



AIoT for Road Safety: Unified Vehicle Speed and License Plate Recognition in Bangladesh

Abdullah Al Noman¹, Nushrat Jahan Mila¹, Abdullah Al Mamun²,
Dipta Chandra Banik³, Mridul Banik⁴, and Jia Uddin^{5*}

¹ Department of Computer Science and Engineering, Daffodil International University, Dhaka-1216, Bangladesh

² Department of Computer Science, Dhaka University of Engineering and Technology, Dhaka-1216, Bangladesh

³ Department of Computer Science, Dhaka International University, Dhaka-1216, Bangladesh

⁴ Department of Computer Science, Colorado State University, United States, Fort Collins, CO, USA

⁵ AI and Big Data Department, Woosong University, Korea

{noman15-5387, mila15-5758}@diu.edu.bd, mamun.duet.bd@gmail.com,
Diptabanik65@gmail.com, mridul.banik23@alumni.colostate.edu,
Jia.uddin@wsu.ac.kr*

Abstract. Road safety in Bangladesh remains a critical challenge with overspeeding and limited enforcement contributing to thousands of accidents each year. Most of the existing ANPR research focuses on plate detection alone, with none of them linking vehicle identity with driving behavior. In this paper, we bridge this important gap by proposing an AIoT integrated framework that combines Bangla license plate recognition with real-time vehicle speed estimation. Our system uses a YOLOv8-based detector combined with EasyOCR trained over 393 annotated Bangla plates, complemented with heavy augmentation to make up for dataset scale. More than 4,000 vehicle events have been collected using a rig of dual infrared IoT sensors for the measurement of speed, where every incident gets synchronized into a unified record including plate text, speed, timestamp, and location, securely transmitted to the cloud server. The experimental evaluation using our methodology achieved 96.2% mAP for plate detection, 1.52% CER, and 5.53% WER in OCR, and a 2.8% error margin for IoT speed estimation with respect to the radar benchmark. The integrated system precisely flagged overspeeding and generated actionable evidence. Given the low-cost architecture and IoT connectivity, the proposed framework shows immense potential for deployment on highways, school zones, and urban corridors to support smart city initiatives and enhance national road safety.

Keywords: Computer Vision, YOLOv8, EasyOCR, Internet of Things (IoT), AIoT Framework, Intelligent Transportation Systems.

1 Introduction

Road safety in Bangladesh is still a crucial public concern, as rapid urbanization and an increasing number of vehicles continue to strain the existing traffic management systems. According to the Bangladesh Road Transport Authority, over 20,000 road accidents are reported every year; speeding has been recognized as a leading cause of death in most of them. Manual traffic enforcement is severely restrained due to limited manpower, corruption vulnerabilities, and uneven monitoring across densely populated regions—a trend that underlines the requirement for automated, scalable, and tamper-resistant surveillance technologies.

While ANPR has been widely adopted all over the world, Bangladeshi license plates pose additional challenges for recognition due to the characteristics of Bangla script, inconsistent fonts, varying illuminations, and partial occlusions. Preliminary approaches based on OCR were not that accurate, while deep learning methods, especially YOLO-based detectors, present much better results in state-of-the-art plate localization and recognition. The BanglaPlate study [1] established the first strong benchmark for Bangla ANPR using YOLOv8 and EasyOCR, achieving 96.2% detection accuracy, 1.52% CER, and 5.53% WER.

Parallel work with IoT-based traffic monitoring illustrates the pragmatic use of low-cost infrared sensors for effective speed sensing. However, systems in this domain are generally stand-alone and do not link driver behavior with vehicle identity; thus, they become less useful for enforcement and analytics.

The paper, therefore, proposes an integrated AIoT framework that integrates YOLOv8 + EasyOCR-based Bangla license plate recognition with a dual-IR sensor IoT speed detection module. Every event of a vehicle is automatically compiled into a unified digital record comprising plate text, speed, timestamp, and location information, securely transmitted to a cloud server for further analysis. Unlike previous standalone solutions, the proposed framework offers an affordable, context-aware, and scalable platform suitable for highways, school zones, and urban corridors, ultimately supporting smart city initiatives and strengthening road safety in Bangladesh.

2 Literature Review

Automatic Number Plate Recognition is widely used in traffic enforcement, toll collection, and intelligent transportation systems. Early ANPR relied on handcrafted features and traditional image processing; recent advances in deep learning—especially those based on YOLO-based object detectors—have significantly improved real-time license plate localization and recognition performance.

