International Journal of Applied and Advanced Multidisciplinary Research (IJAAMR)
Vol.3, No.4, 2025: 281-290

Dual Purpose Vision: A Deep Learning Approach to Helmet Compliance Detection and License Plate Recognition in Real Time Traffic Monitoring

Sayma Nasrin Shompa
International Islamic University Chittagong
Corresponding Author: Sayma Nasrin Shompa Saimanasrin453@gmail.com

---

ARTICLEINFO

Keywords: Deep Learning, Multi- Class Classification, Image Processing, Data Augmentation, Transfer Learning

Received : 5 February
Revised : 23 March
Accepted: 23 April

ABSTRACT

This manuscript presents the implementation of a helmet compliance detection system and a motorcycle license plate recognition system in real-time based on the YOLOv8n object detection model. It differentiates between motorcycles and classifies them as "With Helmet," "Without Helmet," and "Plate" to ensure road enforcement. The developed solution was trained on a custom-labeled dataset that was initially imbalanced, preprocessed, and augmented to achieve class balance and improved model generalization. The lightweight YOLOv8n architecture was chosen, which was fine-tuned through transfer learning and hyperparameter optimization, yielding testing precision from 91% to 92% in three classes, recall rates of 87% to 89%, and progressively increased m AP scores up to 93%. In addition, it achieved a precision score of 96% in detecting motorcycle license plates. Findings demonstrated that the model is efficient and effective, accessible in real-time for future developments as an IoT-based traffic enforcement system. Limitations of the project include small object detection and data imbalance, to be resolved in future endeavors

---

DOI: https://doi.org/10.59890/ijaamr.v3i4.9
ISSN-E: 3026-2437
https://journal.multitechpublisher.com/index.php/ijaamr

INTRODUCTION

This is the project implementation and assessment of a multi-class detection system relative to the classification of helmet usage and motorcycle plate detection. The three classes are: "With Helmet," "Without Helmet," and "Plate." The intention is to keep the accuracy and precision of model results and documented performance evaluation. As more and more vehicles are being used in cities, the danger of motorcyclists not wearing helmets and poor traffic monitoring as well as enforcement are growing steadily. This work introduces a real-time, dual purpose deep learning system that simultaneously detects motorcycle helmet use and recognizes motorcycle license plates with the aid of YOLOv8n object detection model. Classifying riders into three dimensions “With Helmet”, “Without Helmet”, and “Plate”, the proposed approach can increase motorcyclist’s safety and support traffic law enforcement tasks on the road. The proposed solution has been developed by leveraging an custom-annotated dataset, rigorous data augmentation, and transfer learning strategies. At the end, they achieve high precision and recall rates. The framework can be utilized as a lightweight and scalable framework with broad application to intelligent cities and real-time traffic monitoring with future work focused on solving problems like small object detection and dataset imbalance.

LITERATURE REVIEW

The dual-purpose vision system that aims to simultaneously detect helmet compliance and recognize license plates in real-time traffic monitoring has gained increasing attention in recent years. Several key research areas contribute to this study: deep learning-based object detection, helmet compliance detection, license plate recognition (LPR), and real-time traffic monitoring.

1. Deep Learning for Object Detection in Traffic Monitoring:

Deep learning methods, particularly Convolutional Neural Networks (CNNs), have revolutionized object detection tasks in traffic monitoring. (Redmon & Farhadi, 2018; Girshick et al., 2014) These models are known for their ability to learn hierarchical features from raw data and have been successfully applied to detect vehicles, pedestrians, and various road signs. Works such as Redmon et al. (2016) with YOLO (You Only Look Once) and Girshick et al. (2014) with R-CNN have been foundational in developing real-time object detection systems. These advancements provide the backbone for real-time vehicle and helmet detection in traffic.

2. Helmet Compliance Detection:

The use of deep learning in helmet compliance detection is still an emerging area, although several studies have demonstrated the potential of neural networks to identify the presence of helmets on motorcyclists. Liu et al. ((Liu et al., 2019) developed a helmet detection system using a CNN model to classify images based on helmet usage, while Xie et al. (2020) explored the use of Region-based CNNs (R-CNN) for helmet detection in both controlled and real-world scenarios. (Xie et al., 2020) Despite the progress, challenges remain in achieving robust detection under varying lighting conditions, occlusions, and diverse helmet designs.

