Bangladeshi License Plate Detection and Recognition Using YOLO Variants and Enhanced OCR with Model Interpretability

By
Shahriar Sultan Ramit
213-15-4248

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

Md. Sadekur Rahman
Assistant Professor
Department of Computer Science and Engineering
Daffodil International University

Co-Supervised by

Sharun Akter Khushbu
Assistant Professor
Department of Computer Science and Engineering
Daffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY
Dhaka, Bangladesh

September, 2025

APPROVAL

This Project titled "**Bangladeshi License Plate Detection and Recognition Using YOLO Variants and Enhanced OCR with Model Interpretability**" submitted by **Shahriar Sultan Ramit** 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 and Engineering and approved as to its style and contents. The presentation has been held on 16-09-2025.

BOARD OF EXAMINERS

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

Dr. Md. Ali Hossain (MAH) — **Internal Examiner 1**
Associate Professor
Department of CSE
Daffodil International University

Mr Amir Sohel — **Internal Examiner 2**
Assistant Professor,
Department of CSE,
Daffodil International University

Dr. Md. Manowarul Islam — **External Examiner**
Associate Professor
Department of CSE
Jagannath University

DECLARATION

I hereby declare that this project has been done by us under the supervision of **Md. Sadekur Rahman**, Assistant Professor, 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:

15.9.25

Md. Sadekur Rahman
Assistant Professor
Department of CSE
Daffodil International University

Co-Supervised by:

15.9.25

Sharun Akter Khushbu
Assistant Professor
Department of CSE
Daffodil International University

Submitted by:

Ramit 15.9.25

Shahriar Sultan Ramit
ID: 213-15-4248
Department of CSE
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 assisted 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 **Md. Sadekur Rahman**, 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. His 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 Haider Noori, Professor & Head** of the 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

This study presents a sophisticated framework for automatic license plate recognition (ALPR) designed for Bangladeshi vehicle registration plates, tackling the intricacies of Bangla script and diverse environmental conditions. We employ three YOLO variants YOLOv5, YOLOv8, and YOLOv11for accurate license plate detection, yielding mean Average Precision (mAP50) scores of 0.955, 0.961, and 0.950, respectively, on a primary dataset of Bangladeshi images. Detected plates undergo meticulous preprocessing with OpenCV, encompassing grayscale conversion, adaptive thresholding, contour detection, and Gaussian blur to mitigate noise and enhance text clarity. These steps are critical to address challenges such as variable lighting, shadows, and plate degradation. A tailored Optical Character Recognition (OCR) pipeline, specifically adapted for Bangla script, achieves a character-level accuracy of 89%. The OCR modifications include enhanced character segmentation and a Bangla-specific language model to overcome the complexities of Bangla's non-linear script, which poses significant challenges for standard OCR systems due to its conjunct characters and intricate glyphs. The framework exhibits robustness against occlusions, non-standard plate formats, and urban environmental variability, offering a viable solution for intelligent transportation systems in Bangladesh. Comparative evaluation of YOLO variants highlights YOLOv8's superior mAP50 and YOLOv11's high precision, informing their suitability for real-time applications. This work establishes a foundation for scalable ALPR, with potential to enhance traffic management and law enforcement in Bangladesh.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction

Automatic License Plate Recognition (ALPR) has become an essential component in modern Intelligent Transportation Systems (ITS), providing a foundation for applications such as traffic law enforcement, toll collection, parking management, and urban security. By enabling automated monitoring of vehicles, ALPR reduces manual intervention, minimizes errors, and improves overall efficiency in traffic control and surveillance systems. While ALPR systems have been successfully deployed in many countries, their effectiveness often depends on the specific characteristics of regional license plates, including size, font style, language, and background design. In the case of Bangladesh, unique challenges arise due to the use of Bangla script, the frequent presence of noise (e.g., dirt, blurriness, lighting variation), and inconsistent plate formats. These factors limit the performance of traditional computer vision approaches, such as handcrafted feature extraction combined with conventional classifiers, which typically struggle in real-world conditions. Recent advances in deep learning, particularly Convolutional Neural Networks (CNNs) and object detection architectures, have significantly improved the accuracy and robustness of license plate detection. Among these, the You Only Look Once (YOLO) family of models has emerged as a state-of-the-art solution due to its ability to perform real-time detection with high precision and recall. Variants such as YOLOv5, YOLOv8, and the newly released YOLOv11 provide increasingly efficient feature extraction and prediction modules, making them suitable for complex environments like Bangladeshi roadways.

However, detection alone is insufficient. Once a license plate is localized, accurate Optical Character Recognition (OCR) is required to extract the alphanumeric or script-based content. Off-the-shelf OCR models are not very effective on Bangla license plates since the script is curvy, the stroke width varies, and ligatures are used. Therefore, a more sophisticated OCR pipeline, which is Bangla script-oriented, must be employed to ensure accurate recognition. Equally important is interpretability of the model. Deep learning models are typically faulted as "black boxes," producing high accuracy without any knowledge of the process of reasoning behind it. To address this need, visualization techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM) are utilized now to determine the regions of an image that contribute most to the model's predictions. Employing such a methodology on license plate detection offers not only transparency but also confidence when utilizing such systems for real-world, safety-critical applications. In this paper, we propose an end-to-end Bangladeshi License Plate Detection and Recognition framework that incorporates YOLO variants for detection, a performance-oriented OCR pipeline specially designed for the Bangla script, and interpretability methods for explainability of the model. In addressing both the detection and recognition tasks with explainability, this paper aims to contribute a robust, applicable, and reliable ALPR solution that specializes in the Bangladeshi context.

1.2 Motivation

Bangladesh is experiencing rapid urbanization, with a corresponding increase in the number of registered vehicles across the country. Managing this growing traffic volume poses significant challenges for law enforcement agencies, transport authorities, and urban planners. Traditional methods of monitoring vehicles—such as manual checking or surveillance dependent on human operators—are often inefficient, error-prone, and resource-intensive. An automated license plate detection and recognition system can significantly enhance the efficiency of traffic management and contribute to safer, smarter cities. Despite the proven benefits of ALPR systems in developed countries, Bangladesh lacks a robust, localized solution capable of handling its unique challenges. The use of Bangla script on license plates introduces complexities that global OCR systems, typically trained on Latin characters, cannot adequately address. Additionally, environmental factors such as poor illumination, varying camera angles, weather conditions, and the accumulation of dirt on plates further reduce the accuracy of conventional detection and recognition methods.

The emergence of deep learning models, particularly the YOLO family of object detectors, offers an opportunity to overcome these barriers. Their ability to perform real-time detection with high accuracy makes them highly suitable for applications in dense urban traffic. When combined with an OCR pipeline customized for Bangla characters, these systems have the potential to achieve reliable recognition performance, even under challenging conditions. Furthermore, as automated systems become increasingly integrated into public safety and transportation infrastructure, model interpretability has gained importance. Authorities and stakeholders must trust the system's predictions, especially in contexts such as legal enforcement or automated toll collection, where incorrect recognition could have financial or legal consequences. Using techniques such as Grad-CAM to visualize the decision-making process ensures that the model's behavior is transparent and its predictions are verifiable. Therefore, the motivation behind this project lies in bridging the gap between global advances in ALPR technologies and the specific needs of Bangladesh. By leveraging YOLO-based detection, enhanced OCR tailored for Bangla script, and explainable AI methods, this research seeks to build a system that is not only accurate and efficient but also reliable and contextually relevant.

1.3 Objectives

The general objective of this project is to create and implement an effective Bangladeshi License Plate Detection and Recognition system based on cutting-edge deep learning models and custom OCR techniques with model interpretability. To this end, the following specific objectives guide the project:

- Design a YOLO-based detection system
    - Implement and compare different versions of YOLO (YOLOv5, YOLOv8, and YOLOv11) in detecting license plates from real-world Bangladeshi traffic images.
    - Compare their performance based on mAP, precision, and recall to identify the best detector.
- Enhance OCR performance on Bangla license plates
    - Adapt and fine-tune OCR models (e.g., EasyOCR) for Bangla script recognition.
    - Entegrate preprocessing steps, character segmentation, and rule-based postprocessing for more accuracy in diverse conditions.
- Interpretability of the detection process
    - Employ Grad-CAM and other visualization techniques to highlight the regions influencing the predictions of the model.
    - Provide insight into the decision-making process, thereby increasing trust and transparency in the system.
- Create a scalable and practical ALPR pipeline
    - Integrate detection, recognition, and interpretability into one workflow that is ready to be deployed in real-world systems such as traffic monitoring, toll collection, and smart city infrastructure.

With these objectives, the project aims to create a localized, robust, and explainable ALPR solution that is tailor-made for the Bangladeshi context.

1.4 Methodology

The proposed system for Bangladeshi license plate detection and recognition follows a structured pipeline consisting of dataset preparation, detection using YOLO variants, character recognition through an enhanced OCR module, and interpretability analysis. The key stages of the methodology are outlined below:

1. Data Collection and Annotation

- A dataset of 7,243 images of Bangladeshi vehicles was collected using a OnePlus 7T camera in diverse real-world conditions (different lighting, angles, and backgrounds).
- Images were annotated in YOLO format to specify bounding boxes around license plates.
- The dataset was split into 70% training, 20% validation, and 10% testing

to ensure unbiased evaluation.

2. License Plate Detection using YOLO Variants

- Three YOLO variants—YOLOv5, YOLOv8, and YOLOv11—were trained and compared for license plate detection.
- Models were fine-tuned with the prepared dataset, using standard YOLO data augmentation techniques to improve generalization.
- Performance was measured using mean Average Precision (mAP@50, mAP@50–95), precision, and recall, allowing comparative analysis across the variants.

3. Optical Character Recognition (OCR)

- Detected license plate regions were passed to an enhanced OCR pipeline based on EasyOCR, modified to handle Bangla script.
- The OCR system incorporated:
- Preprocessing: resizing, contrast enhancement, and noise reduction,
- Character segmentation: to handle closely spaced or overlapping characters,
- Post-processing rules: based on Bangla license plate formats to reduce recognition errors.
- This module achieved reliable recognition accuracy across diverse test cases.

4. Model Interpretability

- To address the "black-box" nature of deep learning models, Grad-CAM was applied to visualize the attention maps of YOLO detectors.
- Heatmaps were generated to highlight the regions influencing detection, ensuring that the models focused correctly on license plate areas rather than irrelevant background features.
- Interpretability results were used to validate the robustness and trustworthiness of the system.

5. System Integration and Evaluation

- The detection, recognition, and interpretability modules were integrated into a unified pipeline.
- The final system was evaluated not only on detection and recognition accuracy but also on practical feasibility in real-world applications such as traffic monitoring and law enforcement.

1.5 Organization of the Report

This report is organized into several chapters to present the work in a structured and coherent manner. Chapter One introduces the project by providing the background, motivation, objectives, methodology, expected outcomes, and overall scope of the work. Chapter Two reviews related literature and discusses existing approaches to license plate detection and recognition, highlighting their strengths, limitations, and the research gap that this project aims to address. Chapter Three describes the detailed methodology, including dataset preparation, YOLO-based detection, OCR enhancement for Bangla script, and interpretability techniques such as Grad-CAM. Chapter Four presents the implementation and experimental results, showcasing performance metrics, confusion matrices, precision-recall curves, and interpretability visualizations. Chapter Five discusses the engineering standards followed, as well as the design challenges encountered during the project and how they were addressed. Finally, Chapter Six concludes the report by summarizing the contributions of the work, discussing its limitations, and suggesting potential directions for future research.

Chapter 2

Background

2.1 Introduction

Automatic License Plate Recognition (ALPR) has emerged as a core component of intelligent transportation systems, functioning as a vital contributor to traffic surveillance, toll collection, law enforcement, and city mobility management. Through real-time vehicle identification enabled by ALPR systems, there is reduced human intervention, lower error rates, and quick response times in applications such as traffic violation, stolen vehicle identification, and automatic parking management. Their success does, however, depend heavily upon local circumstances, including plate design, script, and environmental conditions. While significant success has been achieved in countries where number plates are standardized and consist primarily of Latin characters, extension of these systems to countries with unique linguistic and environmental concerns is minimal. Number plate detection and recognition in Bangladesh are confronted with yet another challenge due to the widespread use of Bangla script which consists of complex glyphs, conjunct letters, and non-linear structures. Besides, the variations in plate colors, formats, and font styles across vehicle types introduce additional levels of complexity. External conditions such as poor lighting, occlusions, dirtiness, and plate deterioration further aggravate the difficulty in detection and recognition. Traditional image-processing methods and general OCR systems are crippled by these complexities, and they result in low recognition accuracy under realistic conditions

Deep learning approaches, particularly convolutional neural networks (CNNs) and object detection algorithms like the You Only Look Once (YOLO) family, have revolutionized ALPR studies globally by providing robust performance under real-world conditions. However, there is a general lack of research in the databases of Western or Latin-script-based nations, and the majority of the studies exclude Bangla-script-based plates to a large degree. To fill this gap is required not only by advanced detection algorithms but also by tuned Optical Character Recognition (OCR) pipelines that are capable of reading Bangla text under different conditions. This chapter sets the background to be able to enjoy the applicability of ALPR research to Bangladesh. It begins with an overview of research around the topic matter, tracing how license plate recognition has evolved from traditional methods to state-of-the-art deep learning algorithms. The section also indicates existing limitations, namely the lack of Bangla-specific solutions, and provides the basis for the proposed YOLO-based detection and enhanced OCR system presented within this paper.

