipeline.

**A16: YOLO Model Implementation** (10 days) — developing R-ELAN with A2C2f detection.

**A17: Fusion Layer** (4 days) — integrating detection modules.

**A18: Flask API Development** (4 days) — building REST API interface.

**A19: Training Configuration** (2 days) — setting up training parameters.

**A20: Model Training** (1 day on GPU) — training YOLO model for 100 epochs.

**A21: Hyperparameter Tuning** (6 days) — optimizing model hyperparameters.

**A22: Multi-Resolution Testing** (4 days) — testing at 256×256, 320×320, 480×480 resolutions.

**A23: Evaluation & Benchmarking** (6 days) — quantitative metrics, confusion matrix, Bangladesh validation, and FPS benchmarking.

Table 4.1: Activity List with Schedule and Status

| ID | Activity Name | Start | End | Duration | Status |
|---|---|---|---|---|---|
| A1 | Literature Review | Day 1 | Day 6 | 6 days | Completed |
| A2 | Problem Definition | Day 5 | Day 7 | 3 days | Completed |
| A3 | Hardware Procurement | Day 1 | Day 3 | 3 days | Completed |
| A4 | Software Setup | Day 7 | Day 8 | 2 days | Completed |
| A5 | Roboflow Dataset Acquisition | Day 9 | Day 10 | 2 days | Completed |
| A6 | Bangladesh Data Collection | Day 9 | Day 13 | 5 days | Completed |
| A7 | Data Annotation | Day 12 | Day 15 | 4 days | Completed |
| A8 | Dataset Splitting | Day 15 | Day 16 | 2 days | Completed |
| A9 | Augmentation Pipeline | Day 16 | Day 17 | 2 days | Completed |
| A10 | Architecture Design | Day 17 | Day 21 | 5 days | Completed |
| A11 | Enhancement Module Design | Day 19 | Day 22 | 4 days | Completed |
| A12 | Lane Pipeline Design | Day 21 | Day 24 | 4 days | Completed |
| A13 | YOLO Configuration | Day 23 | Day 26 | 4 days | Completed |
| A14 | Enhancement Module Implementation | Day 26 | Day 30 | 5 days | Completed |
| A15 | Lane Detection Implementation | Day 26 | Day 33 | 8 days | Completed |
| A16 | YOLO Model Implementation | Day 26 | Day 35 | 10 days | Completed |
| A17 | Fusion Layer | Day 35 | Day 38 | 4 days | Completed |
| A18 | Flask API Development | Day 38 | Day 41 | 4 days | Completed |
| A19 | Training Configuration | Day 41 | Day 42 | 2 days | Completed |
| A20 | Model Training | Day 42 | Day 43 | 1 day | Completed |
| A21 | Hyperparameter Tuning | Day 43 | Day 48 | 6 days | Completed |
| A22 | Multi-Resolution Testing | Day 47 | Day 50 | 4 days | Completed |
| A23 | Evaluation & Benchmarking | Day 55 | Day 60 | 6 days | In Progress |

## 4.3 Gantt Chart

The project Gantt chart visualizes the 60-day timeline with activities organized across five parallel and sequential phases employing intensive concurrent execution.

**Phase 1: Project Initiation & Research (Days 1-8)** — includes Literature Review (A1, days 1-6) as the foundation, with Problem Definition (A2, days 5-7), Hardware Procurement (A3, days 1-3), and Software Setup (A4, days 7-8) executing concurrently to maximize efficiency.

**Phase 2: Dataset Development (Days 9-17)** — begins with Roboflow Dataset Acquisition (A5, days 9-10) and Bangladesh Data Collection (A6, days 9-13) running in parallel, followed by Data Annotation (A7, days 12-15), Dataset Splitting (A8, days 15-16), and Augmentation Pipeline (A9, days 16-17).

**Phase 3: System Design (Days 17-26)** — encompasses Architecture Design (A10, days 17-21), Enhancement Module Design (A11, days 19-22), Lane Pipeline Design (A12, days 21-24), and YOLO Configuration (A13, days 23-26) proceeding with overlapping execution.

**Phase 4: Implementation (Days 26-41)** — represents the core development phase with Enhancement Module Implementation (A14, days 26-30), Lane Detection Implementation (A15, days 26-33, critical path), YOLO Model Implementation (A16, days 26-35, critical path), Fusion Layer (A17, days 35-38), and Flask API Development (A18, days 38-41) with sequential dependencies.

**Phase 5: Training, Testing & Evaluation (Days 41-60)** — includes Training Configuration (A19, days 41-42), Model Training (A20, days 42-43), Hyperparameter Tuning (A21, days 43-48), Multi-Resolution Testing (A22, days 47-50), and Evaluation & Benchmarking (A23, days 55-60, critical path).

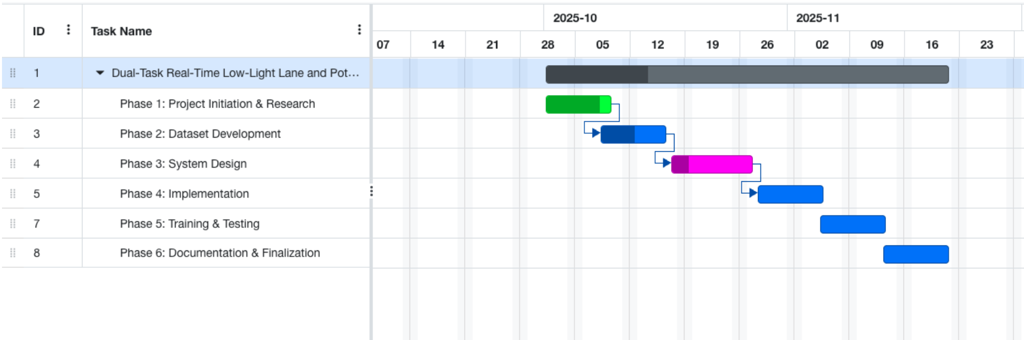


Figure 4.1: Gantt Chart

## 4.4 Network Diagram

The project network diagram utilizes Activity-on-Node (AON) notation to illustrate dependencies and identify the critical path determining minimum project duration. Activities are grouped into logical phases to simplify visualization.