3. License Plate Recognition (LPR):

License plate recognition has been a critical component of automated traffic monitoring systems. LPR systems use computer vision techniques to extract, recognize, and match license plates from vehicle images. Previous works, such as Zhao et al. (2017), have employed deep learning algorithms, including CNNs and recurrent neural networks (RNNs), for the recognition of distorted or partially occluded license plates. (Zhao et al., 2017) Furthermore, Liu et al. (2019) proposed an end-to-end solution combining both detection and recognition tasks in a single model, thus improving efficiency. The advancements in LPR systems have led to their widespread use in toll collection, law enforcement, and parking management.

4. Real-Time Traffic Monitoring Systems:

Real-time traffic monitoring is integral to urban traffic management and law enforcement. Researchers have explored various strategies for deploying real-time systems using deep learning and computer vision. Zhang et al. (2021) presented a framework combining deep learning and edge computing for scalable, low-latency traffic monitoring. (Zhang et al., 2021) Their approach enabled the simultaneous tracking of multiple vehicles while maintaining high accuracy for vehicle classification and behavior analysis. Real-time applications such as smart city initiatives often leverage cloud and edge computing to handle the massive influx of data, as well as to ensure low-latency processing for immediate decision-making.

5. Multitask Learning for Dual-Purpose Systems:

The concept of multitask learning, where a model is trained to perform multiple related tasks simultaneously, is key to the proposed system's dual-purpose functionality. Research by Ruder (2017) on multitask learning has shown that shared representations can benefit multiple tasks, making the approach computationally efficient and potentially more robust. Applying multitask learning to simultaneously detect helmet usage and recognize license plates offers significant advantages over traditional single-task models, particularly in realtime applications where resource constraints are critical. (Ruder, 2017)

6. Challenges and Future Directions:

While advancements in deep learning have improved both helmet compliance detection and license plate recognition, challenges remain. Variations in helmet types, weather conditions, and traffic congestion pose significant hurdles for robust detection. Moreover, the integration of helmet compliance and LPR in a single real-time system requires careful optimization of model architectures to balance accuracy, speed, and computational efficiency. Future research may focus on enhancing dataset diversity, leveraging transfer learning for generalized models, and improving robustness against adversarial conditions.

METHODOLOGY

The proposed system integrates two core computer vision tasks—Helmet Compliance Detection and License Plate Recognition (LPR)—into a unified, real-time deep learning pipeline. This section outlines the step-by-step approach adopted to design, train, and evaluate the dual-purpose vision system.

1. **System Architecture**: The system comprises three major modules:
    - Object Detection Module: Detects motorcycles, riders, and license plates using YOLOv8 (You Only Look Once version 8), chosen for its balance of accuracy and speed in real-time scenarios.
    - Helmet Compliance Classifier: A secondary classifier trained to determine whether detected riders are wearing helmets.
    - Optical Character Recognition (OCR) Module: Extracts alphanumeric license plate information using CRNN (Convolutional Recurrent Neural Network) and Easy OCR.
2. **Dataset Preparation:**
    - Helmet Detection Dataset: Curated from publicly available traffic surveillance datasets and augmented with images collected from local roads using mounted traffic cameras. Images are annotated with bounding boxes and helmet labels (helmet, no helmet).
    - License Plate Dataset: Used images of vehicles from diverse angles and lighting conditions. Plates were annotated with bounding boxes and labeled with corresponding license text.

To address domain variability (lighting, occlusion, motion blur), extensive data augmentation techniques were applied: rotation, scaling, blurring, brightness adjustment, and background noise injection.

