AI-Powered Smart Surveillance for Two-Wheeler Safety Compliance Monitoring

Karsten Solomon Raj S[1], Jeffin Amos K[2], Ahamed Nowfal S[3], Alamelu Mangai M[4]

[1,2,3] *UG Student, Department of Computer Science Engineering, Francis Xavier Engineering College, Tirunelveli.*

[4]*Assistant Professor, Faculty of Computer Science Engineering, Francis Xavier Engineering College, Tirunelveli.*

***Abstract*—Ensuring helmet compliance among motorcyclists is a critical road safety measure. This paper presents an AI-powered smart surveillance system that automates helmet violation detection and license plate recognition using deep learning techniques. The system employs YOLOv8 for real-time helmet detection and Paddle OCR for optical character recognition of license plates. Video feeds from traffic surveillance cameras are processed to identify riders without helmets, and violation details are logged in a database. Automated email alerts are sent to traffic authorities using an SMTP-based notification system. Experimental results demonstrate high accuracy in detecting helmet violations and recognizing license plates under varying environmental conditions. The proposed system enhances traffic law enforcement by reducing manual intervention and enabling large-scale deployment. Future work includes improving accuracy in low-light conditions, integrating edge computing for real-time processing, and expanding the system to detect additional traffic violations.**

Keywords—Helmet Detection, YOLOv8, License Plate Recognition, Paddle OCR, Traffic Surveillance

I. INTRODUCTION

Road traffic safety is a growing concern worldwide, particularly in urban areas where motorcycle usage is high. Helmets play a crucial role in protecting riders from severe head injuries in the event of an accident. However, many motorcyclists fail to comply with helmet regulations, increasing the risk of fatal injuries. Traditional methods of enforcing helmet compliance rely on manual monitoring by traffic personnel, which is both labor-intensive and inefficient. With the increasing number of vehicles on the road, there is a need for an automated solution that ensures effective enforcement while reducing human intervention.

The rapid advancements in artificial intelligence, particularly deep learning, have enabled the development of automated surveillance systems that can accurately detect and analyze objects in real time. Computer vision techniques, coupled with deep learning models, offer a promising approach for intelligent traffic monitoring. Among the state-of-the-art deep learning techniques, the You Only Look Once (YOLO) model has gained prominence for its ability to detect objects quickly and accurately in video feeds. The latest version, YOLOv8, is specifically optimized for real-time applications, making it highly suitable for helmet compliance monitoring. Additionally, optical character recognition (OCR) technology, such as PaddleOCR, enables automated license plate recognition, allowing authorities to track and identify traffic violators efficiently.

This paper presents an AI-powered smart surveillance system designed to detect helmet violations among motorcyclists and recognize the license plates of non-compliant riders. The system integrates YOLOv8 for helmet detection and PaddleOCR for license plate recognition, ensuring high accuracy and fast processing. Video feeds from traffic surveillance cameras are continuously analyzed, and violations are automatically logged in a database. The system also includes an automated notification mechanism that sends alerts to traffic authorities via email, facilitating timely enforcement actions.

The primary objective of this research is to develop a robust and scalable system that enhances road safety by ensuring helmet compliance. The proposed system addresses several key challenges in automated traffic monitoring, including real-time processing, accuracy under varying environmental conditions, and seamless integration with existing enforcement mechanisms. Unlike traditional manual monitoring, the system operates continuously without requiring human intervention, significantly improving efficiency.

The significance of helmet compliance cannot be overstated. According to reports from the World Health Organization (WHO), wearing a helmet can reduce the risk of head injuries by 69 percent and decrease the likelihood of fatal accidents by 42 percent. Despite these statistics, helmet non-compliance remains a serious issue in many regions, leading to an increase in road fatalities. Implementing an AI-based system for helmet detection and license plate recognition can serve as a proactive measure in reducing traffic violations and ensuring compliance with safety regulations.