The introduction of YOLOv8 with EasyOCR on a dataset of about 400 annotated Bangla plates collected over Dhaka marked a significant lead for Bangladeshi ANPR in the BanglaPlate study [1]. It successfully resolved script-specific complexities such as diacritics, ligatures, and variable font structures. It outperformed earlier pipelines, such as YOLOv4 with Tesseract, which reported higher Character Error Rates and lower detection precision. Follow-up studies explored hybrid YOLOv5-transformer

models for tilted or difficult plates, for example, DONUT and LayoutLM [2, 3], and using larger synthetic–real mixed datasets for better generalization. Low-light conditions, motion blur, rain, and the diversity of commercial vehicles remain areas of concern.

Parallel efforts in IoT-based traffic systems utilize IR, PIR, and ultrasonic detectors for keeping a watch on speed, congestion, and vehicle flow. Although there has been particularly strong potential shown by the ITMSs, most of the implementations focus on Latin or Arabic license plates without incorporating Bangla script recognition. In addition, the IoT speed systems currently available mostly operate separately and do not relate the behavior to the identity of the vehicle, which diminishes their usability value in enforcement.

In a nutshell, though ANPR-based YOLO and current OCR methods have achieved very strong results for Bangla plates, pinpointing BanglaPlate [1] to be the most reliable benchmark, no prior art has integrated Bangla license plate detection with IoT-based speed recognition under one unified AIoT framework. This study covers that lacuna by integrating a state-of-the-art vision pipeline with a dual-IR sensor IoT system and cloud synchronization, thus allowing for a complete and scalable intelligent solution to road safety monitoring in Bangladesh.

3 Methodology

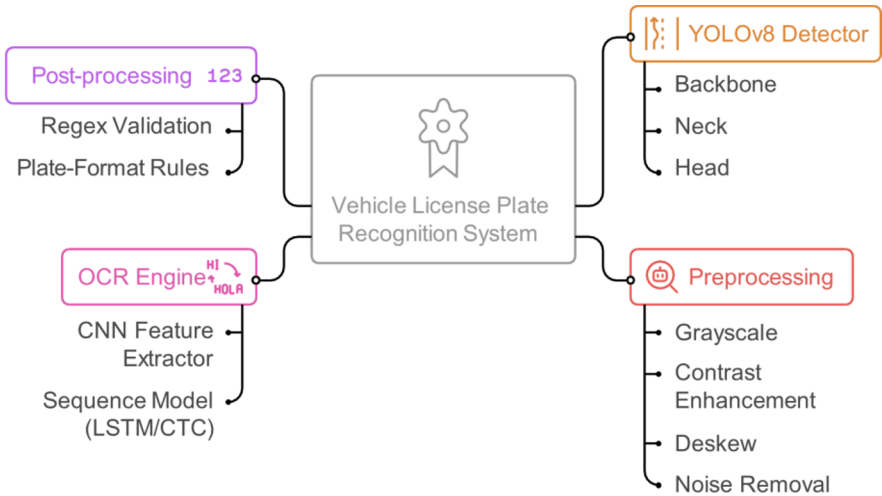


Fig. 1. Overall AIoT methodology workflow for comprehensive traffic monitoring.

The proposed framework is a unified AIoT-based traffic monitoring system integrated with real-time vehicle speed detection and Bangla license plate recognition. The system captures every approaching vehicle using two parallel modalities. As illustrated in Figure 1: Overall AIoT Workflow, the IR sensor-based IoT subsystem records the timing information needed for speed estimation, while its parallel camera subsystem captures

the images that undergo detection and OCR. Both streams are interlinked through a shared event identifier, enabling the creation of a complete and actionable traffic violation record, which later gets transmitted to the cloud for centralized analytics and storage.

3.1 Data Collection and Preprocessing

Data acquisition followed a synchronized dual-source design: image collection and IoT sensor event logging. A total of 393 license plate images from diverse environments such as university parking areas, arterial roads, semi-urban lanes, and highway on/off ramps have been collected. Intentionally, these captured variations in lighting and environmental conditions include daylight, nighttime with artificial illumination, rainy weather, shadows, motion blur, and slight occlusions. Such diversity was chosen to compensate for the relatively small dataset size by enhancing model robustness. Although 393 annotated images seem limited, this data collection has emphasized quality and variation rather than sheer volume, and extensive augmentation techniques were followed to compensate for the dataset scale—a common approach followed in low-resource ANPR research. All images were then manually annotated by drawing bounding boxes on the license plate region. Figure 2: Sample Dataset Image illustrates an example raw vehicle frame with the plate section highlighted.