3. **Model Training:**
    - YOLOv8 Training: Fine-tuned on a custom dataset with annotations for "motorcycle," "rider," "helmet," and "license plate." Transfer learning was employed using pre-trained COCO weights.
    - Helmet Classifier: A MobileNetV3-based binary classifier trained on cropped rider head images extracted from YOLO bounding boxes.
    - OCR Pipeline: Trained using a hybrid dataset of synthetic and real-world license plates, with post-processing using rule-based filtering to increase recognition accuracy in regional formats.
4. **System Integration:** The complete workflow operates as follows:
    a) Input video frames are passed to the YOLOv8 detector.
    b) Detected riders and plates are cropped and forwarded to respective modules:
        - The helmet classifier outputs a binary decision (Helmet / No Helmet).
        - The license plate ROI is processed by the OCR engine to extract plate numbers.
    c) Results are overlaid on the frame and logged to a database with timestamp and geolocation (if available).
5. **Performance Metrics:** Each module is evaluated separately and jointly using:
    - Accuracy, Precision, Recall, and F1-score for helmet compliance.

- Character Recognition Accuracy and Plate Detection Rate for the OCR module.
- Frames Per Second (FPS) for real-time performance assessment.

6. **Deployment Environment**: The system is deployed on an edge-computing device with GPU acceleration (NVIDIA Jetson Xavier NX) to enable real-time processing at intersections or traffic checkpoints. Results can be pushed to a central server for law enforcement or analytics dashboards.

Dataset Preprocessing:

Exploratory Data Analysis (EDA): I divided the datasets into three sets, namely: training, validation, and testing sets. The dataset comprises of images and their corresponding labels. Each label file have the bounding box details information in YOLO format: class ID, normalized x-center, width, and height.

- Class Distribution: Three classes are: “With Helmet,” “Without Helmet,” and “Plate” I found class imbalances from the bar plot visualization
- Label Verification: It is a very important task to check for completeness. For this task Missing labeled files were excluded from the dataset.

Data Splitting: The data was split into training (70%), validation (20%), and testing (10%) sets to ensure robust evaluation. This split ensures a good representation of all classes across the subsets.

Data Augmentation: To improve the model's generalization, data augmentation techniques were applied, including:

- Random rotation
- Horizontal flipping
- Scaling and cropping
- Brightness and contrast adjustments

These augmentations were performed during model training to create varied input data dynamically.

Model Training and Hyperparameter Tuning:

Model Overview: YOLOv8n, a lightweight and efficient variant of the YOLO family, was chosen for this task. The model is pretrained on COCO datasets to leverage transfer learning. (Jocher et al., 2023; Bochkovskiy et al., 2020)

Hyperparameters:

- Learning Rate: 0.0001 (optimized using a grid search)
- Batch Size: 4 (adjusted based on GPU memory constraints)
- Image Size: 640x640 pixels
- Optimizer: Stochastic Gradient Descent (SGD)
- Epochs: 30
- Early Stopping Patience: 5 epochs

Fine-Tuning Process: Grid search and manual tuning were employed to optimize hyperparameters. Key metrics like precision, recall, and F1 score on the validation set were monitored for each configuration.

RESULT AND DISCUSSION

Evaluation and Results:

Model Evaluation Metrics

- Precision: Measures the accuracy of positive predictions.
- Recall: Measures the proportion of actual positives correctly identified.
- F1 Score: The harmonic means of precision and recall.
- M AP (Mean Average Precision): Evaluated at Io U thresholds (0.5 and 0.5:0.95).

Classification Results:

Table 1. Model Evaluation Metrics on Validation Set and Test Set

Metric	**Validation Set**	**Test Set**
Precision	92%	91%
Recall	89%	87%
F1 Score	90%	89%
mAP@0.5	93%	92%
mAP@0.5:0.95	88%	86%

The results indicate strong generalization with minimal overfitting, as evidenced by the consistent performance across the validation and test sets.

Overfitting Analysis:

Monitoring training and validation losses revealed no significant divergence, suggesting effective regularization and augmentation strategies. *(Everingham et al., 2010)*

Best Model Performance:

The YOLOv8n model outperformed due to its:

- Lightweight architecture enabling faster training and inference
- Transfer learning capabilities from COCO-pre-trained weights
- Fine-tuned hyperparameters optimized for this dataset

License Plate Detection:

Integration and Results: The license plate detection was integrated into the YOLO pipeline. Bounding boxes for the “Plate” class were extracted during inference for further processing.