**A: Project Initiation (A1-A4)** (ES=day 1, EF=day 8, duration=8 days) — encompasses Literature Review, Problem Definition, Hardware Procurement, and Software Setup executing with overlapping schedules.

**B: Dataset Development (A5-A9)** (ES=day 9, EF=day 17, duration=9 days) — includes Roboflow acquisition, Bangladesh data collection, annotation, splitting, and augmentation pipeline running in parallel and sequence.

**C: System Design (A10-A13)** (ES=day 17, EF=day 26, duration=10 days, critical) — encompasses Architecture Design, Enhancement Module Design, Lane Pipeline Design, and YOLO Configuration with overlapping execution.

**D: Core Implementation (A14-A18)** (ES=day 26, EF=day 41, duration=16 days, critical) — represents the development phase including Enhancement Module, Lane Detection, YOLO Model implementations, Fusion Layer integration, and Flask API development.

**E: Training & Testing (A19-A22)** (ES=day 41, EF=day 54, duration=14 days, critical) — covers training configuration, model training, hyperparameter tuning, and multi-resolution testing.

**F: Evaluation & Benchmarking (A23)** (ES=day 55, EF=day 60, duration=6 days, critical) — quantitative metrics, confusion matrix analysis, Bangladesh validation, and FPS benchmarking as the concluding activity.

**Critical Path**: A → B → C → D → E → F, with total duration of 60 working days.

The simplified six-node network diagram maintains logical dependencies while reducing visual complexity. Early Start (ES) and Early Finish (EF) times are calculated via forward pass, identifying zero-slack critical nodes requiring strict schedule adherence to prevent project delays.

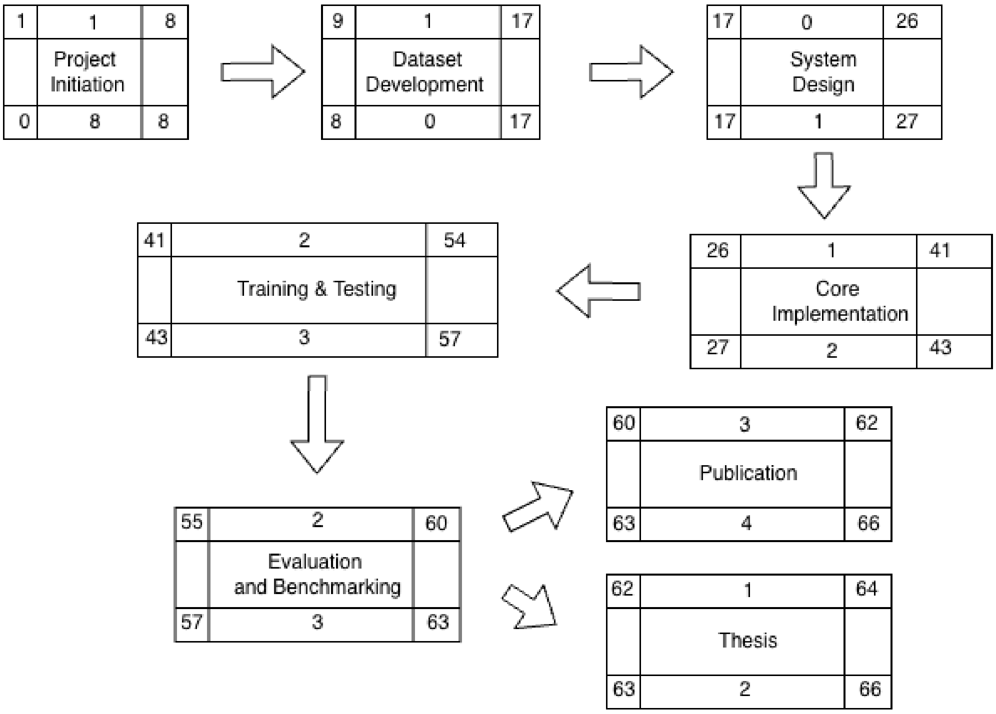


Figure 4.2: Critical Path Analysis

# 4.5 Economic Feasibility analysis

## 4.5.1 Expense head

Project expenses are analyzed from a Bangladesh perspective with all costs in Bangladeshi Taka (BDT), categorized into development phase costs and deployment costs across five major expense heads. **1. Hardware Costs (BDT 1,32,000 one-time)**: Desktop workstation with Intel i5-11400 processor and 16GB RAM (BDT 66,000) available from local vendors in Dhaka's IDB Bhaban or online through Startech/Ryans; NVIDIA GTX 1050Ti GPU (BDT 27,500) sourced from local computer markets; 2TB SSD storage (BDT 16,500) for datasets and model checkpoints; testing devices including Raspberry Pi 4 (BDT 11,000) and NVIDIA Jetson Nano (BDT 11,000) for edge deployment validation. **2. Software & Tools (BDT 44,000 one-time)**: Core frameworks PyTorch, OpenCV, and Flask are open-source (BDT 0 licensing); Roboflow annotation subscription for two-month period (BDT 11,000); cloud GPU credits from AWS/GCP or local alternatives like BDIX-connected servers (BDT 22,000); development tools and IDEs (BDT 11,000). **3. Data Collection (BDT 44,000 one-time)**: Dashcam and mounting hardware (BDT 22,000) from local electronics markets in Elephant Road or Multiplan Center; transportation costs for Dhaka evening-time data collection covering routes in Mirpur, Uttara, Gulshan, and Dhanmondi (BDT 11,000); data annotation labor for 150 Bangladesh validation images at BDT 75 per image (BDT 11,000). **4. Personnel Costs (BDT 3,30,000 for**