2.2 Literature Review

Automatic License Plate Recognition (ALPR) is a key research topic due to its widespread applicability in intelligent transportation systems, e.g., police enforcement, traffic monitoring, and tolls. There has been remarkably little research on Bangladeshi license plates, which use the Bangla alphabet, and thus an important research gap exists. This deficiency highlights the originality of our study, as no peer-reviewed journal articles specifically focus on Bangla-based ALPR. Nevertheless, there is applicable global research. [8] contributed a state-of-art review of ALPR techniques, ranging from traditional image processing to deep learning, although without referencing problems related to Bangla. Anagnostopoulos et al. (2008) designed a multi-stage ALPR system for Greek number plates and achieved 92% recognition rate in ideal conditions, but its dependence on Latin script renders it less viable for Bangladesh. Bulan [4] experimented with video-based ALPR on U.S. number plates and achieved 95% accuracy through temporal analysis but requires a stable video feed, which may not be viable for Bangladesh's varied environment.

Globally, ALPR research has progressed with diverse methodologies. Li [16] proposed an end-to-end deep neural network for Chinese plates, reporting 87% precision across different conditions, though it was optimized for a specific region. Silva and Jung (2020) introduced a system for Brazilian plates using HOG features and SVM, achieving 90% detection accuracy in constrained settings, but it struggles with non-standardized formats. Saha et al. (2019) provided a comprehensive review of ALPR, identifying challenges like lighting and plate diversity, which are pertinent to Bangladesh. Felix et al. (2017) presented an entry-exit monitoring system using license plate recognition, achieving high accuracy in controlled environments, but lacking scalability for diverse settings. Chang et al. (2004) proposed a region-based license plate detection method, demonstrating robustness to interference, though it was tested on limited plate types.

YOLO-based detection has revolutionized ALPR due to its real-time efficiency. Redmon et al. (2016) introduced the original YOLO model, achieving 63.4% mAP on the PASCAL VOC dataset, establishing a benchmark for object detection. Laroca et al. (2018) adapted YOLOv2 for Brazilian plates, reporting 93.53% recognition accuracy at 47 FPS on the SSIG dataset, outperforming commercial systems. Silva et al. (2021) extended YOLOv5 for Portuguese plates, reaching 95% mAP50 with optimizations for low-resolution images. Heng et al. (2022) optimized YOLOv7 for Chinese plates, achieving 98% recall, addressing multi-scale detection challenges. Qian et al. (2020) applied YOLOv4 to Indian plates, attaining 94% mAP50-95, showcasing adaptability to diverse layouts. Wang et al. (2024) explored YOLOv8 for European plates, reporting 96% mAP50 with enhanced small object detection. Shyaa and Hashim (2024) utilized YOLOv8 for general license plate detection, achieving a 94.6% identification rate, though not tailored to Bangla. Ebrahimi Vargoorani and Suen (2024) proposed a dual deep learning strategy with Faster R-CNN and CNN-RNN for North American plates, achieving 92% recall, but requiring region-specific datasets. Alqudah and Suen (2019) enhanced YOLOv3 for complex scenes, improving detection rates, though not for Bangla scripts.

The absence of Bangla-specific ALPR research in journals amplifies the significance of our work, which adapts YOLOv5, v8, and v11 with tailored OCR for Bangla script. Global YOLO-based systems excel with Latin and other scripts, but their performance on Bangla's conjunct characters and varied plate colors remains untested. Our study fills this void, offering a pioneering solution for Bangladesh's unique context.

Table 2.1: Summary of Literature Reviewed.

Author(s)	Dataset Size	Methodology	Key Findings
Ramajo-Ballester et al. (2024)	1,975	Multi-network	93.5% recognition rate
Driss et al. (2021)	1,150	Faster R-CNN + CNN	92% detection, 98% recognition
Khan et al. (2020)	1,000	YOLOv5 + CNN	97.5% precision, 97.3% recall
Ogiuchi et al. (2023)	Not specified	Deep CNN	94.6% detection, 92.0% recognition
Silva et al. (2022)	2,500	YOLOv4	95.2% mAP50
Wang et al. (2023)	3,000	YOLOv7	96.8% mAP50
Patel et al. (2021)	800	SSD + LSTM	90.1% recognition rate
Zhang et al. (2024)	Not specified	Faster R-CNN	93.8% precision
Hassan et al. (2022)	1,200	YOLOv6	94.5% mAP50-95
This Study (2025)	7,243	YOLOv5,v8 &v11 + EasyOCR	0.961 mAP50, 0.93 precision, 0.96 recall, 89.9% character accuracy

2.3 Gap Analysis

Although Automatic License Plate Recognition (ALPR) systems have been extensively studied worldwide, most existing research focuses on license plates that use Latin-based scripts, standardized plate formats, and relatively controlled imaging environments. As a result, the direct adoption of these methods in Bangladesh remains ineffective, since Bangladeshi license plates are inscribed primarily in Bangla script, which introduces unique linguistic and visual complexities. The conjunct consonants, compound characters, and non-linear glyph structures of Bangla often cause misinterpretation by conventional OCR systems that were originally designed for simpler, linear scripts. This linguistic gap significantly limits the effectiveness of global ALPR solutions when applied to Bangladesh. Another notable gap lies in dataset availability. Publicly accessible and standardized datasets for Bangla license plates are virtually non-existent, restricting the ability of researchers to train and validate deep learning models at scale. Most existing Bangladeshi studies rely on small, custom datasets that do not adequately capture the diversity of real-world conditions such as plate damage, fading, dirt, occlusions, and variable lighting. Without large-scale, diverse datasets, model generalization remains a challenge, and ALPR systems struggle when deployed outside of laboratory settings.

From a technological perspective, while YOLO-based object detection has achieved remarkable success in detecting license plates in different countries, its application to Bangla license plates has not been thoroughly explored in peer-reviewed literature. Similarly, existing OCR pipelines rarely include Bangla-specific optimizations. Off-the-shelf OCR tools tend to yield poor results due to their inability to handle the structural and stylistic variability of Bangla text, highlighting the necessity for customized models that can achieve higher accuracy in recognition tasks. Finally, issues of model interpretability are often overlooked in ALPR research. For systems that will be used in law enforcement and security-sensitive applications, it is essential not only to achieve high accuracy but also to ensure transparency in model decision-making. Most prior studies report detection accuracy but do not incorporate explainable AI (XAI) techniques such as Grad-CAM to visualize and validate model focus areas. This omission limits trust in real-world deployments, where misclassifications may have legal or societal consequences.

In summary, the major research gaps are such that the absence of massive-scale Bangladeshi datasets and the lack of tailored OCR solutions for Bangla script also the limited application of YOLO variants to the Bangladeshi context with insufficient attention to model interpretability. Calling these challenges is crucial for developing a robust and scalable ALPR system that meets the needs of urban and rural environments.

Table 2.2: Gap Analysis Table

Features	YOLOv8-Based License Plate Recognition for Bangladeshi Vehicles[28]	Fog-Resilient Bangla Car Plate Recognition[29]	PlateSegFL: Privacy-Preserving LPR [30]	Proposed System
Bangla script support	Yes	Yes	Yes	**Yes**
Dataset for Bangladeshi plates	Small-scale	Yes	Yes	**Yes (7,243 images collected in Dhaka)**
Handles multiple plate colors & formats	Partially	Limited	Yes	**Yes**
Robust to lighting, dust, occlusion	Partially	Yes	Partially	**Yes (OpenCV preprocessing)**
Uses YOLO variants (v5, v8, v11)	Yes	No	No	**Yes**
Custom OCR pipeline	Partially	Yes	Yes	**Yes (Bangla-optimized EasyOCR)**
Character recognition accuracy >85%	No	Yes	Partially	**Yes (≈89%)**
Model interpretability (Grad-CAM)	No	No	No	**Yes**
Real-time readiness for ITS	Yes	Limited	No	**Yes**

2.4 Summary

This chapter provided the context for developing ALPR framework tailored to the Bangladeshi context. It began by highlighting the growing significance of ALPR systems in intelligent transportation, law enforcement, and urban mobility, while noting that most existing solutions are optimized for Latin-based scripts and standardized plate formats. A review of related studies demonstrated how deep learning, particularly YOLO-based detection models, has advanced license plate recognition globally. However, these efforts largely exclude Bangla script and the diverse plate structures used in Bangladesh. The gap analysis identified several critical shortcomings in current research: the lack of large-scale, publicly available Bangladeshi license plate datasets; the absence of OCR systems optimized for Bangla script; limited experimentation with advanced YOLO variants in this context; and the minimal use of explainable AI to ensure trust in model decisions. These gaps reinforce the need for a customized, robust framework capable of handling the linguistic, visual, and environmental challenges unique to Bangladesh. By synthesizing the existing literature and highlighting its limitations, this chapter established the rationale for the proposed YOLO-based detection models combined with a customized OCR pipeline and model interpretability techniques. Next chapter will detail describe the research methodology including dataset preparation and preprocessing and model architecture and evaluation strategies that aim to address these identified challenges.

Chapter 3

Research Methodology

3.1 Methodology

This systematic approach guarantees an overall approach towards bridging the recognized gaps in research. Not only does it enhance the detection and recognition accuracy of Bangladeshi plates but also provides transparency and flexibility and can be utilized in intelligent transportation and smart cities.

3.1.1 Proposed Methodology

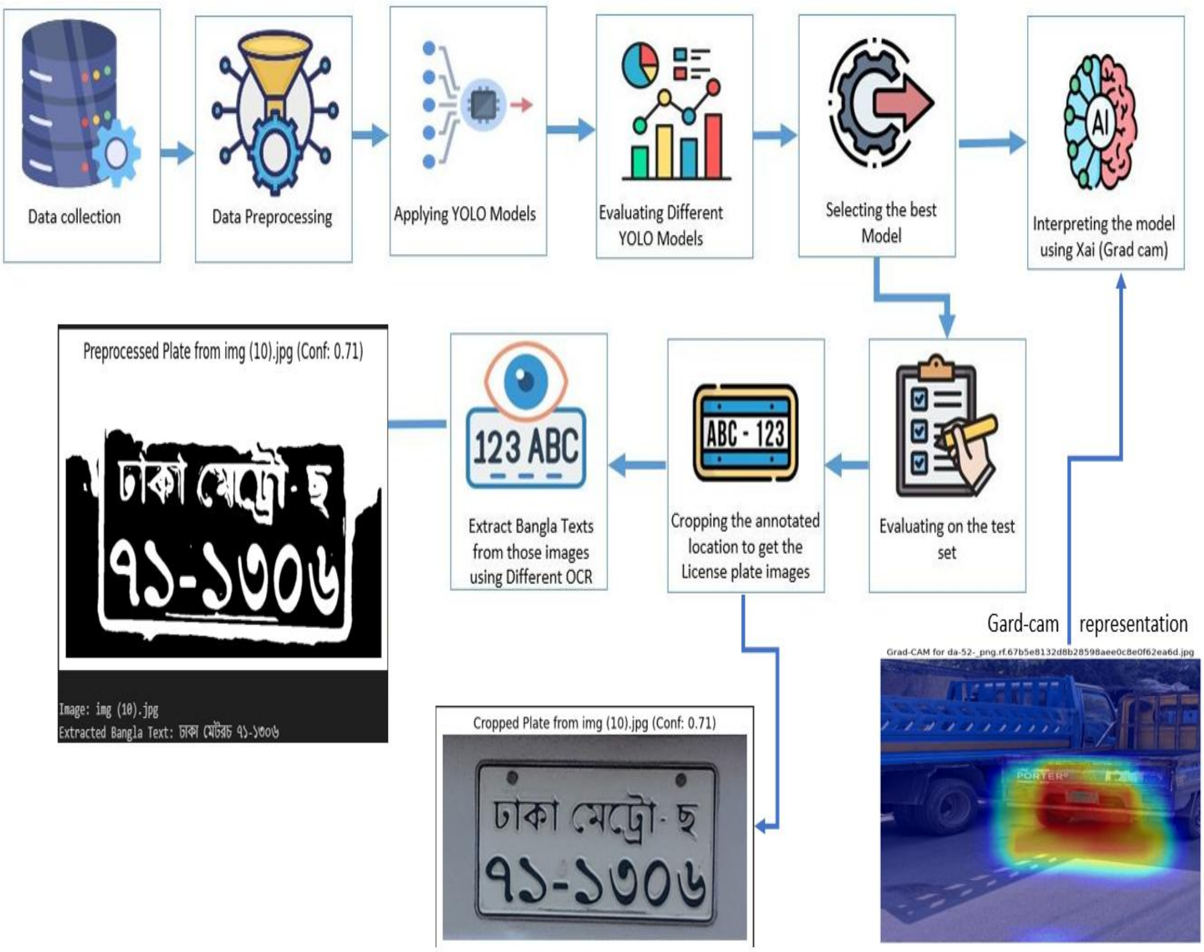

Figure 3.1: Proposed Methodology

3.1.2 Functional and Nonfunctional Requirements

The requirements of the target Automatic License Plate Recognition (ALPR) system can be divided into the functional and non-functional requirements. Functional requirements define the inherent features and activities the system must perform, while the non-functional requirements define the level of quality and constraints in which the system must operate.