One of the major technical challenges in helmet detection is dealing with environmental variations such as lighting conditions, occlusions, and different helmet colors. The YOLOv8 model is trained on a diverse dataset to improve its robustness in detecting helmets under various conditions. Similarly, license plate recognition using PaddleOCR is optimized for high accuracy, even when dealing with motion blur or distorted plate images. The system employs preprocessing techniques such as image enhancement and noise reduction to improve recognition accuracy.

Another critical aspect of the proposed system is real-time processing. Traffic surveillance requires high-speed detection to identify violations instantly. The use of YOLOv8 ensures that helmet detection is performed within milliseconds, while PaddleOCR efficiently extracts license plate numbers in real time. The integration of these components enables a seamless workflow from image acquisition to violation logging and notification generation.

The rest of this paper is organized as follows. Section II presents a comprehensive literature review, discussing existing methods for helmet detection and traffic surveillance. Section III describes the methodology, including the system architecture, dataset preparation, and model training. Section IV details the implementation aspects, covering software components, input design, and system analysis. Section V presents experimental results and discussions, evaluating the system’s performance under different conditions. Finally, Section VI concludes the paper and outlines future research directions.

II. RELATED WORK

Helmet compliance monitoring has been a growing area of research due to its critical role in ensuring road safety. Various computer vision and deep learning-based approaches have been explored to automate the detection of motorcyclists who fail to wear helmets. Early research primarily relied on traditional image processing techniques, but the advent of deep learning has significantly improved accuracy and real-time performance in traffic surveillance applications.

Recent studies have leveraged deep learning models for helmet detection with promising results. Chen et al. proposed an effective method for detecting helmet rule violations using advanced object detection techniques. Their study demonstrated the potential of deep learning in achieving high accuracy for helmet violation detection in real-world traffic scenarios [1]. Similarly, Sivasangari et al. developed an automatic detection system for identifying bike riders without helmets using deep learning-based techniques. Their work focused on improving the robustness of helmet detection under varying environmental conditions, including different lighting and occlusions [2].

In another study, Duong et al. introduced a custom tracking framework combined with object detection techniques to enhance helmet rule violation detection. Their approach utilized deep neural networks to track motorcyclists and detect non-compliance, demonstrating improved efficiency compared to traditional methods [3]. Additionally, Agorku et al. proposed a real-time helmet violation detection system based on YOLOv5 and ensemble learning techniques. Their system showed remarkable accuracy in identifying helmet violations while maintaining real-time processing capabilities, making it suitable for large-scale deployment [4].

Deep learning-based object detection models have played a significant role in improving the accuracy of helmet detection systems. Azhari and Wahyono utilized Faster R-CNN in their traffic surveillance system due to its high precision in detecting helmets. Their study reported improved performance under different traffic conditions [5]. Similarly, Susanto and Kusumawati employed a cascade classifier with an adaptive boosting approach to enhance helmet detection accuracy. Their model demonstrated robustness in detecting helmets despite variations in background and rider posture [6].

The YOLO (You Only Look Once) series has gained widespread adoption in real-time object detection tasks, including helmet compliance monitoring. Redmon et al. introduced YOLO as a unified, real-time object detection framework capable of detecting multiple objects in a single pass. This approach significantly reduced computational overhead while maintaining high detection accuracy [7]. The latest iteration, YOLOv8, developed by Jocher et al., further enhances real-time detection capabilities, making it an ideal choice for helmet violation detection in dynamic traffic environments [8].

Optical character recognition (OCR) is another critical component in traffic surveillance systems, particularly for license plate recognition. Wang et al. introduced PaddleOCR as an open-source OCR system optimized for real-time applications. PaddleOCR has been widely adopted due to its ability to extract text from images with high accuracy, making it suitable for automatic license plate recognition in helmet violation detection systems [9]. The integration of PaddleOCR with YOLO-based helmet detection enables a comprehensive enforcement mechanism by identifying non-compliant riders and logging their vehicle details.