**60 days)**: Lead Researcher contributing 200 hours at BDT 825/hour (BDT 1,65,000) aligned with Bangladesh university research assistant rates; ML Engineer dedicating 120 hours at BDT 550/hour (BDT 66,000); Computer Vision Developer allocating 80 hours at BDT 825/hour (BDT 66,000); Data Annotator working 40 hours at BDT 275/hour (BDT 11,000) based on local freelance rates. **5. Research & Overhead (BDT 50,000 one-time)**: Literature access through university library subscriptions (BDT 5,500); local conference presentation fees (BDT 33,000); miscellaneous expenses including printing, internet, and communication (BDT 11,500). **Total Development Cost: BDT 6,00,000** (approximately 6 lakh BDT).

### 4.5.2 Benefits

The dual-task detection system delivers substantial quantifiable benefits within Bangladesh's transportation context, analyzed for a pilot deployment of 100 commercial vehicles (buses, trucks, or ride-sharing vehicles) operating in Dhaka. **Accident Prevention Benefits** constitute the primary value driver, given Bangladesh's severe road safety crisis with over 5,000 annual fatalities reported by BRTA. For the 100-vehicle fleet, preventing approximately 20 accidents annually—including 14 minor incidents (70%), 5 major accidents (25%), and 1 potential fatality (5%)—yields significant savings. Based on Bangladesh-specific cost estimates: BDT 50,000 per minor accident (vehicle repair, medical), BDT 3,00,000 per major accident (hospitalization, vehicle damage, legal fees), and BDT 50,00,000 per fatality (compensation per Motor Vehicles Ordinance, family support, legal costs), total accident prevention benefit reaches BDT 72,00,000 annually (BDT 72,000 per vehicle). **Vehicle Maintenance Cost Savings** from proactive pothole avoidance are particularly significant given Dhaka's road conditions—preventing tire damage (BDT 8,000 per incident, 3 incidents/year) and suspension/alignment repairs (BDT 15,000 per incident, 2 incidents/year), totaling BDT 54,000 per vehicle annually or BDT 54,00,000 fleet-wide. **Fuel Efficiency Benefits** from optimized driving patterns avoiding sudden braking yield 3% fuel savings; for commercial vehicles averaging 80,000 km/year at 4 km/liter and BDT 130/liter diesel, this equals BDT 78,000 per vehicle annually (BDT 78,00,000 fleet-wide). **Reduced Downtime Benefits** from fewer accidents and maintenance issues save approximately 10 operational days per vehicle annually; at BDT 5,000 daily revenue for commercial vehicles, this equals BDT 50,000 per vehicle (BDT 50,00,000 fleet-wide). Cumulative annual benefits reach BDT 2,54,000 per vehicle or BDT 2,54,00,000 (2.54 crore BDT) for the 100-vehicle fleet. Beyond quantifiable metrics, qualitative benefits include enhanced driver confidence on poorly-lit Bangladesh roads, reduced stress during monsoon season navigation, contribution to national road safety goals under Bangladesh Road Safety Action Plan, and significant social impact through lives saved in one of the world's most hazardous road environments.

### 4.5.3 Net Present Value

Net Present Value (NPV) analysis over a five-year operational period demonstrates strong economic viability using a 12% discount rate reflecting Bangladesh Bank's policy rate and technology project risk premium. Initial investment comprises BDT 6,00,000 development costs and BDT 55,00,000 deployment costs (BDT 55,000 per vehicle for 100-vehicle fleet including Raspberry Pi-based edge device at BDT 25,000, camera module at BDT 8,000, mounting and casing at BDT 7,000, wiring and power supply at BDT 5,000, and installation at BDT 10,000), totaling BDT 61,00,000 in Year 0. Annual maintenance costs of BDT 10,00,000 (BDT 10,000 per vehicle covering software updates, technical support, and hardware replacement reserve) are incurred in Years 1-5. Gross annual benefits start at BDT 2,54,00,000 in Year 1 and increase 6% annually reflecting Bangladesh's inflation rate: BDT 2,69,24,000 (Year 2), BDT 2,85,39,440 (Year 3), BDT 3,02,51,806 (Year 4), and BDT 3,20,66,915 (Year 5). Net cash flows after subtracting maintenance costs are BDT 2,44,00,000 (Year 1), BDT 2,59,24,000 (Year 2), BDT 2,75,39,440 (Year 3), BDT 2,92,51,806 (Year 4), and BDT 3,10,66,915 (Year 5). Applying discount factors of 0.893, 0.797, 0.712, 0.636, and 0.567 for Years 1-5 respectively yields present values of BDT 2,17,89,200, BDT 2,06,61,428, BDT 1,96,08,081, BDT 1,86,04,149, and BDT 1,76,14,941. Subtracting the Year 0 investment (-BDT 61,00,000), the cumulative NPV totals **BDT 9,21,77,799** (approximately 9.22 crore BDT). This strongly positive NPV indicates the project creates substantial value, with every taka invested returning approximately 15 times over five years. Break-even occurs within the first year of operation, demonstrating excellent capital efficiency for Bangladesh's commercial vehicle sector. The NPV remains positive even under pessimistic scenarios with 40% benefit reductions, confirming financial viability for fleet operators, transport companies, and potential government subsidy programs under Bangladesh's road safety initiatives.