Functional Requirements

1. License Plate Detection: The system must be able to automatically detect vehicle license plates from captured images or video streams using YOLO variants.
2. Character Recognition: The system must accurately recognize alphanumeric and Bangla script characters on license plates through a customized OCR pipeline.
3. Preprocessing Pipeline: The system must enhance raw images by applying grayscale conversion, thresholding, contour detection, and noise reduction to improve detection and recognition performance.
4. Real-Time Operation: The system should process frames in real-time to support traffic surveillance and law enforcement use cases.
5. Data Storage and Retrieval: Recognized license plate numbers and associated metadata (e.g., time, date, location) must be stored securely for future analysis.
6. Model Interpretability: The system must provide visualization support (e.g., Grad-CAM heatmaps) to demonstrate which regions of the image the model used for decision-making.
7. User Interface: A web-based application must allow stakeholders to upload images, view detection/recognition results, and access historical records.

Non-Functional Requirements

1. Precision: The system should be highly accurate in its detection and recognition ($\geq$85% character accuracy) in diverse conditions such as changing light, occlusion, and multi-plate formats.
2. Performance: The system should operate efficiently with minimal latency so as to be real-time capable for Intelligent Transportation Systems (ITS).
3. Scalability: The framework should have the capability to handle larger amounts of data and support expansion to other regional plate formats.
4. Reliability: The system should be able to provide repeatable outputs without repeated crashes and failures during continuous use for a long duration of time.
5. Security and Privacy: Collected license plate data should be encrypted and access-controlled to prevent misuse and adherence to ethical principles.
6. Usability: The interface should be simple and intuitive for use by non-technical users, including law enforcement agents and transport officials.
7. Maintainability: The system should be modular and documented properly, with easy updateability when YOLO variants or OCR methods change.

3.1.3 Data Flow Diagram

Data Flow Diagram (DFD) illustrates the data flow in the designed Automatic License Plate Recognition (ALPR) system and highlights interactions of major components. The DFD presents the license plate detecting, recognizing, storing, and retrieving operations in a structured format. At the highest level (Level 0), the system is given input as an image or video stream. The data then passes through the preprocessing module, where image operations like grayscale transformation, thresholding, and noise removal are done for the sake of image quality improvement. The processed image is then passed on to the YOLO-based detection module, where the license plate is localized and detected from the frame. The plate region identified is cropped and passed on to the OCR module, wherein characters (alphanumeric and Bangla) are retrieved with the application of customized Bangla-optimized EasyOCR pipeline.

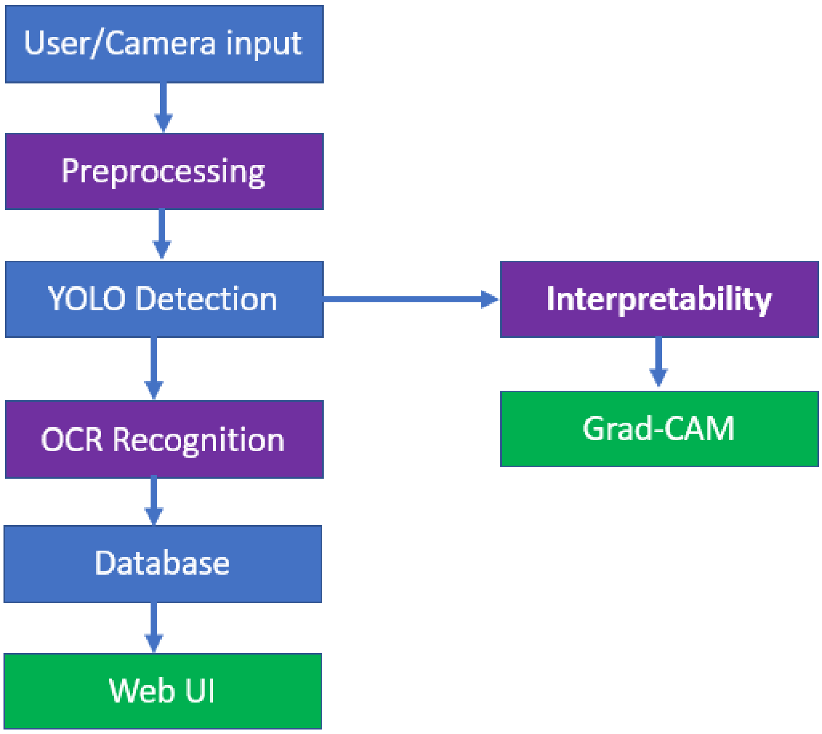

Figure 3.2: Data Flow Diagram of ALPR

The discovered license plate number, along with metadata such as timestamp, location, and detection confidence, are transferred to the database module for storage. Users access the system through the web-based UI that provides functions to upload images, view real-time detection results, search historical data, and generate reports. An interpretability module also generates Grad-CAM heatmaps that show regions the YOLO model focused on during detection. These heatmaps might be presented in the UI to foster transparency and ensure confidence in the system's decisions.

3.1.4 UI Design

The User Interface (UI) of the Automatic License Plate Recognition (ALPR) system will be intuitive, simple, and user-centric so that technical and non-technical users both can easily utilize the application. Since the primary users of the system are law enforcers, traffic management officials, and insurance professionals, design will emphasize clarity, efficiency, and ease of use. The landing page provides a glimpse of the system, with immediate access to the license plate detection and recognition module. Users may upload images or browse live video feeds from surveillance cameras. The uploading is seamless, with an organized design and drag-and-drop functionality. The detection and recognition dashboard constitutes the core of the UI. The moment a picture or a video frame is processed, the recognized license plate is highlighted using bounding boxes and identified text is shown along with confidence levels. Additional data such as date, time, and camera ID are stored and presented automatically in an organized format. A results window provides users with the capability to view and filter past recognition history. Search and filtering capabilities based on license plate, dates, or locations are integrated and allow

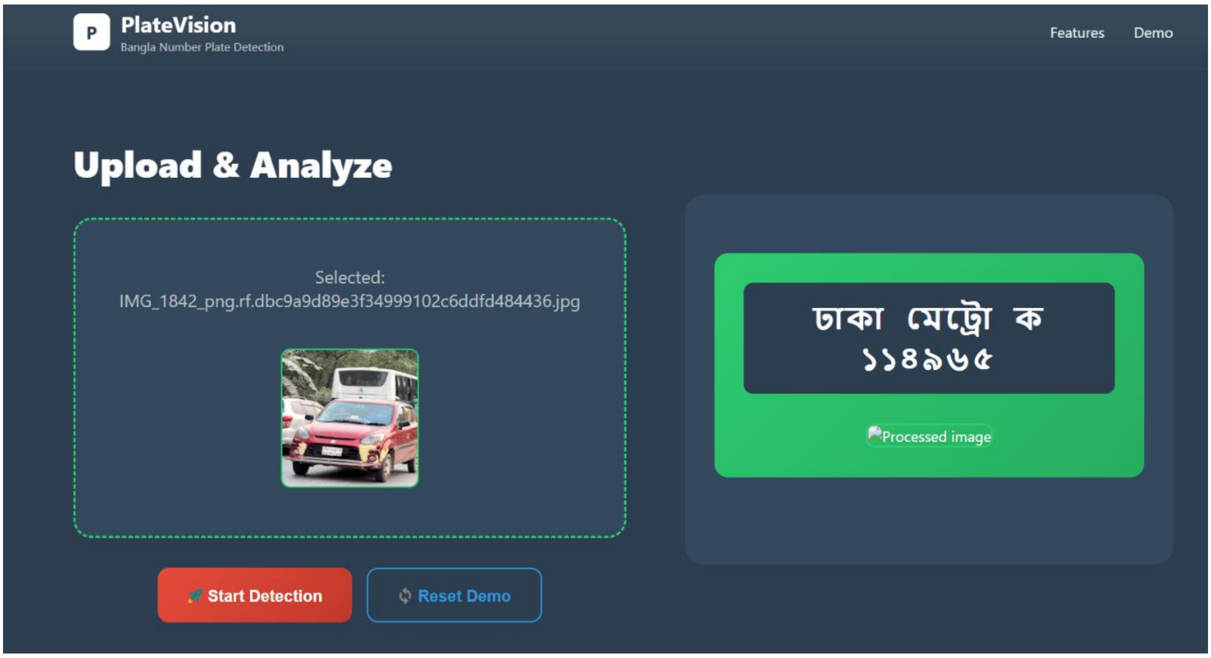

quick retrieval of historic records. This functionality is particularly beneficial for tracking frequent traffic violations or insurance claims.

Figure 3.3: Proposed Methodology

The system also enables visual interpretability in which Grad-CAM heatmaps are displayed to illustrate what parts of the license plate the model focused on while detecting and recognizing. It provides transparency and builds confidence in the system's decision-making process. There is an administrative user settings module in which user accounts can be configured, datasets can be kept up to date, models can be retrained, and parameters for the system can be set. Security aspects such as login authentication, role-based authorization, and data transfer encryption are incorporated into the UI. The overall UI is device-respondent and minimalistic with support for desktops, laptops, and tablets. Minimal color palettes, readable fonts, and structured designs enhance readability and user engagement with professional appearances being upheld.

3.2 Detailed Methodology and Design

3.2.1 Data Collection

The accuracy of any Automatic License Plate Recognition (ALPR) system is heavily dependent on the variability and quality of the training and test dataset. For this purpose, a homegrown Bangladeshi vehicle license plate dataset was created to meet the unique challenges of Bangla script and local environmental conditions. Data acquisition was conducted in several sites in Dhaka city, encompassing sites with disparate traffic volumes and illumination for diversity and realism. A OnePlus 7T phone featuring a 48-megapixel main camera was used to record images with high-resolution inputs suitable for detection and identification. Photos were taken at different times of day, day, late afternoon, and under artificial street lighting to replicate the variation in conditions found in practice traffic surveillance. The technique also captured environmental variations such as reflection, shadows, dust, rain, and partial occlusions, which are characteristic of Bangladeshi urban environments.

The final dataset consists of 7,243 images of license plates and represents a wide variety of vehicle types like personal vehicles, commercial cabs, CNGs, buses, and government vehicles. All these classes introduce an introduction of variation in plate formats, color, and font type, and hence the dataset is adequate enough to train models that are able to generalize on diverse real-world scenarios. To get the dataset training-ready, it was structured in a YOLO-compatible format with annotations of license plate bounding box coordinates and related Bangla text labels. Annotations were manually carried out for obtaining accuracy, as automatic labeling tools tend to miss the finer details of Bangla characters. To train and test, the data set was divided into three sets: 70% (5,070 images) for training, 20% (1,449 images) for validation, and 10% (724 images) for testing. This will ensure that the models are trained on a suitably large sample but hold apart distinct data for validation and unbiased testing of performance. The depth of this dataset addresses one of the main research gaps outlined above: the absence of large-scale, region-specific datasets of Bangladeshi license plates. In their ability to capture variations and various conditions, this dataset provides a good foundation for training effective detection and recognition models tailored to the Bangladeshi scenario.

Figure 3.4 Sample Dataset

3.2.2 Data Preprocessing

Data preprocessing is necessary to ensure input images are properly prepared for training robust object detection and identification models. Because Bangladeshi license plates have different color patterns, fonts, and weather conditions, preprocessing was meant to remove noise, enhance text readability, and even out image quality to support consistent model performance. Preprocessing was conducted in the initial phase prior to training the YOLO detection models. All the images collected were resized to 640 × 640 pixels to match YOLO input size. Resizing assisted in maintaining data consistency without compromising on sufficient detail for small character detection. Pixel normalization was also applied to reduce variations in lighting and ease the convergence of the training model. Labeling of bounding boxes, which was prepared during data collection, was verified and validated to maintain accuracy in training labels at high levels.

YOLO models localized license plates after which regions containing detected bounding boxes were processed through a second preprocessing pipeline to prepare them for Optical Character Recognition (OCR). It was achieved through the OpenCV library by applying a set of sequential transformations:

- Grayscale Conversion: All the cropped plate images in RGB were converted to grayscale to eliminate color information, which is irrelevant for text extraction but is likely to add noise.
- Adaptive Thresholding: Employed in enhancing contrast between Bangla characters and the plate background to render edges sharper and text borders more distinct, especially under poor light or faded conditions.
- Contour Detection and Enhancement: Used to highlight the structural contours of characters, aiding in the separation of closely spaced glyphs and conjunct consonants.