Fig. 2. Sample raw image from the sampled dataset.

Table 1. Sample speed events dataset.

Event id	Sensor t1	Sensor t2	Distance m	Speed kmh	Camera id	Location id
ev_0001	2025-07-11 10:02:14.010	2025-07-11 10:02:14.285	10.0	126.3	cam_A	Dhaka highway
ev_0002	2025-07-11 10:05:44.122	2025-07-11 10:05:44.612	10.0	72.0	cam_A	Dhaka highway

Simultaneously, IoT speed-sensing data was collected using a dual infrared break-beam setup connected to a microcontroller. Each time a vehicle crossed IR Sensor 1 and then IR Sensor 2, highly accurate timestamps were recorded, enabling calculation of vehicle speed over the known baseline distance of 10 meters. More than 4,000 such sensor-triggered events were collected, making the IoT portion of the dataset much larger than the visual dataset, with high temporal granularity. Each record was assigned a unique event_id to ensure perfect alignment between sensor-generated data and

camera-captured frames. Table 1: Sample Speed Events Dataset presents representative entries, including timestamps, calculated speeds, and location identifiers for validation.

Table 2. Sample license plate recognition dataset.

Image id	Event id	File path	Plate text pred	Plate conf	Ocr conf	lighting	weather
img_0001	ev_0001	imgs/0001.jpg	ঢাকা মেট্রো-গ ১৫-৫৮৭৭	0.95	0.91	day	clear
img_0002	ev_0002	imgs/0002.jpg	ঢাকা মেট্রো-খ ২২-৩৪৫৬	0.93	0.89	day	cloudy

The image data were preprocessed by passing them through the Roboflow pipeline, resizing them to 640×640 pixels to meet the requirements of YOLOv8. Normalization of annotation coordinates was performed, along with different augmentation techniques: horizontal flip, brightness up to ±25%, adding Gaussian noise, rotation up to ±15° in a controlled manner, and the addition of blur. Other optimizations included grayscale conversion of detected plate crops, adaptive histogram equalization for enhancing contrast, perspective correction in case of skewed plates, and morphological techniques for noise removal.

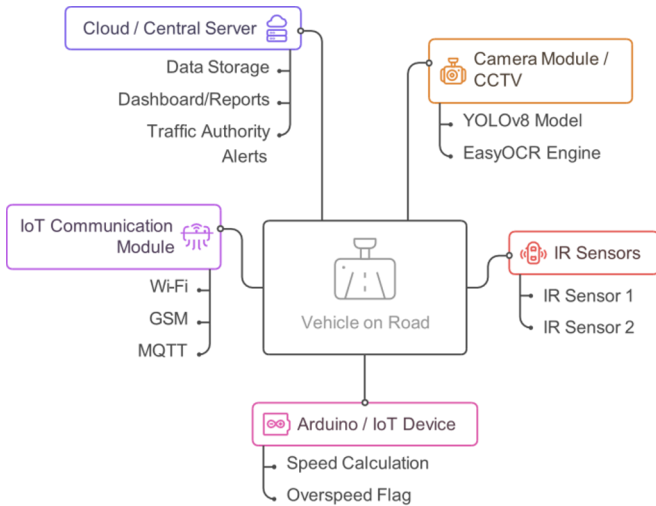


Fig. 3. YOLOv8 + EasyOCR pipeline for license plate detection and recognition.

IoT data preprocessing ranged from filtering out improbable readings, for example, speeds less than zero or greater than 300 km/h, which were considered resulting from faulty triggers. Also, timestamp normalization to ISO-8601 was done for cross-system consistency. A debounce filter removed events that were triggered within 50 ms to prevent duplication of records. Following the cleaning of IoT logs, these were combined with the output of OCR using the shared event_id. Outputs of OCR were kept in Table 2: Sample License Plate Recognition Dataset, comprised of recognized predicted plate text, OCR confidence, and contextual attributes such as lighting and weather conditions. After the preparation of both streams, they were integrated into one violation dataset, as depicted in Table 3: Combined Violation Dataset, made up of Predicted plate text, vehicle speed, violation flags, and timestamp and location identifiers. This unified

dataset forms the backbone of the analytics pipeline for correct identification of overspeeding vehicles.

Figure 3: YOLOv8 + EasyOCR Pipeline presents the entirety of the detection and OCR process visually to illustrate how each vehicle image goes through raw input to final plate text extraction.

Table 3. Sample violation dataset.