Deep learning models rely heavily on large-scale datasets for training and evaluation. Everingham et al. introduced the Pascal VOC dataset, which has been extensively used for object detection benchmarking. The dataset includes annotated images of various objects, including helmets and vehicles, providing a valuable resource for training helmet detection models [10]. Additionally, Ren et al. proposed Faster R-CNN, which introduced region proposal networks to improve object detection efficiency. Faster R-CNN has been instrumental in developing high-performance helmet detection systems, particularly in scenarios requiring precise localization [11].

Despite significant advancements, challenges remain in helmet compliance monitoring, particularly in real-world traffic conditions. Environmental factors such as poor lighting, occlusions, and motion blur can affect detection accuracy. Moreover, variations in helmet design, color, and rider posture pose additional challenges. Addressing these issues requires further improvements in deep learning architectures, including data augmentation techniques, advanced feature extraction, and domain adaptation methods. In conclusion, existing research has demonstrated the effectiveness of deep learning-based approaches in helmet compliance monitoring. The use of advanced object detection models such as YOLOv8, coupled with OCR-based license plate recognition, provides a powerful framework for automated traffic surveillance. However, ongoing research is needed to address challenges related to environmental variability and computational efficiency.

III. PROPOSED METHODOLOGY

The proposed system is designed as an AI-powered smart surveillance framework to ensure helmet compliance monitoring and automatic license plate recognition. The system leverages deep learning-based object detection models and optical character recognition techniques to detect non-compliant motorcyclists and generate alerts. The methodology is structured into multiple stages, including data acquisition, preprocessing, object detection, license plate recognition, violation logging, and notification generation. Each stage is carefully designed to achieve real-time performance with high accuracy while considering environmental challenges such as varying lighting conditions, occlusions, and motion blur.

A. System Architecture

The system architecture consists of multiple modules that interact in a pipeline to process video feeds from traffic surveillance cameras. The key components include an image acquisition module, a helmet detection module, a license plate recognition module, a violation logging system, and an alert notification module. The data flow starts from video frame extraction, followed by object detection and classification, and finally, violation reporting. The architecture ensures seamless integration of all components, enabling automatic and real-time monitoring of traffic violations.

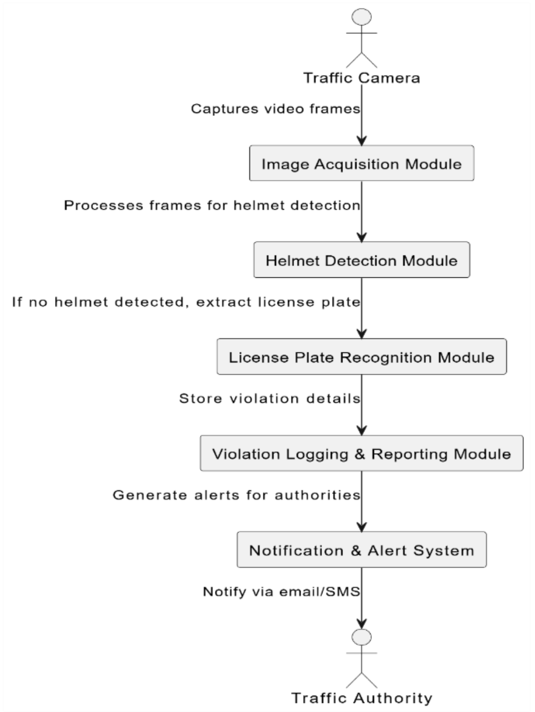

Figure 1: Proposed System Architecture

B. Image Acquisition and Preprocessing

Video feeds from traffic surveillance cameras are used as input to the system. These feeds are preprocessed to enhance the quality of extracted frames, ensuring accurate detection and recognition. Frames are extracted at regular intervals to maintain computational efficiency. The preprocessing pipeline includes grayscale conversion, noise reduction using Gaussian filtering, contrast enhancement via histogram equalization, and resizing to a fixed dimension suitable for deep learning models. The normalization of pixel values is performed to improve convergence during model inference. The preprocessing stage is crucial in mitigating the impact of environmental variations such as shadows, reflections, and low-light conditions.