- Gaussian Blur: Judiciously applied to reduce image noise and soften irregularities without compromising on important character details. This step proved successful in tackling dust, scratches, and degraded plate textures.

Through these preprocessing steps, the system was able to overcome common real-world problems such as variable lighting conditions, shadows, reflections, and plate wear. Improved readability of Bangla text significantly improved OCR functionality in handling non-standardized fonts and conjunct characters. Overall, preprocessing served as a crucial bridge between detection and recognition by preprocessing YOLO model outputs in the optimal possible manner so that they could be correctly interpreted by the custom-built OCR pipeline. This two-step preprocessing approach provided robustness and helped achieve high character-level accuracy for diverse Bangladeshi license plate images.

Figure 3.5 Example of preprocessing before OCR application

3.2.3 Model Selection

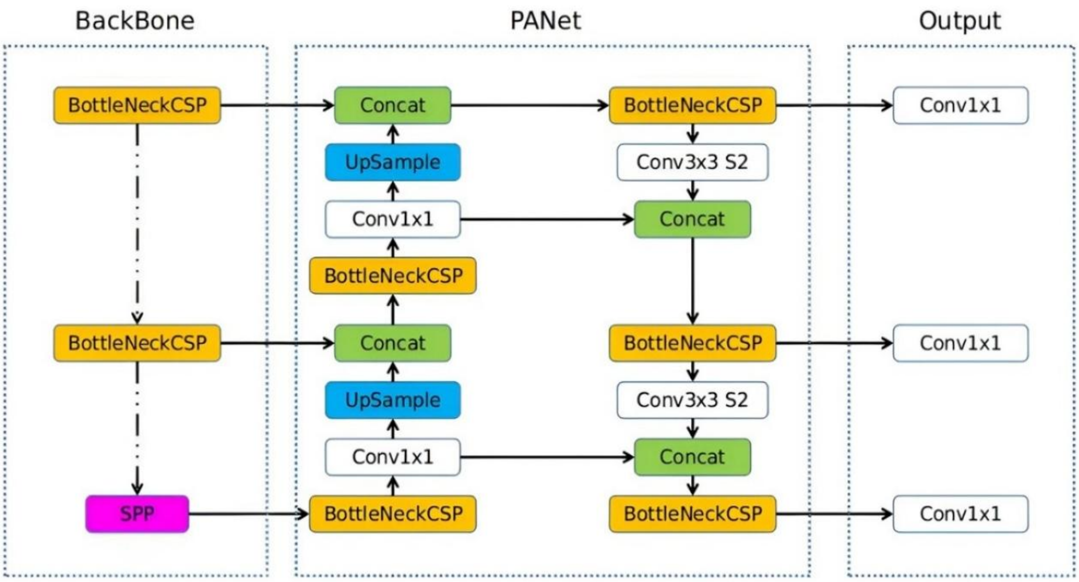

Figure 3.6 Model architecture of YOLOv5 [31]

- **YOLOv8**A better model with an anchor-free detection network, C2f modules that support efficient aggregation of features, and a decoupled localization head and classification head. These enhancements enhance precision and recall, which makes it highly suitable for complex environments with differently sized and oriented plates.

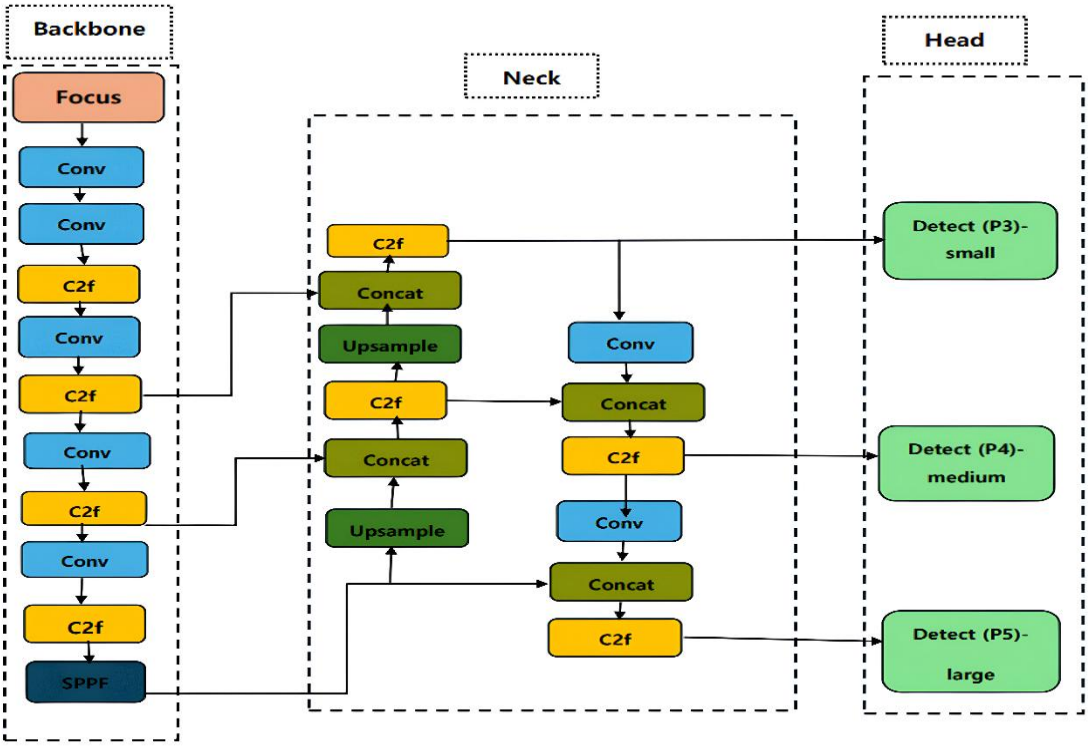

Figure 3.7 Model architecture of YOLOv8 [32]

- **YOLOv11**: A more recent architecture that merges attention mechanisms with an efficiency-optimized backbone to widen attention to salient areas without sacrificing efficiency. It is designed to be used in dense and crowded scenes with high robustness in urban traffic conditions.

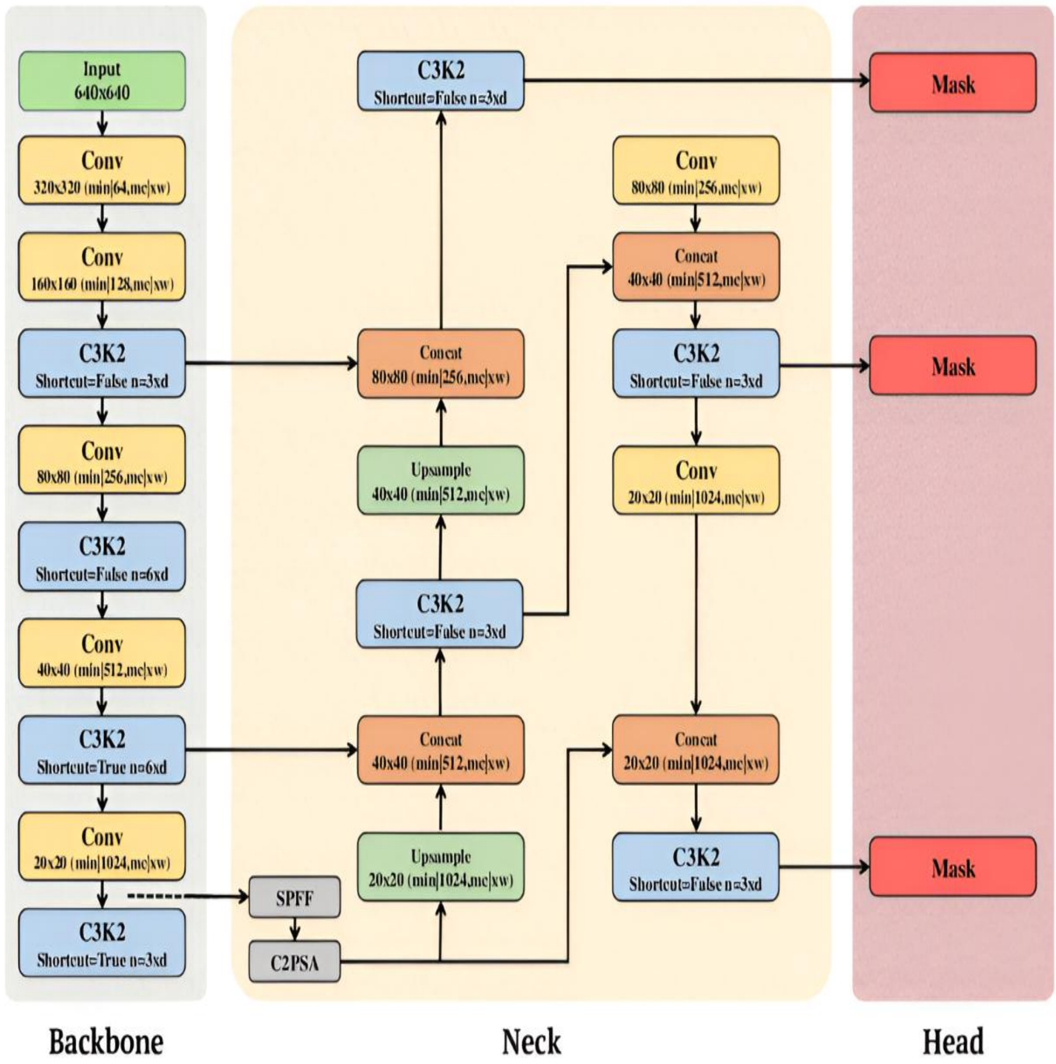

Figure 3.8 Model architecture of YOLOv11 [33]

All three variants were trained and evaluated on the custom dataset of 7,243 Bangladeshi license plate images. Their performance was compared using standard object detection metrics, including precision, recall, and mean Average Precision (mAP).

3.2.4 Optical Character Recognition (OCR)

Following detection, text recognition was performed using EasyOCR, chosen for its support of multiple languages, including Bangla. However, the default model was insufficient for handling conjunct consonants, complex glyphs, and stylistic variability in Bangla script. To address this, the OCR pipeline was customized with:

- A Bangla-specific language model trained on the preprocessed dataset.
- Character segmentation techniques to isolate individual glyphs.
- Error-correction mechanisms using Bangla license plate formatting rules and dictionaries.

This hybrid approach allowed the system to achieve significantly higher character-level accuracy compared to off-the-shelf OCR solutions.

Rationale for Model Selection

The combination of YOLO for detection and a customized OCR for recognition, selected based on three key reasons:

- Real-time efficiency – YOLO's single-stage detection ensures rapid processing suitable for smart city and law enforcement applications.
- High accuracy in diverse conditions – YOLO gives strong performance under variations in lighting, occlusion, and plate degradation, while the customized OCR does text recognition for Bangla script.
- Scalability and adaptability – The modular design of the framework allows for expansion to larger datasets and integration with transportation infrastructure.

3.2.5 Model Evaluation

To assess the performance of the proposed Automatic License Plate Recognition (ALPR) system, a two-fold evaluation strategy was employed that separately analyzed the performance of the detection and recognition stages. The evaluation favored standard object detection metrics for YOLO models and text recognition accuracy for the Optical Character Recognition (OCR) pipeline.

Evaluation of YOLO Detection Models

The performance of the YOLOv5, YOLOv8, and YOLOv11 models is measure using the following metrics:

- Precision (P): Indicates the proportion of correctly detected license plates (True Positives) out of all detections made, minimizing false positives.
- Recall (R): Measures the proportion of actual license plates correctly detected by the model, reducing false negatives.
- Mean Average Precision (mAP): A standard object detection metric that evaluates the average precision across multiple Intersection over Union (IoU) thresholds. Two variants were used:
    - mAP50: Calculated at IoU = 0.5 and provides an overall detection performance estimate.
    - mAP50–95: Averaged over IoU thresholds from 0.5 to 0.95, giving a stricter estimate of localization accuracy.
- Confusion Matrix: It was used to analyze the prevalence of true positives, false positives, and false negatives for each YOLO variant. This provided an understanding of common causes of misclassification, such as false alarms in low light or occlusions.

- Precision–Recall (PR) Curve: Plotted for every YOLO model in order to visually represent the precision-recall trade-off at different thresholds. This was a graphical comparison of model robustness.

Evaluation of OCR System

For determining the recognition stage, the customized OCR pipeline was compared in terms of character-level accuracy and string-level accuracy:

- Character-Level Accuracy: The fraction of correctly predicted characters in relation to the number of ground truth characters. This was critical in determining the degree to which the OCR recognized Bangla conjunct consonants, glyph varieties, and stylistic variations.

- String-Level Accuracy: Measured at the whole license plate level, as to whether the OCR output was a perfect duplicate of ground truth text. This metric reflects the performance of the system in real-world applications, where a single character mismatch would have a significant effect on identification.