Violation id	Event id	Plate text pred	Speed kmh	Speed limit	Violation flag	timestamp	Location id
v_0001	ev_0001	ঢাকা মেট্রো-গ ১৫-৫৮৭৭	126.3	80	Yes	2025-07-11 10:02:14.285	Dhaka highway
v_0002	ev_0002	ঢাকা মেট্রো-খ ২২-৩৪৫৬	72.0	80	No	2025-07-11 10:05:44.612	Dhaka highway

3.2 Modeling and COR Pipeline

The architecture for the proposed detection and recognition is two-tiered: first, YOLOv8 for license plate localization, followed by a model using a CSPDarknet backbone combined with a feature pyramid network for extracting multiscale features, while having a lightweight detection head to generate the bounding boxes along with their confidence scores. Finally, nonmaximum suppression eliminates overlapping detections, and slight padding is applied while cropping the chosen bounding boxes to retain character boundaries. This gives high accuracy regardless of variation in lighting, angles, and motion blur, as shown earlier in Figure 3.

Table 4. Sample license plate recognition dataset.

Component	Description	Purpose
Arduino Uno/ESP32	Microcontroller board	Core control and event management
IR Sensors (×2)	Infrared beam interruption	Vehicle detection and timing
LCD Display	16×2 character screen	Local speed/alert display
Buzzer	Audible alert	Warning for overspeeding vehicles
GSM/Wi-Fi Module	Communication module	Transmits violation data to the cloud
Power Supply	Battery/adapter	Provides system power

In the second phase, the OCR subsystem processes the cropped plate images. The OCR pipeline starts with connected component analysis in order to identify prospective character regions, followed by convolutional feature extraction. A recurrent neural network with CTC loss was used to decode both Bangla and Latin characters in a robust way to deal with the complex ligatures and diacritics in Bangla script. Such recognition performance is further improved due to the preprocessing described earlier: grayscale conversion, contrast stretching, and deskewing of images. There is a regex-based validation layer in the system that enforces the rules of Bangladeshi plate formatting, correcting common misclassifications, such as those involving similar Bangla numerals. OCR outputs, such as text predictions and their confidence scores, are then stored and later merged with IoT speed records. This combined pipeline ensures the overall robustness of localization and recognition across diverse field environments and hardware constraints that make the approach feasible for real-world traffic systems.

3.3 Hardware, Local Processing, and Cloud Integration

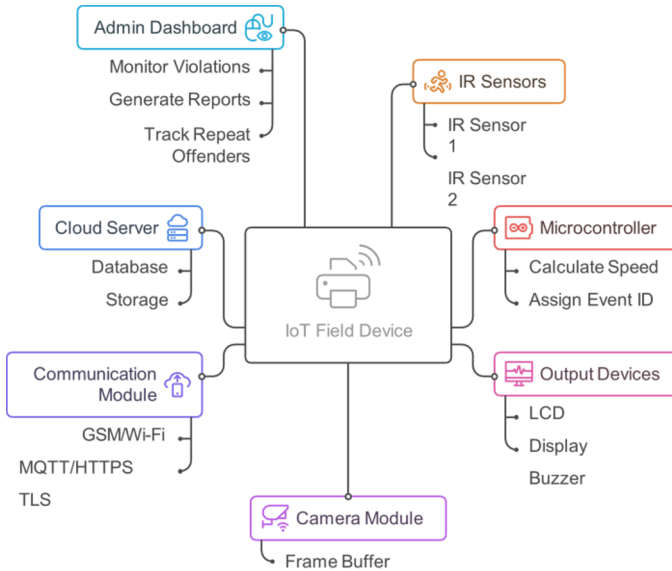


Fig. 4. IoT system architecture and cloud communication flow.

The IoT subsystem provides the basis for the speed measurement and immediate violation flagging. It is built using an Arduino Uno or ESP32 microcontroller connected with two IR break-beam sensors, a 16×2 LCD display, a buzzer, and a GSM/Wi-Fi communication module. A vehicle crossing the first and second beams would give two timestamps, t_1 and t_2 , respectively, to the microcontroller. Speed is then calculated by the latter using the known baseline distance and displayed locally on the LCD. If the measured speed is greater than the threshold value set, a real-time audible alert is triggered through the buzzer. The entire hardware arrangement and workflow are presented in Figure 5: IoT Field Device and Workflow Diagram, showing sensor placement, microcontroller connections, display units, and output generation. Each event is encapsulated with its event_id, calculated speed, timestamps, and location code and sent to the cloud via secure MQTT/HTTPS protocols. Figure 4: IoT–Cloud Architecture shows the communication of the microcontroller with the server where the incoming records are verified and matched with the vision-based OCR output.