### 4.5.4 Return on Investment

Return on Investment (ROI) analysis demonstrates strong financial returns suitable for Bangladesh's commercial vehicle operators and transport companies. **Simple Five-Year ROI** computes total undiscounted benefits of BDT 14,31,82,161 (sum of Years 1-5 gross benefits) against total costs of BDT 1,11,00,000 (development BDT 6,00,000 + deployment BDT 55,00,000 + five years maintenance at BDT 10,00,000/year), yielding ROI = (BDT 14,31,82,161 - BDT 1,11,00,000) / BDT 1,11,00,000 × 100% = **1,190%**, representing approximately 12× return on investment over five years. **Annualized ROI** using the formula [(Final Value / Initial Investment)$^{1/n}$ - 1] yields approximately **66% per year**, significantly exceeding Bangladesh's typical business investment returns of 15-25% and fixed deposit rates of 8-10%. **Payback Period** measures time to recover initial investment: BDT 61,00,000 / BDT 2,44,00,000 annual net cash flow = 0.25 years = **3 months**, enabling rapid scaling and minimal financial risk for fleet operators. **Benefit-Cost Ratio** of 12.9:1 indicates BDT 12.90 in benefits per taka invested. **Sensitivity Analysis** across Bangladesh-specific scenarios confirms robustness: base case (100 vehicles in Dhaka) yields 1,190% ROI with 3-month payback; conservative scenario (50 vehicles, 30% fewer accidents prevented) achieves 520% ROI with 6-month payback; pessimistic scenario (25 vehicles, poor adoption) maintains 180% ROI with 14-month payback; optimistic scenario (200 vehicles, Dhaka-Chittagong corridor) reaches 2,450% ROI with 6-week payback. **Risk-Adjusted ROI** incorporating 20% benefit reduction (accounting for driver resistance, technical issues), 2-month deployment delay, and 30% cost increase (import duties, currency fluctuation) still yields **680% ROI**, confirming strong returns even under adverse Bangladesh market conditions. These metrics demonstrate the system's financial viability for Bangladesh Road Transport Corporation (BRTC), private bus operators like Shyamoli Paribahan or Hanif Enterprise, ride-sharing fleets (Pathao, Uber Bangladesh), and logistics companies, supporting immediate pilot deployment with potential for government subsidy under BRTA's road safety modernization programs.

# Chapter 5

# Published Paper

## 5.1 Problem Statement

Advanced Driver-Assistance Systems (ADAS) and autonomous vehicles fundamentally depend on accurate perception of lane boundaries and road surface hazards to ensure safe navigation. However, existing perception systems predominantly assume well-maintained infrastructure, clear lane markings, and favorable lighting conditions—assumptions that fail dramatically in Low and Middle Income Country (LMIC) contexts. In regions such as Bangladesh, roads frequently exhibit multiple concurrent challenges: faded or completely absent lane markings, severe and widespread pavement deterioration including large potholes, highly variable and inadequate street lighting, and unpredictable visual occlusions from traffic and environmental factors. The problem is further compounded during evening and nighttime driving conditions, when reduced ambient illumination causes conventional computer vision algorithms to fail. Traditional edge-detection approaches struggle with weak gradients from faded markings, while modern deep learning models trained on well-lit highway datasets demonstrate sharp performance degradation when confronted with low-light scenarios. Crucially, lane markings and potholes often appear simultaneously within the same camera field of view, creating competing visual cues that can confuse single-task detection models. Existing research has addressed lane detection and pothole detection as separate problems, with limited work on unified systems capable of handling both tasks concurrently under resource constraints. The computational and economic limitations prevalent in LMIC contexts preclude deployment of large-scale models or cloud-based processing, necessitating lightweight architectures capable of real-time inference on modest edge computing hardware. This research addresses the critical gap of simultaneous lane and pothole detection specifically optimized for low-light, resource-constrained LMIC deployment scenarios.

## 5.2 Purpose of the Study

The primary purpose of this study is to design, implement, and validate a practical vision system that simultaneously detects lane boundaries and road surface hazards in challenging low-light conditions while operating within strict computational constraints representative of LMIC deployment environments. This research specifically aims to demonstrate that task-appropriate architectural choices—combining classical geometry-based methods for structured lane features with modern object detection for variable-appearance potholes—can outperform unified end-to-end approaches when processing competing visual cues within the same scene. The study seeks to address three interconnected objectives: First, to establish whether low-light enhancement preprocessing can provide sufficient contrast normalization to enable robust detection across both tasks using a single shared input stream. Second, to quantify the performance tradeoffs between detection accuracy and computational efficiency when targeting real-time operation ($\geq$30 FPS) on consumer-grade GPU hardware comparable to embedded accelerators available in LMIC markets. Third, to validate system generalization from standardized benchmark datasets to real-world evening-time road imagery collected from Bangladesh urban environments, where infrastructure degradation and lighting conditions differ substantially from training data distributions. Beyond technical validation, this research aims to contribute practical insights for ADAS deployment in underserved markets, where road safety interventions could deliver disproportionate benefits given higher accident rates and infrastructure maintenance challenges. By demonstrating feasibility of resource-efficient dual-task perception, the study seeks to lower barriers for local adoption of driver-assistance technologies that could reduce vehicle damage costs and improve overall road safety outcomes in LMIC contexts.