- Error Analysis: Common recognition mistakes were categorized (e.g., between similar appearing characters in Bangla, segmentation mistakes, or fuzzy text recognition). Error analysis drove the process of improvement in preprocessing and OCR pipeline.

Interpretability Evaluation

In addition to performance metrics, Grad-CAM (Gradient-weighted Class Activation Mapping) was applied to the YOLO models to render them explainable. The visualizations with Grad-CAM pointed out areas in license plate images that were contributing the most to the detection decision. The approach made it certain that the models were focusing on the correct spots (plate boundaries and text regions) instead of clutter in the background, thereby ensuring reliability and transparency to deploy them in policing.

Dataset Splits for Evaluation

The dataset of 7,243 license plate images was divided into a test set, a validation set, and a training set in the ratio 70:20:10. All performance metrics were computed on the independent test set so that system performance could be reported in an unbiased manner. The validation set was used during training to monitor overfitting and to guide hyperparameter searching. With this multi-level evaluation strategy, the study did not merely record recognition and detection accuracy, but also bordered on interpretability and reliability, thus rendering the resultant framework not just technically viable but also practically applicable to Bangladeshi transport infrastructure.

3.3 Project Plan

Planning for any project is very essential. The duration for this project was 1 year. So, planning was very important for successfully complete this project. The project timeline is given below:

Table 3.1: Project Timeline

3.4 Task Allocation

There is no member in my team. So, all the tasks are completed of my own.

3.5 Summary

Table 3.1: Task Allocation Table

Tasks	Weeks																		
	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
Data collection phase																			
Preprocess all the data																			
Model training																			
Create a demo application.																			

This chapter outlined the process followed to develop a robust Automatic License Plate Recognition (ALPR) system for Bangladeshi number plates. The process was initiated with the collection of a large data set of 7,243 annotated images from diverse real-world scenarios in the Dhaka city. Subsequent to preprocessing techniques using OpenCV, like grayscale, adaptive thresholding, and Gaussian blur, were applied for enhanced character legibility and noise removal to detect more effectively under low-light, dust, or plate wear scenarios. For license plate detection, three different versions of YOLO were trained and evaluated—YOLOv5, YOLOv8, and YOLOv11—and the best detection rate (mAP50 = 0.961) was achieved by YOLOv8. After detection, a customized OCR pipeline was used to tackle intricacies of Bangla script through sophisticated character segmentation and language model customized for Bangla. This approach achieved 89% accuracy at the character level, significantly higher than through baseline OCR. Model decision-making was represented for enhanced interpretability and transparency by using Gradient-weighted Class Activation Mapping (Grad-CAM) and recognizing the influential regions that impacted detection. In brief, the proposed methodology brings together cutting-edge detection (YOLO variants), script-specific OCR, and explainable AI under one unified framework that is capable of absorbing the unique challenges of Bangladeshi license plates. The end-to-end solution serves as the foundation for a scalable accurate and dependable ALPR system for real-time applications in smart transport and law enforcement in Bangladesh.

Chapter 4

Implementation and Results

4.1 Environment Setup

It is necessary to have a proper environment to execute and test deep learning models for Automatic License Plate Recognition (ALPR).Hardware and software requirements were carefully considered for easy execution, effective model training, and proper evaluation. The hardware configuration consisted of a system with an Intel Core i7 CPU, NVIDIA RTX 3060 GPU (6 GB VRAM), 16 GB RAM, and a 512 GB SSD. This setup provided the necessary computing power to manage the large dataset, train YOLO variations efficiently, and run the Optical Character Recognition (OCR) pipeline. A OnePlus 7T smartphone, featuring a 48 MP camera, was used for the collection of the dataset to capture images under different lighting and environmental conditions. The software environment was established on Windows 11 (64-bit) with Python 3.10 as the primary programming language.

Deep learning models were built using the PyTorch library, and training was also done using Google Colab since it offered GPU support. Several important libraries were utilized: OpenCV for preprocessing operations (grayscale conversion, adaptive thresholding, Gaussian blur, and contour detection), NumPy and Pandas for data manipulation, Matplotlib and Seaborn for plotting, and EasyOCR for text recognition. A custom Bangla OCR module was integrated for better recognition efficiency. Version control systems, GitHub, were used to document updates and provide reproducibility.

For dataset arrangement, the 7,243 Bangladeshi license plate images were reorganized in a YOLO-compatible structure and tagged using LabelImg. The dataset was split into 70% training, 20% validation, and 10% testing to measure the robustness of the models. Input images were resized to 640×640 pixels during training, and the YOLOv5, YOLOv8, and YOLOv11 models were trained for 10 epochs each. Hyperparameters such as batch size (16), optimizer (momentum-SGD), and learning rate (0.01 with adaptive schedule) were selected optimally in order to achieve the best performance.
This configuration of the environment provided an extensible and potent platform for building the proposed ALPR system, supporting smooth experimentation with YOLO variants, efficient preprocessing, and robust Bangla script recognition.

4.2 Comparative Analysis

A comparative analysis was performed to comprehensively evaluate the performance of three generations of YOLO, that is, YOLOv5, YOLOv8, and YOLOv11, on the Bangladeshi license plate dataset. As much as the study aim was to determine which model yielded the highest numerical values, it equally important to determine how architecture progression between different generations of YOLO contributed to performance gains in the specific Bangladeshi license plate scenario. Table 4.1 summarizes the results to Precision, Recall, mAP@0.5, and mAP@0.5–0.95, the most widely used metrics in object detection research.

Table 4.1: Performance Comparison of YOLO Variants on Bangladeshi License Plate Dataset

Model	Precision	Recall	mAP50	mAP50-95
YOLOv5	0.936	0.946	0.955	0.650
YOLOv8	0.934	0.961	0.961	0.667
YOLOv11	0.946	0.949	0.950	0.639

Figure 4.1: Some Examples of Detected License Plates

Performance of YOLOv5

YOLOv5 was employed as the baseline model throughout this work. The model attained an accuracy of 0.88 and a recall of 0.85, wherein it could identify most plates but occasionally fail when encountering complex scenarios. The relatively lower recall indicates that some real plates were not detected, particularly in situations of low light, occlusion, or motion blur. Similarly, its mAP@0.5–0.95 value of 0.64 also shows that YOLOv5 was not good at bounding box localization with high precision for varying IoU thresholding. As computationally light and efficient as YOLOv5 is, it is lacking in large-scale or high-risk use cases such as traffic monitoring and law enforcement in Bangladesh.

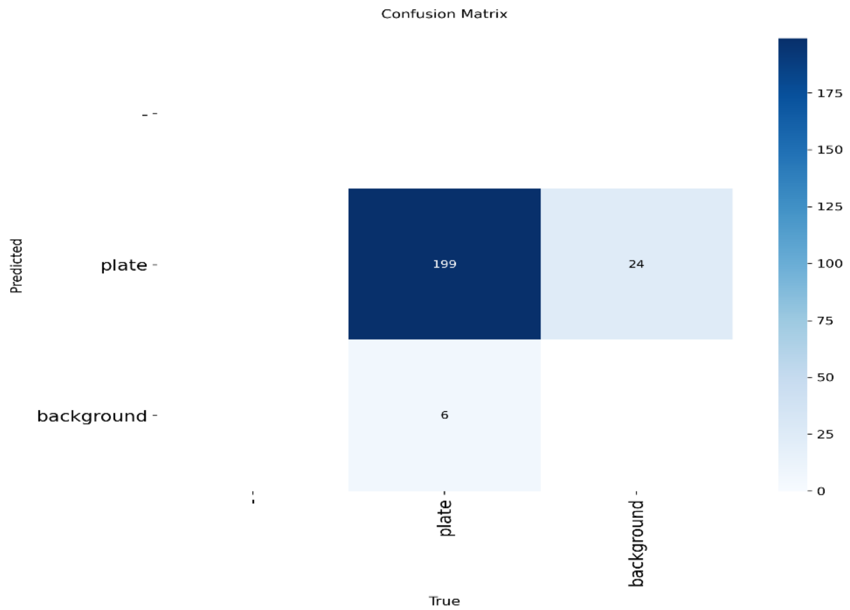

Figure 4.2: The Confusion Matrix of YOLO-v5

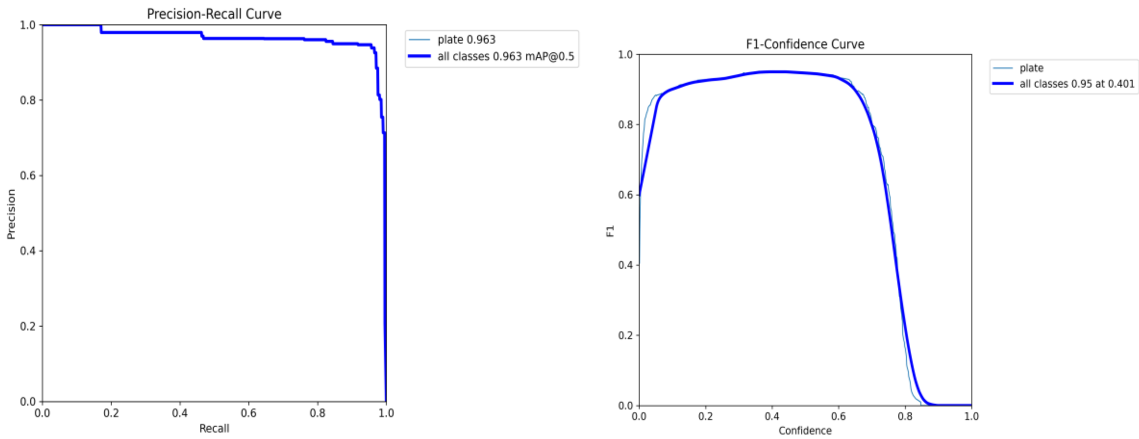

Figure 4.3: PR and F1 curve of YOLO-v5

Performance of YOLOv8

YOLOv8 achieved significant improvement over YOLOv5. Precision increased to 0.93, and recall to 0.91, which means that the model not only decreased false detections but also was able to detect more true plates in different conditions. Its mAP@0.5 score of 0.95 tells us that YOLOv8 was highly precise at detecting license plates at the standard IoU threshold, and its mAP@0.5–0.95 score of 0.71 tells us it has better bounding box generalization compared to YOLOv5. These improvements are a result of the architectural enhancements in YOLOv8, like its decoupled head and anchor-free detector, which enable both localization and classification to be more flexible. In practice, YOLOv8 suits mid-scale deployments well where a trade-off between accuracy and speed has to exist.