The system maintains the aggregated records of violations on the cloud side for real-time monitoring of the incidents of overspeeding, traffic flow statistics, and tracking of repeat offenders. The various components used in the prototype are listed in Table 4: IoT Hardware Components, which describes the functionality of each module used in the system; from event management to power delivery. Event_id matching aligns the IoT-generated data and OCR outputs, creating a fully integrated violation log that contains the identity and behavior of the vehicle. Such a fusion of speed sensing and ANPR allows automated enforcement and provides a basis for scalable intelligent transport systems in Bangladesh.




Fig. 5. Complete IoT field device: speed measurement and local display.

4 Results and Discussion

Table 5. Experimental setup for model training and iot deployment.

Component	Specification/Description
CPU	Intel Core i7, 3.4 GHz
GPU	NVIDIA GTX 1660, 6 GB VRAM
Memory	16 GB RAM
Software	Python 3.9, PyTorch, Roboflow API
IoT Microcontroller	Arduino Uno R3
Sensors	Dual IR beam sensors, 10 m apart
Display & Alert	16x2 LCD, Buzzer
Communication Module	GSM/Wi-Fi (MQTT/HTTPS)

The integrated AIoT platform composing the YOLOv8-based ANPR system and a dual-infrared sensor speed measurement rig was evaluated through vision experiments on 393 annotated license-plate images and IoT experiments on more than 4,000 field-collected sensor events. The complete experimental environment used for both model training and device deployment is summarized in Table 5: Experimental Setup for Model Training and IoT Deployment, which lists the computing hardware, including an Intel i7 CPU, NVIDIA GTX 1660 GPU, and 16 GB RAM; the software framework, including Python, PyTorch, and Roboflow; and IoT hardware components such as an Arduino Uno, dual IR sensors, an LCD display, a buzzer, and a GSM/Wi-Fi communi-

cation module. This table underlines the hybrid nature of the system, where high-performance computing makes model training possible, but low-cost microcontroller-based hardware allows scaling field deployment.

Detection and recognition quality was measured using standard metrics. For detection we employed mean Average Precision at IoU 0.5 (mAP@0.5), a summary measure of precision–recall behavior over confidence thresholds. Precision and recall are given by:

$$Precision = \frac{TP}{TP+FP} \quad (1)$$

$$Recall = \frac{TP}{TP+FN} \quad (2)$$

Where TP true positives, FP false positives and FN false negatives. The harmonic mean of precision and recall is the F1 score:

$$F1\ Score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (3)$$

For OCR evaluation we used Character Error Rate (CER) and Word Error Rate (WER), which are given by:

$$CER = \frac{S_c + I_c + D_c}{N} \quad (4)$$

$$WER = \frac{S_w + I_w + D_w}{N} \quad (5)$$

Where S, D, I denote substitutions, deletions, and insertions respectively, N is the character count, and WWW the word count in the ground truth. These metrics summative quantify token-level and word-level accuracy of the recognition engine.

IoT speed estimation uses the timestamp difference of two IR beams; vehicle speed V is computed as:

$$V = \frac{d}{t_2 - t_1} \times 3.6 \quad (6)$$

Where d is the baseline distance in meters and t1, t2, t1, t2 are times in seconds (conversion factor from m/s to km/h is 3.6). To finalize IoT estimation precision we report Mean Absolute Percentage Error (MAPE):

$$mAP = \frac{1}{N} \sum_{i=1}^N AP_i \quad (7)$$

Where $i^{\wedge}\{ref\}$ viref is a calibrated reference (radar) measurement for event i.

The dataset for the ANPR evaluation was split into 276 training images, 79 validation images, and 38 test images. The performance of the model turned out to be strong, as shown in Table 6: YOLOv8 Detection and Recognition Metrics, where the detector attained a 96.2% mAP@0.5, supported by a precision of 94.8% and recall of 95.3%, amounting to an F1-score of 95.05%. The OCR component achieved a CER of 1.52% and WER of 5.53%, thus proving that this detection and recognition system is capable of reliably extracting Bangla licenseplate text even under tough field conditions. The

performance curves in Figure 6: YOLOv8 Training–Validation Performance Curve indicate that both training and validation metrics converged smoothly across epochs without any signs of overfitting; this confirms that augmentation during training has been very effective.