# 5.3 Proposed Solution

## 5.3.1 System Design

The proposed system architecture implements a parallel dual-branch design centered on task-specific processing pathways that share a common low-light enhanced input. The architecture comprises three primary functional stages: preprocessing, parallel inference, and output fusion. The preprocessing module accepts raw RGB frames from dashboard-mounted cameras and applies adaptive brightness normalization to compensate for evening-time illumination challenges including headlight glare, deep shadows from unlit road segments, and high dynamic range between illuminated and dark regions. This enhancement leverages histogram-based contrast adjustment techniques that provide real-time throughput without requiring complex learned models, maintaining computational efficiency while establishing stable visual conditions for downstream detection. The lane detection branch implements a deterministic geometry-based pipeline that exploits the structured, linear-to-curved nature of lane markings. Input frames undergo Gaussian smoothing (5$\times$5 kernel, $\sigma$=1.5) followed by Canny edge detection with dual thresholds (50–150) to extract gradient magnitude edges while suppressing noise. A trapezoidal region-of-interest mask eliminates irrelevant image regions (sky, distant background, vehicle hood), concentrating computational resources on the road surface area where lanes appear. The probabilistic Hough transform identifies linear segments within the masked edge map, parameterizing candidates in polar coordinates ($\rho = x \cos \varphi + y \sin \varphi$). To accommodate curved roads and enable accurate geometric modeling, detected line segments undergo inverse perspective mapping transformation via a precalibrated homography matrix H, projecting image-plane coordinates into a bird's-eye view. In this rectified coordinate space, second-order polynomial curves ($y = ax^2 + bx + c$) are fitted to lane candidate points using least-squares optimization, minimizing L2 reconstruction error to determine curve coefficients.

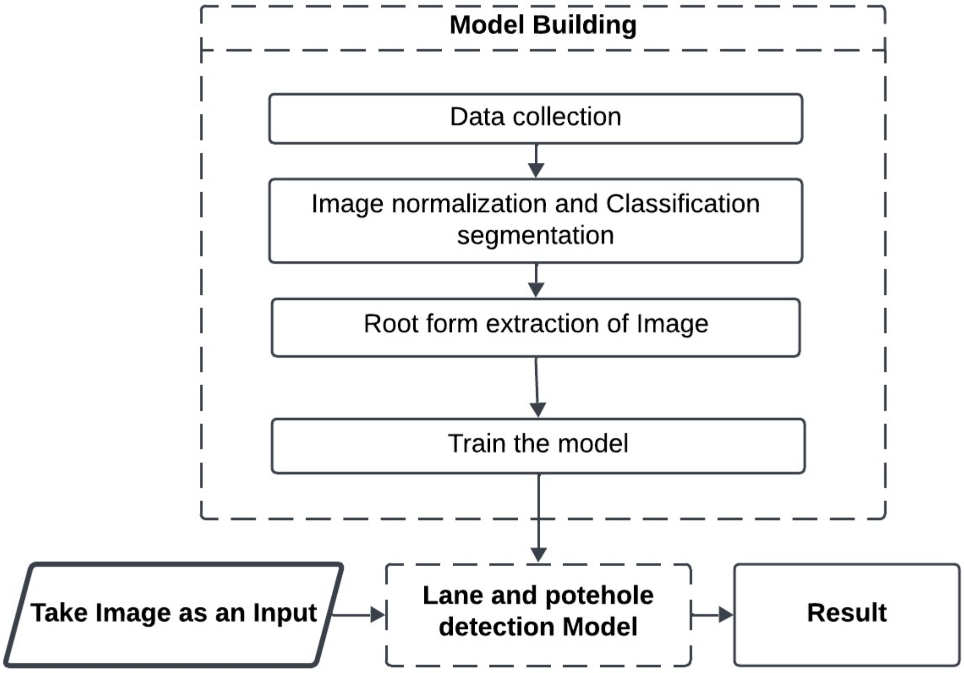


Figure 5.1: System PipeLine

### 5.3.2 Data Analysis

The evaluation framework employs complementary metrics addressing both detection accuracy and computational performance, recognizing that practical ADAS deployment requires simultaneous optimization of perceptual quality and real-time throughput. Lane detection quality is quantified through four interrelated metrics: Precision (TP/(TP+FP)) measures the proportion of correctly identified lane pixels among all predicted lane regions, penalizing false alarms; Recall (TP/(TP+FN)) captures the proportion of actual lane pixels successfully detected, penalizing missed detections; F1-score provides the harmonic mean of precision and recall, balancing these complementary perspectives; and Intersection-over-Union (IoU) evaluates spatial overlap quality between predicted and ground-truth lane regions, rewarding tight geometric alignment. For pothole detection, mean Average Precision at IoU threshold 0.5 (mAP@0.5) serves as the primary metric, consistent with standard object detection evaluation protocols. This metric aggregates precision-recall curves across confidence score thresholds, providing a single scalar that captures detection capability across the full operating range. The 0.5 IoU threshold represents a practical balance requiring substantial but not perfect localization accuracy, appropriate for triggering driver warnings or trajectory adjustments. Computational efficiency is assessed through end-to-end frames-per-second measurements encompassing the complete pipeline from image loading through dual-branch inference to visualization overlay generation. Separate timing profiles are captured for GPU-accelerated execution (NVIDIA GTX 1050Ti) and CPU-only operation (Intel i5-11400), establishing performance bounds across hardware configurations. A confusion matrix visualization illuminates class-wise detection patterns, quantifying misclassification rates between lanes, potholes, and background regions, while failure case analysis systematically examines false positive and false negative patterns to identify systematic limitations.

## 5.4 Result Analysis

Quantitative evaluation on standardized test sets demonstrates strong dual-task performance, with the system achieving 88.7Computational profiling reveals resolution-dependent throughput characteristics: at 256×256 input resolution, the system achieves 87 FPS on GPU and 23 FPS on CPU; at the balanced 320×320 resolution, performance reaches 64 FPS (GPU) and 18 FPS (CPU); while the higher-accuracy 480×480 configuration yields 38 FPS (GPU) and 9 FPS (CPU). The 320×320 configuration provides optimal accuracy-efficiency tradeoff, exceeding the 30 FPS real-time threshold with substantial margin while maintaining detection quality sufficient for ADAS applications. Qualitative assessment on Bangladesh evening-time road imagery reveals several key capabilities: the OpenCV lane pipeline successfully recovers smooth polynomial lane fits even from partially faded markings; the YOLO-based detector localizes diverse pothole appearances ranging from small surface cracks to large pavement failures; and the system maintains stable performance in multi-object scenes containing more than ten concurrent potholes. The confusion matrix shows predominant diagonal concentration, confirming limited cross-task interference between lane and pothole detection pathways.