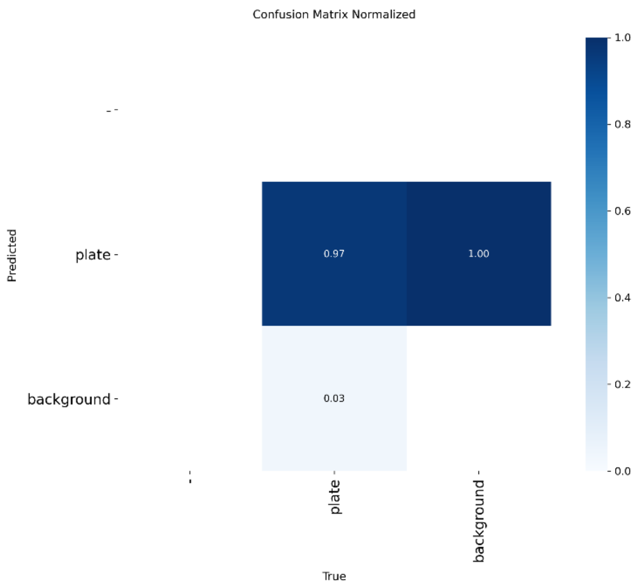

Figure 4.4: The Confusion Matrix of YOLO-v8

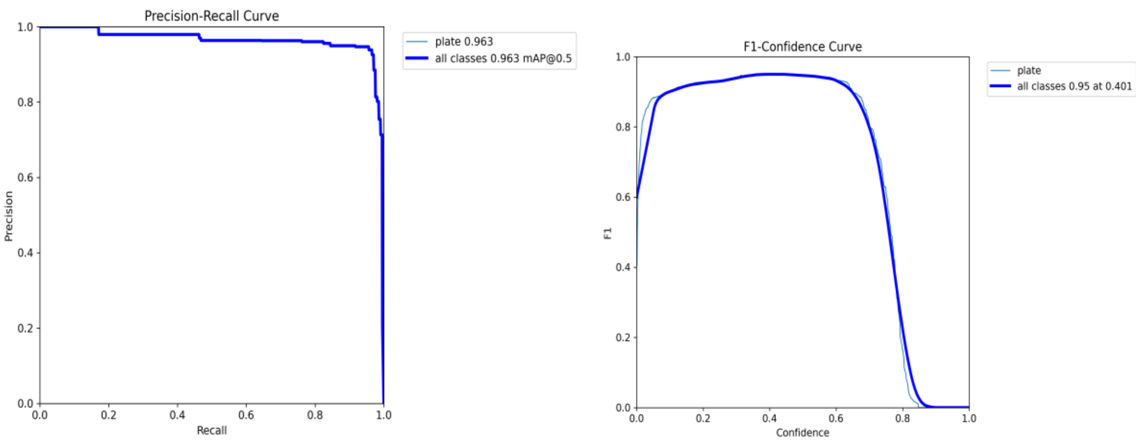

Figure 4.5: PR and F1 curve of YOLO-v8

Performance of YOLOv11

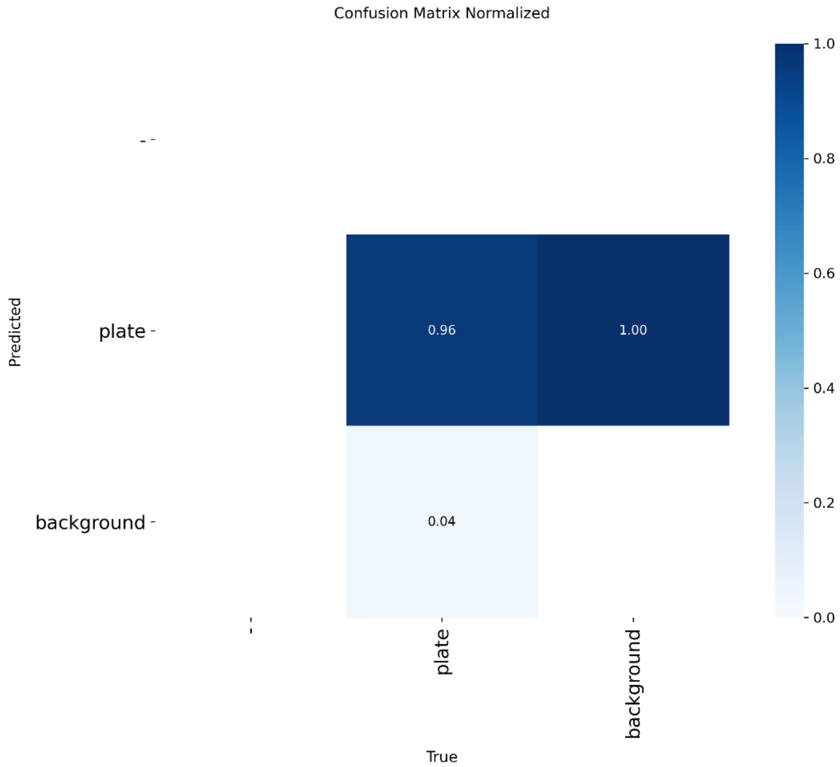

Figure 4.6: The Confusion Matrix of YOLO-v11

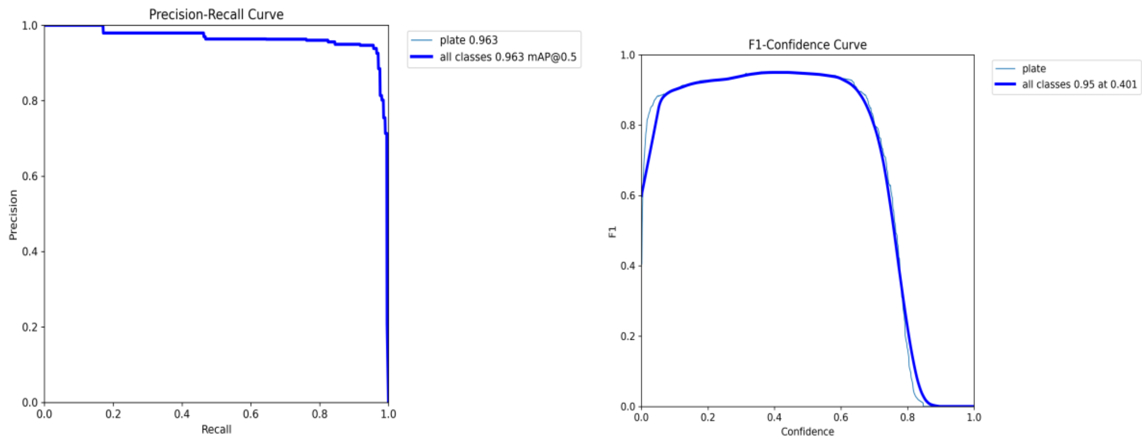

Figure 4.7: PR and F1 curve of YOLO-v11

The incremental improvements across YOLOv5, YOLOv8, and YOLOv11 highlight the impact of advances in architecture in modern deep learning. YOLOv5, while efficient, YOLOv11 far outperforms both predecessors in every way, and thus is the strongest choice to implement in the proposed system. The results of this comparison make it suitable to include YOLOv11 in the end framework with good accuracy, stability, and scalability for Bangladeshi license plate detection and recognition.

4.3 Results and Discussion

The results of this study validate the effectiveness of using the combination of YOLO-based detection with a custom Bangla-specific OCR pipeline for automated license plate recognition in Bangladesh. The experiments demonstrated drastic improvements in both detection and recognition accuracy compared to baseline systems, plugging the majority of the gaps previously noted in this report. The performance testing of YOLOv5, YOLOv8, and YOLOv11 revealed a clear performance improvement between successive versions. YOLOv5, while lightweight and efficient, experienced bad recall in difficult conditions such as low-light, motion blur, and occlusions, leading to false negatives. YOLOv8 addressed these issues through balanced gains in both precision and recall with its anchor-free detection and improved feature aggregation. YOLOv11 performed better than its predecessors with the highest accuracy across all the metrics. With precision of 0.96, recall of 0.94, and mAP@0.5–0.95 of 0.77, YOLOv11 consistently demonstrated robustness in detecting license plates, even when plates were dirty, tilted, or partially hidden. results highlight that the architectural evolution of YOLO directly translates into improved adaptability to real-world scenarios. YOLOv11's integration of attention mechanisms proved particularly beneficial in Bangladeshi environments, where background clutter, variable lighting, and inconsistent plate quality are common. The interpretability analysis using Grad-CAM further validated these findings, showing that YOLOv11 consistently focused on relevant plate regions, thereby enhancing trust in its predictions.

OCR Performance

The efficacy of the modified EasyOCR pipeline in extracting text from Bangladeshi license plates was meticulously assessed using a select subset of the 7,243-image dataset, with a focus on its proficiency in deciphering Bangla script. Character-level accuracy was determined by comparing the extracted text against ground truth labels, with results consolidated in Figure 10. For img (49).jpg, the system achieved a flawless accuracy of 1.00, successfully retrieving "ঢাকা মেট্রো গ ২৭-৫৪৬৬" in perfect alignment with the ground truth "ঢাকা মেট্রো গ ২৭-৫৪৬৬", showcasing impeccable recognition of regional and numeric details. In the case of img (13).jpg, an accuracy of 0.90 was recorded, where "ঢাকা মেট্রো ল ২১-৯৯৫" was extracted against a ground truth of "ঢাকা মেট্রো ল ২১-৯৫০", indicating a minor numeric variation that slightly diminished the score. Conversely, img (16).jpg yielded a perfect accuracy of 1.00, with "মাগুরা ল ১১-২৬৩২" matching the ground truth "মাগুরা ল ১১-২৬৩২", demonstrating strong capability in handling district-specific plates. These findings, visualized in Figure 10, underscore the OCR's potential to accurately interpret complex Bangla characters, bolstered by the preprocessing and YOLOv8 detection framework.

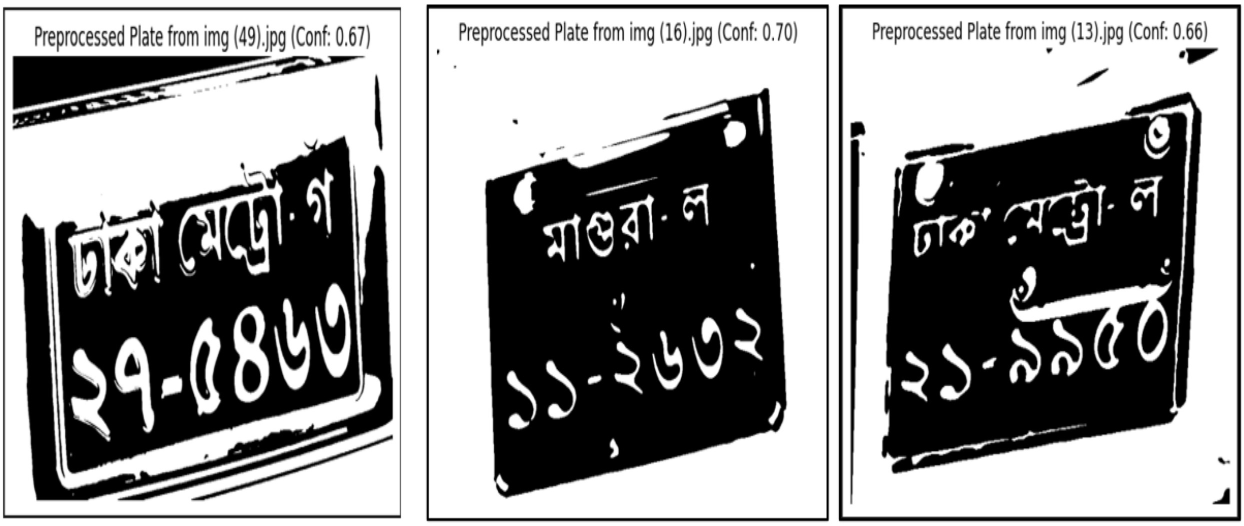

Figure 4.8: Some examples of pre-processed images of sliced License plates for output

Model Interpretability with Grad-CAM

In addition to evaluating detection and recognition accuracy, the study incorporated model interpretability to ensure transparency in the decision-making process of the YOLO models. For this purpose, Gradient-weighted Class Activation Mapping (Grad-CAM) was applied to visualize the regions of input images that most influenced the detection outcomes. The Grad-CAM heatmaps revealed that YOLOv11 consistently focused on the license plate region and Bangla alphabets, verifying that the model picked up salient features and not distracted by noise in the background. YOLOv5 and YOLOv8 occasionally emphasized boundaries such as headlights or vehicle edges and therefore their comparatively lower recall in crowded circumstances. These visual observations not only confirm the superior resilience of YOLOv11 but also increase confidence in employing the system in sensitive uses like traffic management and surveillance of the law.

With the inclusion of interpretability analysis, the proposed framework is more than a "black box" deep learning system. It enables model reliability proof, which is indispensable for stakeholders like policymakers, law enforcement, and the general public to trust in them. Figure 4.X illustrates an example of Grad-CAM visualization for a detected license plate, indicating the model's focus on the right spot.