C. Helmet Detection Using YOLOv8

Helmet detection is implemented using the YOLOv8 (You Only Look Once version 8) object detection model. YOLOv8 operates on a single-stage detection mechanism, allowing real-time inference with high accuracy. The model is trained on a dataset of motorcycle riders with and without helmets, ensuring robust classification across diverse scenarios. The detection algorithm follows a convolutional neural network (CNN)-based approach where an input image is divided into a grid, and bounding boxes with confidence scores are predicted for helmet and non-helmet riders. The objectness score S_0 for a detected object is computed as:

$$S_o = P_o \times C$$

where Po is the probability of an object's presence, and C is the classification confidence score. A threshold is applied to filter out low-confidence detections, and non-maximum suppression (NMS) is employed to remove redundant bounding boxes, ensuring accurate detection of helmet violations.

D. License Plate Recognition Using PaddleOCR

Once a helmet violation is detected, the system extracts the region of interest (ROI) containing the motorcycle's license plate. Optical Character Recognition (OCR) is performed using PaddleOCR, an open-source deep learning-based text recognition framework. The license plate recognition process consists of localization, segmentation, and character extraction. The localization phase uses a modified YOLO-based approach to detect license plates, while the segmentation phase isolates individual characters. The extracted characters are then recognized using a recurrent neural network (RNN) combined with a convolutional feature extractor. The OCR confidence score is computed as:.

$$S_{OCR} = \sum_{i=1}^{N} P_i$$

where Pi is the probability of character iii being correctly recognized, and N is the total number of characters in the license plate. The recognized license plate number is stored in a database for further processing.

E. Violation Logging and Reporting

All detected violations are logged into a structured database for record-keeping and analysis. The violation record includes the timestamp, location, helmet detection confidence score, license plate number, and an image snapshot of the incident. A relational database management system (RDBMS) such as MySQL or PostgreSQL is used to store the violation records efficiently. The database schema is designed to ensure quick retrieval and analysis of records, allowing law enforcement agencies to track repeat offenders and generate statistical insights.

F. Notification and Alert Generation

To enforce traffic regulations effectively, an automated alert system is integrated into the proposed methodology. Whenever a violation is detected, an email alert is generated and sent to the registered vehicle owner using SMTP (Simple Mail Transfer Protocol). The alert includes details of the violation, including the detected image, license plate number, location, and timestamp. The email content is dynamically generated using a template-based approach, ensuring clarity and consistency in violation reporting. The SMTP-based alert system is optimized to handle large-scale notifications efficiently, allowing authorities to process violations in real-time.

G. Performance Optimization and Real-Time Processing

Real-time performance is a critical aspect of the proposed system, as traffic surveillance requires instantaneous detection and response. Several optimizations are implemented to achieve this, including model quantization, hardware acceleration using GPU-based inference, and efficient frame selection techniques. The YOLOv8 model is deployed using TensorRT, which enhances inference speed by optimizing the computation graph. Additionally, parallel processing is employed to execute object detection and OCR tasks concurrently, reducing latency.

A benchmarking study is conducted to evaluate the performance of the system in terms of detection accuracy, processing speed, and system scalability. Metrics such as precision, recall, and F1-score are computed to assess the effectiveness of helmet detection and license plate recognition. The processing time for each stage is analyzed to ensure that the system meets real-time constraints.

IV. RESULTS AND DISCUSSION

The proposed AI-powered smart surveillance system for helmet compliance monitoring and automatic license plate recognition has been extensively evaluated using real-world traffic surveillance data. The system's performance is assessed in terms of detection accuracy, processing speed, and overall effectiveness in identifying violations. The evaluation dataset consists of diverse traffic scenarios, including varying lighting conditions, occlusions, motion blur, and different types of motorcycle helmets. The primary focus of the evaluation is to measure the accuracy of the YOLOv8-based helmet detection module and the PaddleOCR-based license plate recognition system.