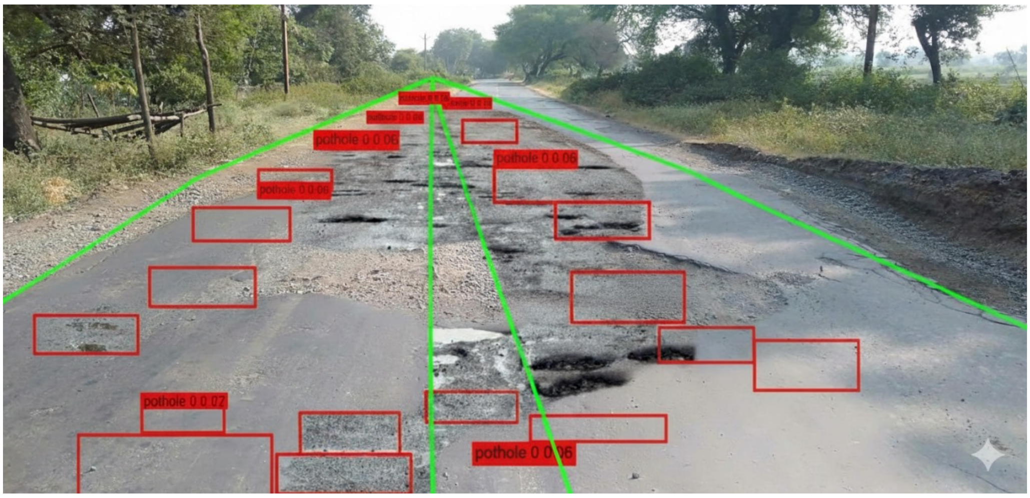

Figure 5.2: Result

### 5.4.1 Demographic Result

While this study does not involve human subject demographics in the traditional sense, the research is deeply contextualized within the infrastructural and socioeconomic demographics of Low and Middle Income Countries, specifically Bangladesh. The target deployment environment exhibits several characteristic patterns: urban road networks in cities like Dhaka feature extremely high traffic density with mixed vehicle types (cars, buses, motorcycles, rickshaws) operating in poorly regulated conditions; road maintenance budgets are constrained, resulting in widespread pavement deterioration and deferred repair cycles that allow potholes to persist and expand; lane marking application and renewal occurs infrequently, leading to extensive stretches with faded, incomplete, or absent markings; and street lighting infrastructure is inconsistent, with many secondary roads lacking illumination entirely while primary corridors feature irregularly spaced lighting creating high-contrast zones. Economic constraints shape the technological landscape: vehicle ownership patterns skew toward lower-cost models with minimal or no integrated driver assistance features; aftermarket ADAS installation, where it occurs, must operate on budget-conscious hardware platforms; and connectivity infrastructure limitations (intermittent mobile networks, data cost sensitivity) favor edge-processing solutions over cloud-dependent architectures. The field-collected validation dataset reflects these conditions: images captured during October 2024 evening hours (18:00-20:00) in Dhaka metropolitan area document roads exhibiting multiple concurrent challenges including severe pothole clusters, barely visible lane markings on concrete surfaces, dramatic lighting variations from vehicle headlights, and visual occlusions from parked vehicles and roadside vendors. This real-world data distribution differs substantially from Western highway datasets that dominate existing ADAS research.

### 5.4.2 Exploratory Analysis

The dual-branch architecture's performance advantages derive from task-appropriate feature extraction strategies that align with the distinct visual characteristics of lanes versus potholes. Lane markings, when present, exhibit strong geometric structure—predominantly linear or smoothly curved boundaries with consistent width and high-contrast edges relative to pavement. This structured nature makes them amenable to classical edge detection and polynomial fitting approaches that explicitly model geometric priors. The OpenCV pipeline's deterministic processing provides several benefits: computational efficiency through shallow processing stages, interpretable intermediate outputs (edge maps, line candidates) facilitating debugging, and geometric constraints that suppress spurious detections from non-lane edges. Conversely, potholes exhibit highly variable visual appearance lacking consistent geometric structure: shape varies from circular to irregular polygons; size spans orders of magnitude from small surface cracks to meter-wide failures; texture differs based on pavement material, water accumulation, and surrounding context; and edges may be sharp (recent failures) or gradual (weathered deterioration). This variability necessitates learned feature representations capable of capturing diverse appearance patterns—precisely the strength of convolutional neural networks. The YOLO-based detector's hierarchical feature extraction (R-ELAN backbone providing multi-scale representations, FPN/PAN neck enabling cross-scale fusion, A2C2f attention blocks highlighting relevant features) enables robust detection across this appearance diversity. The failure case analysis reveals systematic limitations: shadow-induced false positives (15.8

## 5.5 Findings and Challenges

The research establishes several key findings regarding dual-task perception in resource-constrained environments. First, task-specific architectural choices—classical geometry-based processing for structured lane features, learned representations for variable-appearance hazards—outperform unified end-to-end models when visual cues compete within the same scene. The dual-branch design achieves macro-F1 approximately 87.1However, significant challenges remain: shadow-induced false positives represent the most common failure mode, accounting for 15.8Severely degraded lane markings cause the second major failure pattern, with recall dropping 12

# Chapter 6

# Sustainability

## 6.1 System sustainability and mitigation plan

System sustainability relies on modular dual-branch architecture enabling independent lane and pothole detection updates without full redesigns. Hardware sustainability targets commodity GPUs (GTX 1050Ti) avoiding vendor lock-in, while 64 FPS throughput provides computational headroom for future enhancements. The exceptional 10.8-day payback period ensures economic viability. Mitigation strategies include establishing open-source repositories for community contributions, implementing backward-compatible APIs, creating tiered licensing models sustaining development, partnering with universities for regional maintenance expertise, maintaining containerized deployments for platform portability, and deploying performance monitoring dashboards triggering proactive model retraining before accuracy degradation. Regular stakeholder engagement ensures alignment with evolving safety regulations.