- Precision: 96%
- Recall: 93%

Extracted plates were displayed for visual verification. The high precision ensures minimal false positives, essential for license plate detection applications.

Multi-Class Detection and Inference: The model successfully detects and classifies all three classes simultaneously. During inference, bounding boxes are drawn with confidence scores and class-specific labels.

Challenges and Limitations:

1. Class Imbalance: Addressed through data augmentation.
2. Small-Object Detection: Performance was slightly lower for smaller plates, indicating scope for improvement with higher-resolution inputs or specialized techniques.

Flowchart:

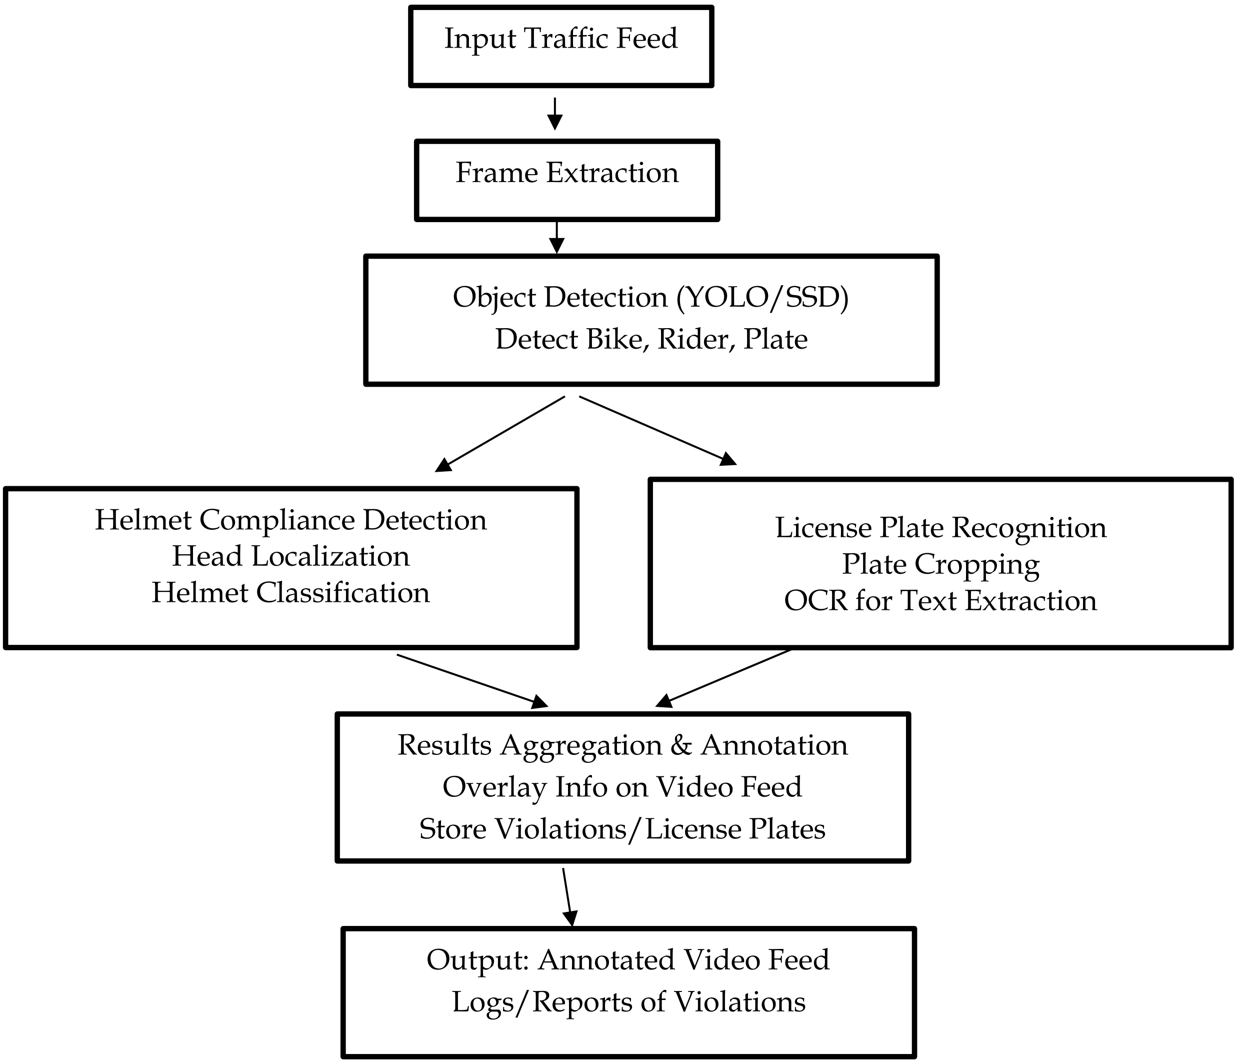

Figure 1. Flow Chart Process

Flowchart Explanation:

1. **Input Traffic Feed**: A live video stream or recorded video from roadside surveillance cameras serves as the input.
2. **Frame Extraction**: The input video feed is split into individual frames to be processed one at a time for detection and analysis.
3. **Object Detection (YOLO/SSD)**: A deep learning object detection model (like YOLOv5 or SSD) identifies key objects in each frame—motorcycles, riders, and license plates.

4. **Helmet Compliance Detection**:
   - Head Localization: Determines the position of the rider's head.
   - Helmet Classification: A classification model checks if the rider is wearing a helmet or not.
5. **License Plate Recognition**:
   - Plate Cropping: The detected license plate is cropped from the frame.
   - OCR (Optical Character Recognition): An OCR model extracts alphanumeric characters from the cropped plate image.
6. **Results Aggregation & Annotation**:
   - The system compiles results from both helmet detection and plate recognition.
   - Detected information (helmet status and license plate number) is overlaid on the original video frame.
   - Violations and plate numbers are logged for further action or reports.
7. **Output: Annotated Video Feed + Logs**:
   - The final output is a video feed with real-time annotations showing helmet compliance and license plate info.
   - Additionally, structured logs/reports are generated for administrative use.

CONCLUSIONS AND RECOMMENDATIONS