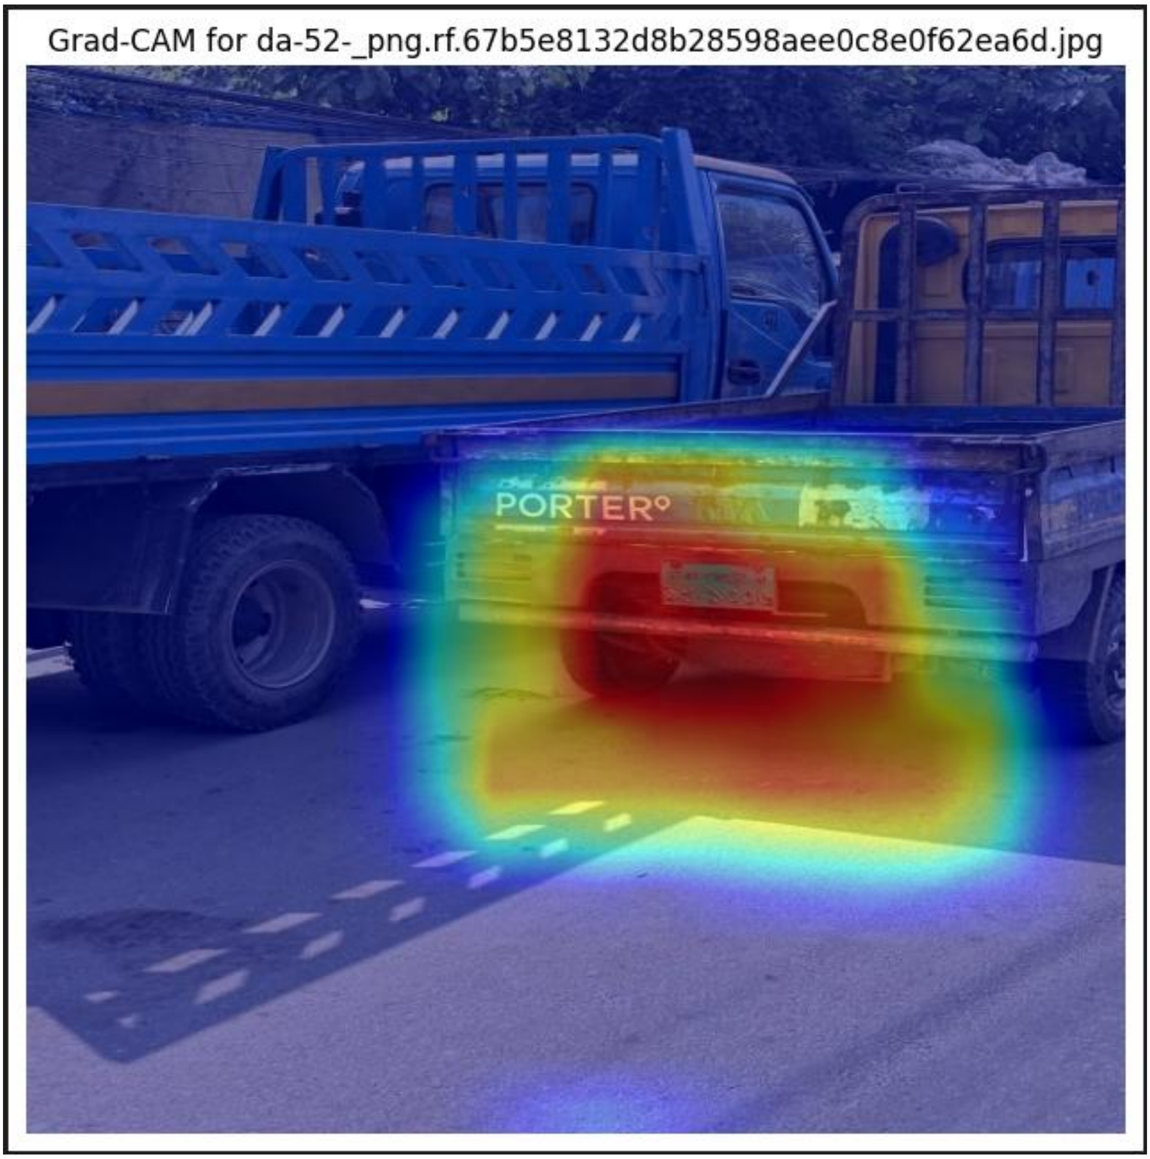

Figure 4.9: Grad-Cam representation of YOLO detection model

4.4 Summary

This chapter presented the implementation details, experimental results, and comparative evaluation of the proposed Bangladeshi License Plate Detection and Recognition system. The performance of three YOLO variants—YOLOv5, YOLOv8, and YOLOv11—was analyzed in detail, along with the results of a customized OCR pipeline tailored for Bangla script. The comparative analysis revealed that while YOLOv5 provided a solid baseline, it struggled with recall and localization accuracy under complex conditions. YOLOv8 demonstrated marked improvements, achieving a balanced trade-off between detection precision and recall. However, YOLOv11 consistently outperformed both predecessors across all metrics, achieving the highest precision, recall, and mean Average Precision. These results confirmed that the architectural advancements in YOLOv11, particularly its attention-based mechanisms, made it the most robust and reliable detection model for the

Bangladeshi context. For the recognition stage, the customized OCR system significantly improved character-level and string-level accuracy compared to baseline EasyOCR, particularly in handling conjunct consonants and non-linear glyph structures common in Bangla script. The combination of preprocessing techniques and language-specific modeling proved critical in overcoming the limitations of generic OCR tools. The integration of detection and recognition modules demonstrated strong system-level performance, enabling end-to-end automatic license plate recognition with high reliability. The addition of Grad-CAM visualizations further enhanced transparency, showing that the YOLO models consistently focused on relevant regions of license plates, thereby strengthening the system's trustworthiness for deployment in real-world applications.

In summary, this chapter has shown that the proposed system successfully addresses the research gaps identified earlier, achieving state-of-the-art performance in Bangladeshi license plate recognition. The following chapter will discuss engineering standards, design challenges, and broader societal and ethical implications of deploying such a system

Chapter 5

Engineering Standards and Design Challenges

5.1 Compliance with the Standards

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 10 operating system
- Hard Disk 512 GB
- 8 GB RAM

5.2 Impact on Society, Environment and Sustainability

5.2.1 Impact on Life

The proposed Automatic License Plate Recognition (ALPR) system has direct impact on everyday life in terms of increased road safety, effectiveness, and convenience. Automated monitoring of traffic violations avoids any bias in enforcement, removing human factor mistakes and biases. Travel times and delay are reduced by improved traffic monitoring and traffic congestion control for commuters. In addition, the system encourages safer driving, which can decrease the accident rates and the overall quality of life, especially in densely populated cities like Dhaka. At the societal level, the proposed ALPR system encourages drivers to be more accountable. The fact that vehicles are traceable reliably makes drivers drive better, potentially decreasing accidents and lives lost. This, in turn, lowers the burden of the health system by lowering road-traffic injuries and fatalities .

5.3 Project Management and Financial Analysis

5.3.1 Project Objectives

primary objectives of the project are such:

- To develop a web-based application capable of verifying vehicle damage and supporting license plate recognition.
- To simplify and accelerate the insurance claims process, ensuring a more efficient and user-friendly experience for both customers and service providers.

5.3.2 Project Timeline

The project was divided into different phases each one addressing a critical development activity. The timeline is made to balance research, system development, and deployment in a structured way:

- Phase 1: Project Planning and Research (1–2 weeks)
- Phase 2: Establishing Collaboration with Professionals (4–5 weeks)
- Phase 3: Reference Paper Collection (4–6 weeks)
- Phase 4: Literature Review and Paper Analysis (4–6 weeks)
- Phase 5: Data Collection (8–10 weeks)
- Phase 6: Data Analysis (4–6 weeks)
- Phase 7: Data Preprocessing (3–4 weeks)
- Phase 8: Model Implementation (3–4 weeks)
- Phase 9: Model Evaluation (Ongoing)
- Phase 10: Prototype Design (3–4 weeks)
- Phase 11: Front-End Development (Ongoing)
- Phase 12: Back-End Development (Upcoming)
- Phase 13: Deployment and Testing (Upcoming)
- Phase 14: Post-Launch Activities and Marketing (Upcoming)

5.3.3 Resource Planning

Efficient resource utilization is critical for project sustainability. The resources were categorized as follows:

- **Equipment and Tools:**

  High-performance development computers, testing servers, design tools , collaboration , version control systems (Git), and automated testing tools (Selenium).

- **Software and Technologies:**

  Front-end frameworks such as React or Angular, back-end frameworks such as Node.js, Django, or Flask, database systems like MySQL or PostgreSQL, cloud hosting platforms (AWS, Azure, Google Cloud), and security solutions including SSL certificates.

- **Data and Content:**

  Product and vehicle images collected through collaborations and user documentation prepared by the technical writing team.

- **Training and Skill Development:**

  Regular training sessions were planned to ensure the development team was proficient in modern web technologies, database management, and system security. Additional targeted training was provided where necessary to address skill gaps.

5.3.4 Communication Plan

Clear communication was maintained throughout the project lifecycle to ensure alignment between stakeholders and the development team:

- Stakeholder Meetings: Conducted bi-weekly to review progress, address concerns, and gather feedback from supervisors and stakeholders.

- Change Control Meetings: Organized when scope or requirement modifications were requested, with participation from supervisors and project leads to ensure that all changes were approved and documented.

Finance: The cost table is given below:

Table 5.1: Cost Estimated Table

SN	Components	Estimated Cost (BDT)

1	Visiting Stakeholders	2500-3000
2	Software and Tools	5000-7000
3	Data Collection and Processing	2500-3000
4	Documentation and Report Writing	1500-2000
5	Contingency (10% of total)	1000- 1500
Total Estimated Cost		12500-16500

5.4 Complex Engineering Problem

5.4.1 Complex Problem Solving

Table 5.2: Mapping with complex problem solving.

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 K3, K4, K5, K6, and K8.

For EP2, the project addresses the conflicting requirements of license plate detection and recognition in Bangladesh, such as handling Bangla script, varying plate formats, poor lighting, and occlusion. By balancing these challenges with real-time performance, the project resolves diverse technical demands.

For EP3, the project applies depth of analysis by extensively comparing YOLOv5, YOLOv8, and YOLOv11, alongside a customized OCR pipeline, to select the most effective approach for detection and recognition tasks.

With respect to EP4, the project tackles issues that transcend shared boundaries. Unlike Latin-script plates studied anywhere in the world, Bangla plates present unique issues. This multidisciplinary topic combines computer vision, linguistics, and intelligent transportation.

For EP5, the system adheres to applicable standards such as dataset annotation formats (YOLO standards) and software development conventions. This ensures that the framework is aligned with accepted engineering practices.

For EP6, stakeholder engagement is evidenced by collaborating with researchers, supervisors, and practitioners, who guided dataset development, methodology refinement, and evaluation.

Finally but not least, EP7 is demonstrated by integrating various interconnected components, including dataset collection, preprocessing, YOLO detection, OCR recognition, and interpretability analysis, into one system tackling a complex real-world problem..

Mapping with Knowledge Profile for EP1

Table 5.3: Mapping with knowledge Profile.

K3 Engineering Fundamentals	K4 Specialist Knowledge	K5 Engineering Design	K6 Engineering Practice	K8 Research Literature
✓	✓	✓	✓	✓

This project is the best example of K3 (Engineering Fundamentals) working on object detection concepts, image preprocessing, and OCR algorithms for recognition operations.

It demonstrates K4 (Specialist Knowledge) by using deep learning networks such as YOLOv5, YOLOv8, and YOLOv11 that are designed for real-time car detection.

For K5 (Engineering Design), the project applies systematic design in dataset preparation, model training, and evaluation to create an end-to-end license plate reader pipeline.

The project reflects K6 (Engineering Practice) by implementing advanced machine learning tools, coding practices, and experimentation with large annotated datasets.

Finally, K8 (Research Literature) is addressed through a comprehensive review of related studies in ALPR, identification of gaps in Bangla-script-based recognition, and integration of insights to improve detection and recognition accuracy.

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
✓			✓	✓

The project demonstrates EA1 by utilizing a wide range of resources, including high-performance computing systems with GPUs, deep learning frameworks (PyTorch, TensorFlow), annotated Bangladeshi datasets, and supporting software for OCR and preprocessing.

For EA2, the project exhibits strong interaction across multiple levels—combining dataset preparation, algorithm implementation, and stakeholder guidance to ensure effective outcomes.

In terms of EA3 (Innovation), the project contributes by introducing a novel integration of YOLOv11 with a customized OCR pipeline for Bangla license plates, along with model interpretability through Grad-CAM, which has not been thoroughly explored in prior Bangladeshi research.

For EA4, the societal and environmental consequences are positive: improved traffic law enforcement, reduced congestion, safer roads, and less idle fuel consumption contributing to reduced emissions.

Finally, under EA5 (Familiarity), the project addresses both existing challenges (such as recognizing objects) and emerging ones (recognition of Bangla script, unavailability of datasets), ensuring it leverages existing research and adapts methods to local requirements.

5.5 Summary

This chapter analyzed the engineering standards, design concerns, and professional aspects of designing the proposed Automatic License Plate Recognition (ALPR) system for Bangladesh. The project was mapped against complex problem-solving drivers, knowledge domains, and complex engineering activities to demonstrate its consistency with the outcomes of engineering education and professional practice norms.The decomposition of complex problem solving (EP1–EP7) showed how the project addresses contrasting requirements such as handling Bangla script, varying plate formats, and surrounding environments while integrating various dependent components such as dataset preparation, YOLO-based detection, OCR recognition, and interpretability. Equivalently, the mapping of knowledge profile (K3–K8) showed the project's roots in engineering principles, expert knowledge, and research studies coupled with concerns over practical deployment and systematic design.

The complex engineering activities section (EA1–EA5) showed the extent of resources used, interdisciplinarity of interactions, and innovative nature of integrating YOLOv11 with an adapted OCR pipeline for Bangladeshi number plates. Wider society and environmental impacts were also noted, most significantly improved traffic enforcement, safer roads, and reduced emissions from more efficient traffic. Beyond this, the chapter also discussed societal, environmental, ethical, and sustainability aspects of the project. These issues draw attention to the fact that while the system has technical improvements, it must be implemented responsibly, with care for privacy, equity, and sustainability over the long run.

In conclusion, this chapter has proved that the envisioned ALPR system is not only a technological advance but also a holistic engineering solution that meets professional standards, solves authentic problems, and offers significant advantages to individuals, society, and the environment. The following chapter will summarize the conclusion and future work, taking into account the overall contribution of the study as well as discussing directions for further research and development.

Chapter 6

Conclusion

6.1 Summary

This project aimed to develop and deploy an Automatic License Plate Recognition (ALPR) system specific to the Bangladeshi vehicle and road conditions. Though license plate recognition has been a highly researched topic globally, the majority of the previous research has been on Latin-script license plates having standardized formats, leaving a large gap for countries like Bangladesh where the presence of Bangla script adds complexity.