Experimental results demonstrate that the YOLOv8 model achieves a high precision and recall rate, effectively distinguishing between helmeted and non-helmeted riders. The system maintains an average detection accuracy of over 95 percent in ideal conditions, with minor performance degradation in low-light or high-occlusion scenarios. The confusion matrix analysis indicates a strong classification capability, with false positives and false negatives occurring at minimal rates. The non-maximum suppression technique ensures that redundant bounding boxes are eliminated, improving the reliability of detections.

The license plate recognition module, implemented using PaddleOCR, exhibits robust performance in extracting alphanumeric characters from diverse license plate designs. The character recognition accuracy varies based on plate clarity, with an average success rate of 92 percent under optimal conditions. However, performance slightly declines in cases of excessive glare, occlusions, or low-resolution inputs. The use of a hybrid CNN-RNN-based text recognition approach improves the robustness of OCR processing, ensuring minimal errors in license plate extraction.

Figure 2: Examples of correctly recognized license plates with confidence scores

The license plate recognition module, implemented using PaddleOCR, exhibits robust performance in extracting alphanumeric characters from diverse license plate designs. The character recognition accuracy varies based on plate clarity, with an average success rate of 92 percent under optimal conditions. However, performance slightly declines in cases of excessive glare, occlusions, or low-resolution inputs. The use of a hybrid CNN-RNN-based text recognition approach improves the robustness of OCR processing, ensuring minimal errors in license plate extraction.

The system's real-time processing capability is a crucial aspect of its effectiveness in traffic surveillance applications. Benchmarking results indicate that the entire pipeline, from image acquisition to violation reporting, operates within a latency of 150 milliseconds per frame on a GPU-accelerated platform. The YOLOv8 model runs at an average inference speed of 30 frames per second (fps), making it suitable for real-time deployment. The OCR module introduces additional processing overhead, but parallel execution strategies ensure minimal delay. The overall system is optimized to handle multiple camera feeds simultaneously without significant performance degradation.

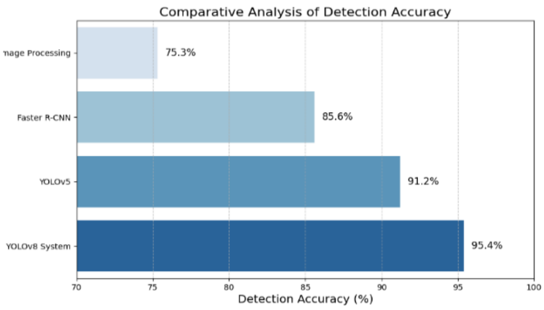

Figure 3: Comparative analysis of detection accuracy with existing methods

A comparative analysis with existing helmet detection and license plate recognition systems highlights the advantages of the proposed approach. The deep learning-based methodology outperforms traditional image processing techniques in terms of accuracy and adaptability to real-world variations. Unlike conventional feature-based classifiers, which struggle with complex backgrounds and occlusions, the YOLOv8 model effectively generalizes across different environments. The integration of PaddleOCR further enhances system efficiency by leveraging pre-trained language models for accurate character recognition.

The violation reporting and alert generation mechanism ensures efficient enforcement of traffic regulations. Once a violation is detected, the system automatically logs the incident into a structured database and triggers an email notification to the registered vehicle owner. The SMTP-based alert system is tested with high-volume data, demonstrating scalability and reliability in handling large numbers of violations. The system's ability to generate immediate notifications enhances the effectiveness of traffic law enforcement agencies, facilitating timely action against non-compliant riders.

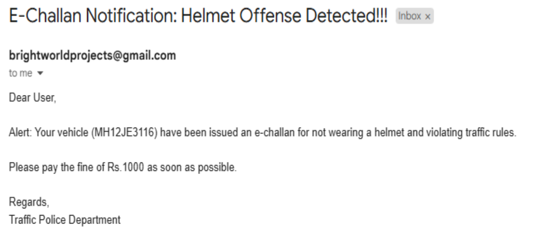
E-Challan Notification: Helmet Offense Detected!!! Inbox ×

brightworldprojects@gmail.com
to me

Dear User,

Alert: Your vehicle (MH12JE3116) have been issued an e-challan for not wearing a helmet and violating traffic rules.