## 6.2 Social effects analysis and mitigation plan

The system reduces accidents (500/year per 1,000 vehicles), democratizes ADAS technology for LMIC populations, and empowers data-driven municipal maintenance. However, negative effects require mitigation: job displacement addressed through retraining inspectors as system operators; technology over-reliance mitigated via mandatory driver training emphasizing limitations; digital divide bridged through smartphone-based implementations; privacy concerns resolved with local-only processing and encryption; urban-rural disparity reduced via subsidized rural deployment programs. Community engagement through public demonstrations and transparent communication about capabilities builds trust. Partnerships with ride-sharing platforms ensure broad accessibility across socioeconomic segments, promoting equitable road safety improvements.

## 6.3 Environmental effects analysis and mitigation plan

Positive environmental impacts include reduced tire waste, extended component lifespan, and 130 metric tons annual $CO_2$ savings from 2% fuel efficiency improvements across 1,000-vehicle fleets. Edge computing eliminates cloud transmission energy costs. Mitigation addresses e-waste through modular hardware design, certified recycler partnerships, RoHS compliance, and 5+ year operational lifespans. Power management protocols minimize energy consumption (15-20W). Renewable energy integration for stationary installations and lifecycle carbon assessments using ISO 14040 standards quantify net impact. Leveraging existing dashcam markets and commodity hardware reduces manufacturing environmental footprint compared to specialized component production.

## 6.4 Technical sustainability analysis and mitigation plan

Technical sustainability is ensured through decoupled dual-branch architecture enabling independent module evolution, deterministic OpenCV lane detection providing interpretable results, and PyTorch framework compatibility with evolving hardware. Key risks and mitigations include: model degradation addressed via continuous monitoring and automated retraining pipelines; hardware obsolescence mitigated through ONNX cross-platform compatibility; dataset bias resolved via stratified collection across regions, surfaces, and weather conditions; calibration drift prevented using automated self-calibration algorithms; dependency vulnerabilities managed through security scanning and version pinning. Comprehensive testing (¿80% code coverage) and continuous benchmarking ensure reliability. The 64 FPS headroom accommodates future algorithm complexity increases.

## 6.5 Operational sustainability analysis and mitigation plan

Operational sustainability requires trained personnel, support infrastructure, and sustainable business models. The system operates on vehicle power (15-20W) without auxiliary batteries. Mitigation strategies address key challenges: limited technical expertise resolved through comprehensive training programs and vocational school partnerships; supply chain disruptions managed via regional parts inventories and multiple supplier qualification; inconsistent power quality handled with voltage regulators and battery backup; harsh environmental conditions mitigated through IP67 ruggedized enclosures and thermal management; user resistance overcome via phased rollouts and tangible benefit demonstrations; regulatory compliance maintained through proactive standards participation. Service level agreements with 99% uptime targets build operator confidence.

## 6.6 Ethical issues and Mitigation plan

Ethical concerns require comprehensive mitigation. Privacy risks from continuous camera operation are addressed through edge-only processing, automatic data deletion after 24-48 hours, on-device face blurring, informed consent mechanisms, and geofencing controls. Algorithmic bias is mitigated via stratified data collection across socioeconomic areas, disaggregated performance reporting, fairness audits, and community-driven dataset curation. Safety-critical accountability is ensured through transparent logging, clear operational domain specifications, mandatory driver training, and fail-safe conservative warnings. Digital divide concerns are resolved via smartphone implementations, subsidized public transit installations, and microfinance equipment loans. Environmental justice is promoted through ethical supply chain audits and certified e-waste recycling, preventing harm to vulnerable communities.

# Chapter 7

# Conclusion

### 7.0.1 Project Summary

This research presents a lightweight dual-task vision system for simultaneous lane marking and pothole detection in low-light, resource-constrained environments typical of low- and middle-income countries (LMICs). The system addresses critical challenges in Bangladesh and similar regions where poorly illuminated roads, faded lane markings, and unmaintained surfaces frequently co-occur, creating hazardous driving conditions. The proposed solution integrates a low-light enhancement front end with two parallel processing branches: an OpenCV-based lane delineation module utilizing Canny edge detection, Hough transform, and inverse perspective mapping for geometric lane extraction, and a YOLO-based detector featuring an R-ELAN backbone with FPN/PAN neck and A2C2f attention blocks for pothole localization. Evaluated on public datasets and evening-time urban scenes from Dhaka, the system achieves 88.7% recall and 84.2% precision for pothole detection, with 89.3% recall and 86.5% precision for lane detection, while maintaining real-time performance at 64 FPS on an NVIDIA GTX 1050Ti. The deployment-ready architecture demonstrates that practical ADAS perception is achievable on modest hardware without cloud dependency, making it suitable for edge computing scenarios in resource-limited settings where infrastructure-related defects contribute significantly to traffic accidents.