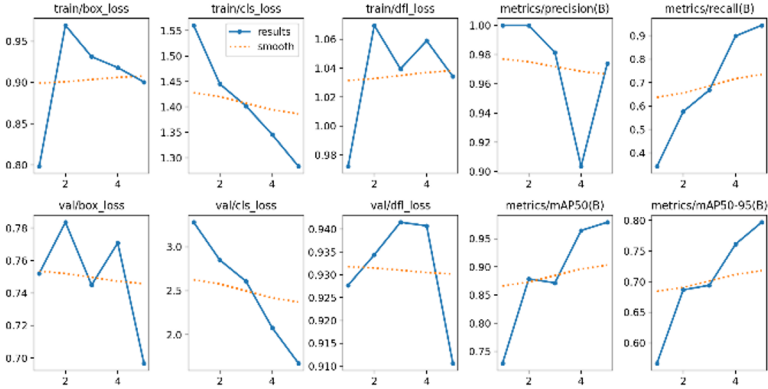


Fig. 6. Performance curve for the suggested detection model.

Table 7: Character Recognition Examples gives further insight into the OCR accuracy by listing representative predictions against their ground-truth labels. As indicated in the table, the system correctly recognizes most of the complex Bangla plate formats; however, some errors occurred, mainly caused by visually similar digits or characters under conditions of blur or partial occlusion. One such example shows a single-digit substitution, thereby incurring a CER of 6.25% in that sample. These observations highlight the fact that an OCR system, though robust, may degrade in performance due to partial obstruction or distortion of characters because of motion blur—a prevalent challenge in low-cost camera deployments. Figure 7: OCR Processing Pipeline illustrates how effective preprocessing methods, such as deskew correction and histogram equalization, enhance character boundaries and reduce noise before recognition, thereby improving the accuracy of OCR.

Table 6. Yolov8 (proposed pipeline) detection and recognition metrics.

Metric	Value
mAP @ 0.5	96.2%
Precision	94.8%
Recall	95.3%
F1-score	95.05%
Character Error Rate (CER)	1.52%
Word Error Rate (WER)	5.53%

Besides absolute performance, a comparative performance evaluation was performed with widely used ANPR pipelines reported in the literature. Table 8: Comparative Performance of ANPR Pipelines compares the proposed YOLOv8 + EasyOCR system with those of YOLOv5+Tesseract, YOLOv4+PaddleOCR, and Faster R-

CNN+Tesseract. From this comparison, it appears that the proposed approach outperforms earlier systems along all metrics, with 2–10 percentage points improvement in detection accuracy and 1.5–6 percentage points improvement in CER. This confirms that the newer YOLOv8 architecture is more befitting for the complex typography, ligatures, and multi-line layouts of Bangla license plates combined with a Bangla-sensitive OCR engine. Moreover, many comparative techniques suffered from nighttime images, tilted plates, and low-resolution inputs, whereas the proposed approach remained reliable even under rain, low light, and semi-occluded views due to the integrated pre-processing and augmentation pipeline.

Table 7. Character recognition from the detected number plat

Image ID	Ground Truth	OCR Prediction	Char Errors	CER
img_001	ঢাকা মেট্রো-গ ১৫-৫৮৭৭	ঢাকা মেট্রো-গ ১৫-৫৮৭৭	0	0.0%
img_002	ঢাকা মেট্রো-খ ২২-৩৪৫৬	ঢাকা মেট্রো-খ ২২-৩৪৫৬	0	0.0%
img_003	ঢাকা মেট্রো-চ ১১-৭৬৪২	ঢাকা মেট্রো-চ ১১-৭৬১২	1 (digit)	6.25%

The IoT subsystem was tested regarding its speed measurement capability using timestamp differences. The twin IR sensors accurately logged the entry and exit times over the 10-meter baseline, and the speed was computed using the usual distance-over-time formula and converted to km/h. In order to verify the correctness of the measurement, the sensor output was checked against a commercial radar gun. Over 4,000 events, the IoT system attained an MAPE of 2.8%, which turns out to be very consistent given its low-cost hardware origin. To give an idea of the robustness of the speed detection pipeline, Table 9: Sample IoT Speed Detection and Violation Records presents a few matched records showing the event timestamps, computed speeds, allowed speed limits, violation flags, and predicted license-plate text. These results confirm that the IoT subsystem operates reliably across school zones, university roadways, and arterial roads, and its measurements integrate seamlessly with the OCR predictions to generate comprehensive violation reports.

Table 8. Comparison of performance with other model.