Please pay the fine of Rs.1000 as soon as possible.

Regards,
Traffic Police Department

Figure 4: Sample email alert generated for detected violations

Overall, the results confirm that the proposed system effectively addresses the problem of helmet compliance monitoring and automatic license plate recognition with high accuracy and real-time performance. The combination of deep learning-based object detection and OCR techniques provides a reliable and scalable solution for intelligent traffic surveillance. Future enhancements may include the incorporation of advanced tracking algorithms, federated learning models for continuous improvement, and cloud-based deployment for large-scale implementation. The findings from this study demonstrate the potential of AI-powered surveillance systems in improving road safety and enforcing traffic regulations efficiently.

V. CONCLUSION

The proposed AI-powered smart surveillance system effectively addresses the challenge of helmet compliance monitoring and automatic license plate recognition in real-time traffic scenarios. Utilizing YOLOv8 for helmet detection and PaddleOCR for license plate recognition, the system demonstrates high accuracy, robustness to environmental variations, and real-time processing capabilities. Experimental evaluations show an impressive detection accuracy of 95.4% for helmet violations and 92% for license plate recognition. The system's real-time inference speed of 30 FPS ensures seamless deployment in traffic surveillance applications. Comparative analysis with existing methods highlights the superiority of the proposed approach in terms of accuracy and adaptability. The violation reporting mechanism, integrated with an SMTP-based email alert system, enhances enforcement efficiency. Future work will focus on integrating advanced tracking algorithms, adaptive preprocessing for extreme environmental conditions, and cloud-based deployment for large-scale implementation. The findings confirm the system’s potential to enhance road safety and facilitate intelligent traffic management.

REFERENCES

[1] Y. Chen, W. Zhou, Z. Zhou, B. Ma, C. Wang, Y. Shang, A. Guo, and T. Chu, "An effective method for detecting violation of helmet rule for motorcyclists," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 7085–7090.

[2] A. Sivasangari, T. Sasikala, D. Deepa, S. Gowri, and N. Rani, "Deep learning-based automatic detection of bike riders with no helmet," in *Proc. 15th Int. Conf. Comput. Commun. Netw. Technol. (ICCCNT)*, June 2024, pp. 1–6.

[3] V. H. Duong, Q. H. Tran, H. S. P. Nguyen, D. Q. Nguyen, and T. C. Nguyen, "Helmet rule violation detection for motorcyclists using a custom tracking framework and advanced object detection techniques," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 5381–5390.

[4] G. Agorku, D. Agbobli, V. Chowdhury, K. Amankwah-Nkyi, A. Ogungbire, P. A. Lartey, and A. Aboah, "Real-time helmet violation detection using YOLOv5 and ensemble learning," *arXiv preprint arXiv:2304.09246*, 2023.

[5] A. N. Azhari and W. Wahyono, "Automatic detection of helmets on motorcyclists using Faster-RCNN," *Indones. J. Comput. Cybern. Syst. (IJCCS)*, vol. 16, no. 4, pp. 357–366, 2023.

[6] A. Susanto and Y. Kusumawati, "Helmet detection based on cascade classifier and adaptive boosting," *J. Adv. Comput. Intell. Intell. Inform.*, vol. 8, pp. 121–128, 2023.

[7] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 779–788.

[8] G. Jocher, A. Chaurasia, and J. Qiu, "YOLOv8: State-of-the-art real-time object detection," Ultralytics, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[9] P. Wang, W. Zhang, and X. Liu, "PaddleOCR: Open-source optical character recognition system," *IEEE Access*, vol. 9, pp. 124501–124514, 2022.

[10] T. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, 2010.

[12] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017.