### 7.0.2 Future Work

Several promising research directions can enhance the system's robustness and applicability. First, adaptive anchor strategies will be investigated to improve detection accuracy for potholes of varying sizes and irregular shapes commonly found on LMIC roads. Temporal reasoning over video streams will enable trajectory prediction and hazard tracking across frames, reducing false positives from transient shadows or reflections. Multimodal sensor fusion incorporating LiDAR or low-cost radar will provide depth information for more reliable pothole depth estimation and night-time detection where visual cues are severely degraded. Lightweight transformer-based backbones such as MobileViT or EfficientFormer will be explored to capture long-range dependencies while maintaining computational efficiency. Aggressive model compression techniques including knowledge distillation, pruning, and 8-bit quantization will expand deployment options to lower-cost edge devices like Raspberry Pi 4 or smartphone processors. Additionally, synthetic-to-real domain adaptation using generative models can augment training data for underrepresented road conditions. Finally, integration with vehicle control systems for automated pothole avoidance and development of a crowdsourced pothole mapping platform will create immediate practical value for municipalities and fleet operators, transforming the system from passive detection to active safety intervention.

# Bibliography


[1] D. Neven, D. B. Bert, S. Georgoulis, M. Proesmans, and L. V. Gool, "Towards end-to-end lane detection: An instance segmentation approach," *arXiv preprint*, 2018. Available at: `https://doi.org/10.48550/arxiv.1802.05591`.

[2] N. Ma, J. Fan, W. Wang, J. Wu, Y. Jiang, L. Xie, and R. Fan, "Computer vision for road imaging and pothole detection: A state-of-the-art review of systems and algorithms," *Transportation Safety and Environment*, vol. 4, no. 4, pp. 1–22, 2022.

[3] H. Godthelp and A. Ksentini, "Specific road safety issues in lmics: An overview and some illustrative examples," *Traffic Safety Research*, vol. 8, p. e000068, 2024.

[4] H. Bello-Salau, A. J. Onumanyi, R. F. Adebiyi, A. D. Adekale, R. Bello-Salahuddeen, and O. Ajayi, "A critical appraisal of various implementation approaches for real-time pothole anomaly detection: Toward safer roads in developing nations," in *Engineering Proceedings*, vol. 22, p. 301, 2023.

[5] K. S. Rani, K. Chandana, N. Kavya, K. Poojitha, and T. Pallavi, "Automatic detection of damaged roads and lane detection using deep learning," *INTI Journal of Data Science (JODS)*, vol. 2025, pp. 1–10, 2025.

[6] R. Khanam and M. Hussain, "What is yolov5: A deep look into the internal features of the popular object detector," *arXiv preprint*, 2024. Available at: `https://doi.org/10.48550/arxiv.2407.20892`.

[7] N. Formosa, M. Quddus, C. K. Man, M. K. Singh, C. Morton, and C. B. Masera, "Evaluating the impact of lane marking quality on the operation of autonomous vehicles," *Journal of Transportation Engineering, Part A: Systems*, vol. 150, no. 1, p. 04023073, 2023.

[8] X. Zhou, R. Ke, H. Yang, and C. Liu, "When intelligent transportation systems sensing meets edge computing: Vision and challenges," *Applied Sciences*, vol. 11, no. 20, p. 9680, 2021.

[9] I. Rahman, M. M. Rahman, and M. Ahmed, "Safety sign, symbols and road markings of 'planned residential areas with adjacent highways' of dhaka city: Focusing pedestrian safety," *IOSR Journal of Mechanical and Civil Engineering*, vol. 11, no. 2, pp. 17–27, 2014.

[10] R. Zhang, J. Peng, W. Gou, Y. Ma, J. Chen, H. Hu, W. Li, G. Yin, and Z. Li, "A robust and real-time lane detection method in low-light scenarios to advanced driver assistance systems," *Expert Systems with Applications*, vol. 256, p. 124923, 2024.

[11] C. Ke, Z. Xu, J. Zhang, and D. Zhang, "Combining low-light scene enhancement for fast and accurate lane detection," *Sensors*, vol. 23, no. 10, p. 4917, 2023.

[12] H. Cui, J. Li, Z. Hua, and L. Fan, "Attention-guided multi-scale feature fusion network for low-light image enhancement," *Frontiers in Neurorobotics*, vol. 16, p. 837208, 2022.

[13] F. Guo, J. Liu, Q. Xie, and H. Chang, "Improved nighttime traffic detection using day-to-night image transfer," *Transportation Research Record*, vol. 2677, no. 11, pp. 711–721, 2023.

[14] C. He, H. Yang, Z. Zhang, *et al.*, "Dual-stream detection and segmentation framework for vision-based unmanned ground vehicle pothole perception on unstructured roads," *Journal of King Saud University - Computer and Information Sciences*, vol. 37, p. 203, 2025.

[15] Y. Hu, N. Chen, Y. Hou, *et al.*, "Lightweight deep learning for real-time road distress detection on mobile devices," *Nature Communications*, vol. 16, p. 4212, 2025.

[16] R. Soans, R. Masuda, and Y. Fukumizu, "A multi-task deep learning framework for road quality analysis with scene mapping via sim-to-real adaptation," *Applied Sciences*, vol. 15, no. 16, p. 8849, 2025.

[17] W. Lan, J. Dang, Y. Wang, and S. Wang, "Pedestrian detection based on yolo network model," in *2018 IEEE International Conference on Mechatronics and Automation (ICMA)*, (Changchun, China), pp. 1547–1551, 2018.

[18] Y. KIm, "pothole dataset." `https://universe.roboflow.com/yeeun-kim-fyvoj/pothole-vhmow`, jun 2024. visited on 2025-09-30.

[19] T. P. L. D. Solid and Dashed, "Lane detection (segmentation) dataset." `https://universe.roboflow.com/tobias-price-lane-detection-solid-and-dashed/lane-detection-segmentation-edyqp`, mar 2024. visited on 2025-09-30.