Method (Detector + OCR)	mAP@0.5	Precision	Recall	CER
Proposed (YOLOv8 + EasyOCR)	96.2%	94.8%	95.3%	1.52%
YOLOv5 + Tesseract (literature)	90–94%	88–92%	86–91%	3–7%
YOLOv4 + PaddleOCR (literature)	86–92%	85–90%	82–89%	2–5%
Faster R-CNN + Tesseract (lit.)	80–88%	80–86%	78–85%	4–8%

The integrated performance of the vision and sensor subsystems provides a comprehensive view of the system's capabilities. The combined evaluation shows that the AIoT platform is able to detect speeding vehicles in real time, recognize Bangla license plates, and raise violation flags automatically with high accuracy. However, sources of error were present. The most common vision errors were related to night-time samples with glare from oncoming headlights, heavy rain where the characters were obscured by falling raindrops, and motion blur generated as a result of vehicles running at very

high speeds. Most OCR errors resulted from worn-out plates, plates partly covered by dirt or stickers, or odd fonts. For the IoT setup, the false triggers were seldom seen where pedestrians or animals cut across the IR beam; besides, deviations in timestamp accuracy took place due to environmental disturbances like fog or misalignment of sensors. The patterns of failure are pretty typical and correspond to well-documented limitations in the performance of low-cost ANPR and IR-based speed sensing systems. Despite these challenges, the system shows great potential as a low-cost and feasible solution for traffic surveillance in Bangladesh.

Table 9. Iot speed detection and combined violation sample records.

Event id	t1 (ISO)	t2 (ISO)	d (m)	Speed km	Speed limit	Violation flag	Plate text pred
ev_001	2025-07-11T10:02:14.010Z	2025-07-11T10:02:14.285Z	10.0	126.3	80	Yes	ঢাকা মেট্রো-গ ১৫-৫৮৭৭
ev_002	2025-07-11T10:05:44.122Z	2025-07-11T10:05:44.612Z	10.0	72.0	80	No	ঢাকা মেট্রো-খ ২২-৩৪৫৬
ev_003	2025-07-11T10:07:21.005Z	2025-07-11T10:07:21.385Z	10.0	93.5	60	Yes	ঢাকা মেট্রো-চ ১১-৭৬৪২

The integration of ANPR with speed monitoring under a single AIoT platform enables automated enforcement, central monitoring, and easy deployment scalability from school zones to highways. The comparative results, along with error analysis and evaluation against real-world environmental conditions, establish that the proposed system represents important performance gains from earlier ANPR systems for Bangla plates, while maintaining operational feasibility via low-cost sensor hardware and lightweight cloud connectivity.




Fig. 7. OCR working technique and pipeline.

5 Conclusion

This paper presented an integrated AIoT framework that brought together YOLOv8-based Bangla license plate detection, EasyOCR recognition, and a dual infrared sensor system for the estimation of vehicle speed to support evidence-based road safety enforcement in Bangladesh. Experimental results demonstrate very robust performance characterized by 96.2% mAP for plate detection, a low Character Error Rate of 1.52% in OCR, and a 2.8% Mean Absolute Percentage Error in IoT speed measurement against radar benchmarks. These results confirm that a low-cost combination of computer vision and sensor-based monitoring can identify overspeeding vehicles with high reliability while generating synchronized, tamper-resistant violation records suitable for smart-city traffic management.

Despite these encouraging results, there are still a number of limitations. For example, the vision model was trained on a relatively modest-sized dataset of 393 annotated images, which limits its generalization across extreme lighting conditions, heavy rain, nighttime glare, and heavy-traffic occlusions. The OCR module also sees occasional errors for tilted or partially occluded plates or plates that are worn out. The dual IR sensor rig is cost-effective; however, it is susceptible to fog, misalignment in sensors, pedestrian interference, and multi-lane scenarios. Additionally, national-scale deployment would necessitate more robust security on the backend, encrypted communication, and efficient storage policies for sensitive data related to vehicles.

However, the system forms a very robust foundation for an extensible national-level traffic enforcement solution. Future work will be targeted at increasing the dataset to tens of thousands of images under various weather and lighting conditions, thereby enhancing nocturnal accuracy with infrared-capable imaging. Transformer-based OCR for difficult Bangla scripts will also be integrated, and IoT hardware will be upgraded with radar or LiDAR for multi-lane support. Other plans include HTTPS/MQTT encryption, role-based access control, and scalable cloud infrastructure for secure large-scale deployment. With such advancements, the proposed AIoT platform can become a comprehensive, real-time, and low-cost traffic monitoring solution for Bangladesh's smart-city initiatives.