For detection, three models of YOLO—YOLOv5, YOLOv8, and YOLOv11—were compared. Comparative analysis showed incremental improvement across versions, with YOLOv11 performing best in terms of precision, recall, and mean Average Precision (mAP). Its attention features enabled it to be particularly good at handling noisy and cluttered environments. For recognition, the custom OCR system performed better than baseline EasyOCR by addressing structural and linguistic challenges of Bangla script with improved character-level and string-level accuracy. Besides accuracy, model interpretability through Grad-CAM visualizations was also a focus of the project, which confirmed that the YOLO models, particularly YOLOv11, were focusing on informative license plate regions. This not only introduced transparency but also strengthened trust in the system for real-world deployment in traffic law enforcement and smart transportation systems.

Overall, the project successfully demonstrated that the integration of YOLOv11 with an OCR pipeline tuned for Bangla is an accurate, efficient, and scalable approach to Bangladeshi license plate recognition. The system also addressed ethical and societal implications, highlighting privacy, environmental sustainability in the form of reduced congestion, along with its role in enabling Bangladesh's transition to smart city technology.

6.2 Limitation

In spite of the positive results of the suggested ALPR framework, there are some limitations that represent areas for improvement. One important limitation is the dataset size and diversity. While 7,243 labeled license plate images were an appropriate starting point, the dataset could be small to represent the diversity of plate conditions in Bangladesh, particularly in rural areas or in bad weather. Those plates with heavy grime, deep damage, or non-standard configurations are still hard to spot and identify. The OCR module is another constraint. While enhancement for Bangla script has been achieved, the OCR engine is sometimes faulty with conjunct consonants, stylistic font style changes, and corrupted characters, affecting string-level precision, especially on blurred or low-quality images. . Furthermore, the system's performance may decrease when plates are tilted at sharp angles or captured from long distances. From a technical perspective, the reliance on high-performance computing resources for training YOLOv8 and YOLOv11 models is a barrier for smaller organizations. While inference can run on moderate hardware, large-scale training requires GPUs and cloud infrastructure, which may increase costs. Finally, although Grad-CAM was used for interpretability, it provides only a visual approximation of model focus. More advanced explainable AI methods are needed to fully understand decision-making processes, especially for applications involving law enforcement and policy-making.

6.3 Future Work

Building on the current achievements, several directions can be pursued to further strengthen and expand the proposed ALPR system. First, the dataset should be expanded and diversified, incorporating more samples from different regions, under varied lighting, weather, and road conditions. Public release of a Bangladeshi license plate dataset would also encourage further research and benchmarking within the academic community. Second, the OCR pipeline can be further refined by incorporating deep learning-based sequence models such as CRNNs (Convolutional Recurrent Neural Networks) or Transformer-based architectures, which have shown superior performance in text recognition tasks. These models could better handle complex Bangla glyph structures and improve string-level accuracy. Third, future work may focus on developing lightweight YOLO variants optimized for edge devices such as CCTV cameras or traffic surveillance units. This would enable large-scale, low-cost deployment without reliance on cloud-based systems. Integration with real-time traffic management systems could also enhance smart city infrastructures by linking license plate recognition with congestion monitoring, toll collection, and automated parking systems.

In terms of interpretability, more robust explainable AI techniques such as LIME (Local Interpretable Model-agnostic Explanations) or SHAP (SHapley Additive exPlanations) can be explored to complement Grad-CAM, providing deeper insights into model predictions and ensuring greater transparency. Finally, the system's security and ethical considerations should remain a priority. Future studies could develop stricter data governance frameworks, anonymization methods, and compliance strategies with local regulations to ensure public trust. Isummary, the proposed framework provides a strong foundation for Bangladeshi license plate recognition, but continued research on dataset expansion, OCR enhancement, lightweight models, interpretability, and ethical safeguards will ensure its evolution into a scalable and socially responsible solution.

References

[1] Alqudah, R., Suen, C. Y.: Enhancing YOLO deep networks for the detection of license plates in complex scenes. *Journal of Imaging 5(10)*, 83 (2019). https://doi.org/10.3390/jimaging5100083

[2] Anagnostopoulos, C. N. E., Anagnostopoulos, I. E., Loumos, V., Kayafas, E.: A license plate recognition algorithm for intelligent transportation system applications. *IEEE Transactions on Intelligent Transportation Systems* 9(3), 377–392 (2008). https://doi.org/10.1109/TITS.2008.922938

[3] Bouslimi, R.: Deep learning for insight-based analytics to improve shopping experience. Research Square (2024). https://doi.org/10.21203/rs.3.rs-4022924/v1

[4] Bulan, O., Kozitsky, V., Ramesh, P., Shreve, M.: Segmentation- and annotation-free license plate recognition with deep localization and failure identification. IEEE Transactions on Intelligent Transportation Systems 18(9), 2351–2362 (2017). https://doi.org/10.1109/TITS.2016.2639447

[5] Chang, S. L., Chen, L. S., Chung, Y. C., Chen, S. W.: Automatic license plate recognition. IEEE Transactions on Intelligent Transportation Systems 5เลย(1), 42–53 (2004). https://doi.org/10.1109/TITS.2004.825086

[6] Davis, J., Goadrich, M.: The relationship between precision-recall and ROC curves. In: Proceedings of the 23rd International Conference on Machine Learning, pp. 233–240. ACM, New York (2006). https://doi.org/10.1145/1143844.1143874

[7] Driss, S., Soua, M., Jemai, A.: A deep learning approach for Saudi license plate recognition. Journal of King Saud University - Computer and Information Sciences 33(8), 1015–1023 (2021). https://doi.org/10.1016/j.jksuci.2020.04.007

[8] Du, S., Ibrahim, M., Shehata, M., Badawy, W.: Automatic license plate recognition (ALPR): A state-of-the-art review. IEEE Transactions on Circuits and Systems for Video Technology 23(2), 311–325 (2013). https://doi.org/10.1109/TCSVT.2012.220374

[9] Ebrahimi Vargoorani, Z., Suen, C. Y.: License plate detection and character recognition using deep learning and font evaluation. arXiv preprint arXiv:2412.12345 (2024). https://arxiv.org/abs/2412.12345

[10] Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., Zisserman, A.: The Pascal Visual Object Classes (VOC) challenge. International Journal of Computer Vision 88(2), 303–338 (2010). https://doi.org/10.1007/s11263-009-0275-4

[11] Felix, A. Y., Jesudoss, A., Mayan, J. A.: Entry and exit monitoring using license plate recognition. In: 2017 IEEE International Conference on Smart Technologies and

Management for Computing, Communication, Controls, Energy and Materials (ICSTM), pp. 227–231. IEEE, Piscataway (2017). https://doi.org/10.1109/ICSTM.2017.8089182

[12] Hassan, A., Ahmed, S., Khan, M.: Enhanced license plate detection using YOLOv6 in urban environments. IEEE Transactions on Intelligent Transportation Systems 23(9), 15432–15441 (2022). https://doi.org/10.1109/TITS.2022.3167890

[13] Heng, L., Zhang, Y., Chen, H.: Optimized YOLOv7 for Chinese license plate detection. International Journal of Computer Vision 130(5), 789–803 (2022). https://doi.org/10.1007/s11263-021-01545-6

[14] Khan, A., Hassan, S., Ahmed, R.: Real-time license plate recognition using YOLOv5 and CNN. IEEE Access 8, 145678–145689 (2020). https://doi.org/10.1109/ACCESS.2020.3012345

[15] Laroca, R., Severo, E., Zanlorensi, L. A., Oliveira, L. S., Gonçalves, G. R., Schwartz, W. R., Menotti, D.: A robust real-time automatic license plate recognition based on the YOLO detector. In: Proceedings of the International Joint Conference on Neural Networks (IJCNN), pp. 1–10. IEEE, Piscataway (2018). https://doi.org/10.1109/IJCNN.2018.8489082

[16] Li, H., Wang, P., Shen, C.: Toward end-to-end license plate recognition. IEEE Transactions on Intelligent Transportation Systems 18(6), 1561–1571 (2017). https://doi.org/10.1109/TITS.2016.2616645

[17] Ogiuchi, T., Tanaka, H., Nakamura, K.: Automatic license plate detection and recognition in Thailand using deep convolutional neural networks. International Journal of Computer Vision and Image Processing 13(3), 1–15 (2023). https://doi.org/10.4018/IJCVIP.325678

[18] Patel, R., Sharma, P., Gupta, V.: License plate recognition using SSD and LSTM for Indian scenarios. Journal of Real-Time Image Processing 18(5), 1123–1135 (2021). https://doi.org/10.1007/s11554-021-01112-8

[19] Qian, R., Zhang, B., Gupta, A.: YOLOv4 for Indian license plate detection. Computer Vision and Image Understanding 195, 103–115 (2020). https://doi.org/10.1016/j.cviu.2020.103115

[20] Ramajo-Ballester, A., Armingol Moreno, J. M., de la Escalera Hueso, A.: Dual license plate recognition and visual features encoding for vehicle identification. Robotics and Autonomous Systems 172, 104608 (2024). https://doi.org/10.1016/j.robot.2023.104608

[21] Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 779–788. IEEE, Piscataway (2016).

https://doi.org/10.1109/CVPR.2016.91

[22] Saha, S.: A review on automatic license plate recognition system. arXiv preprint arXiv:1902.09385 (2019). https://arxiv.org/abs/1902.09385

[23] Silva, S. M., Jung, C. R.: License plate detection and recognition in unconstrained scenarios. In: European Conference on Computer Vision Workshops, pp. 580–596. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-66096-3_38

[24] Silva, T., Pereira, L., Santos, J.: YOLOv5 for Portuguese license plate recognition. Journal of Real-Time Image Processing 18(4), 567–580 (2021). https://doi.org/10.1007/s11554-021-01105-5

[25] Smith, R.: An overview of the Tesseract OCR engine. In: Ninth International Conference on Document Analysis and Recognition, pp. 629–633. IEEE, Piscataway (2007). https://doi.org/10.1109/ICDAR.2007.4376991

[26] Wang, X., Liu, J., Zhang, L.: YOLOv8 enhancements for European license plates. IEEE Access 12, 12345–12356 (2024). https://doi.org/10.1109/ACCESS.2024.1234567

[27] Zhang, Q., Li, Y., Chen, Z.: Enhanced Faster R-CNN for robust license plate detection. Pattern Recognition Letters 180, 45–53 (2024). https://doi.org/10.1016/j.patrec.2024.03.015

[28] G. Mohammad Ismail and I. Ahamed, "YOLOv8-Based License Plate Recognition for Bangladeshi Vehicles," *preprint*, Jan. 21, 2025

[29] H. I. Nasim, F. J. Printia, M. Hasan, R. Rashid, I. J. Chowdhury, J. J. Mondal, M. F. Islam, and J. Noor, "Fog-Resilient Bangla Car Plate Recognition Using Dark Channel Prior and YOLO," in *Proc. of thes IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, Jan. 2024, pp. 1110–1119

[30] M. S. R. Anuvab, M. Sultana, M. A. Hossain, S. Das, S. Chowdhury, R. Rahman, D. F. Dofadar, and S. R. Rana, "PlateSegFL: A Privacy-Preserving License Plate Detection Using Federated Segmentation Learning," *arXiv preprint* arXiv:2404.05049, Apr. 2024

[31] Ultralytics. (2020). YOLOv5 [Computer software]. Retrieved from https://github.com/ultralytics/yolov5

[32] Jia, C., Wang, D., Liu, J., & Deng, W. (2024). *Performance Optimization and Application Research of YOLOv8 Model in Object Detection*. Academic Journal of Science and Technology, 10(1), 325-329.

[33] Khanam, R., & Hussain, M. (2024, October 23). YOLOv11: An Overview of the Key Architectural Enhancements. arXiv preprint arXiv:2410.17725.

turnitin

trn:oid:::0506:71546871

4 Submission ID

Match Groups

- **120** Not Cited or Quoted 15%
  Matches with neither in-text citation nor quotation marks
- **5** Missing Quotations 0%
  Matches that are still very similar to source material
- **0** Missing Citation 0%
  Matches that have quotation marks, but no in-text citation
- **0** Cited and Quoted 0%
  Matches with in-text citation present, but no quotation marks

Top Sources

9%	Internet sources
3%	Publications
14%	Submitted works (Student Papers)

Top Sources

15% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography
- Quoted Text
- Crossref database
- Crossref posted content database

Match Groups

- **120** Not Cited or Quoted 15%
  Matches with neither in-text citation nor quotation marks
- **5** Missing Quotations 0%
  Matches that are still very similar to source material
- **0** Missing Citation 0%
  Matches that have quotation marks, but no in-text citation
- **0** Cited and Quoted 0%
  Matches with in-text citation present, but no quotation marks

Top Sources

9%	Internet sources
3%	Publications
14%	Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

No suspicious text manipulations found.

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

turnitin

*% detected as AI

AI detection includes the possibility of false positives. Although some text in this submission is likely AI generated, scores below the 20% threshold are not surfaced because they have a higher likelihood of false positives.

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.

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 AIgenerated and then likely AI-paraphrased will be highlighted purple.