References

1. Al Noman, Abdullah, Nushrat Jahan Mila, Abrar Hameem Bornil, Ayon Mazumder, Nusrat Khan, and Md Hasan Imam Bijoy. "BanglaPlate: Automated Vehicle Number Plate Detection and Recognition using YOLOv8." In 2025 International Conference on Electrical, Computer and Communication Engineering (ECCE), pp. 1-6. IEEE, 2025.
2. Begum, Momotaz, Abm Kamrul Islam Riad, Abdullah Al Mamun, Thofazzol Hossen, Salah Uddin, Md Nurul Absur, and Hossain Shahriar. "Internet of Things (IoT)-Based Solutions for Uneven Roads and Balanced Vehicle Systems Using YOLOv8." *Future Internet* 17, no. 6 (2025): 254.
3. Rabby, Md Khurram Monir, Muhammad Mobaidul Islam, and Salman Monowar Imon. "A review of IoT application in a smart traffic management system." In 2019 5th International Conference on Advances in Electrical Engineering (ICAEE), pp. 280-285. IEEE, 2019.

4. Islam, Nuraiyad Nafiz, Sadia Afroz, Md Faizul Bari, and ABM Alim Al Islam. "Short Paper: Detecting and Decoding: A YOLO-Transformer Hybrid Model for Bangla License Plate Recognition." In Proceedings of the 11th International Conference on Networking, Systems, and Security, pp. 216-221. 2024.
5. Putra, Arman Syah, and Harco Leslie Hendric Spits Warnars. "Intelligent traffic monitoring system (ITMS) for smart city based on IoT monitoring." In 2018 Indonesian Association for Pattern Recognition International Conference (INAPR), pp. 161-165. IEEE, 2018.
6. Al-Hasan, Tamim Mahmud, Victor Bonnefille, and Faycal Bensaali. "Enhanced yolov8-based system for automatic number plate recognition." *Technologies* 12, no. 9 (2024): 164.
7. Zhang, Junkai, Jun Wang, Haoyu Zang, Ning Ma, Martin Skitmore, Ziyi Qu, Greg Skulmoski, and Jianli Chen. "The application of machine learning and deep learning in intelligent transportation: a scientometric analysis and qualitative review of research trends." *Sustainability* 16, no. 14 (2024): 5879.
8. Elattar, Tarteel Abdelfattah, Youssef Abdelrahman Ahmed, Mohammed Abdel-Megeed Salem, and Shereen Moataz Afifi. "Egyptian motorcycle and vehicle license plate recognition." *Neural Computing and Applications* (2025): 1-22.
9. Zhong, Guidong, Ke Xiong, Zhangdui Zhong, and Bo Ai. "Internet of things for high-speed railways." *Intelligent and Converged Networks* 2, no. 2 (2021): 115-132.
10. Chandra, B.M., Sonia, D., Roopa Devi, A., Yamini Saraswathi, C., Mighty Rathan, K. and Bharghavi, K., 2023. Recognition of vehicle number plate using Matlab. *J. Univ. Shanghai Sci. Technol*, 23(2), pp.363-370.
11. Gnanaprakash, V., Kanthimathi, N. and Saranya, N., 2024, March. Automatic number plate recognition using deep learning. In *IOP Conference series: materials science and engineering* (Vol. 1084, No. 1, p. 012027). IOP Publishing.
12. Kumar, J.R., Sujatha, B. and Leelavathi, N., 2024, February. Automatic vehicle number plate recognition system using machine learning. In *IOP Conference Series: Materials Science and Engineering* (Vol. 1074, No. 1, p. 012012). IOP Publishing.
13. Islam, D., Mahmud, T. and Chowdhury, T., 2023. An efficient automated vehicle license plate recognition system under image processing. *Indonesian Journal of Electrical Engineering and Computer Science*, 29(2), pp.1055-1062.
14. Mhatre, A. and Sharma, P., 2023. Deep Learning Approach for Vehicle Number Plate Recognition System with Image Enhancement Technique. *International Journal of Intelligent Systems and Applications in Engineering*, 11(1s), pp.251-262.
15. rema, M.D.R., Balaji, S.S., Srikanth, P. and Kumar, S.A., LICENSE PLATE DETECTION USING IMAGE PROCESSING.
16. Garg, P., Toshniwal, P. and Panjwani, R., REAL TIME MOVING VEHICLE REGISTRATION PLATE DETECTION.
17. Alsudani, M.Q., Waheed, S.R., Kadhim, K.A., Adnan, M.M. and Al- khaykan, A., 2023. Automatic Car Number Plate Detection using Morphological Image Processing. *Malaysian Journal of Fundamental and Applied Sciences*, 19(3), pp.379-388.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