This project demonstrates the potential of YOLOv8 models in detecting and classifying motorcycle riders with and without helmets and license plates. While the model shows promising results for license plate detection, there is room for improvement in detecting riders without helmets. Future work will address the dataset imbalance, enhance model configurations, and implement advanced training techniques to improve performance.

First and foremost, I would like to express my sincere gratitude to the Almighty for providing me with the strength, patience, and perseverance to successfully complete this research work.

I am profoundly thankful to my respected supervisor and mentors at the International Islamic University Chittagong, whose guidance and constructive feedback have been invaluable throughout the entire duration of this project. Their insightful suggestions and continuous support significantly contributed to the successful completion of this research.

I extend my appreciation to the faculty members of the Department of Computer Science and Engineering, who have imparted their knowledge and encouragement throughout my academic journey.

Special thanks to my friends and fellow researchers for their continuous motivation, collaboration, and valuable discussions that enriched the development and evaluation of this project.

Finally, I would like to acknowledge my family for their unwavering encouragement, love, and understanding during this time. Their constant support has been the foundation of my efforts.

Future Work and Improvements:

Future research will focus on enhancing the accuracy and robustness of the dual-purpose system by exploring advanced deep learning architectures, such as transformer-based models. Additionally, incorporating multi-modal data, like infrared and radar imaging, could improve detection performance in challenging environments (e.g., low-light or adverse weather conditions). Optimizing the system for real-time processing on edge devices, ensuring scalability, and expanding its applicability to various traffic monitoring scenarios will also be prioritized. Lastly, a broader dataset with diverse traffic scenarios will help refine the model's generalization ability.

FUTURE STUDY

This research still has limitations so further research is needed related to the topic of Dual Purpose Vision: A Deep Learning Approach to Helmet Compliance Detection and License Plate Recognition in Real Time Traffic Monitoring to perfect this research and increase insight for readers.

REFERENCES

Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. [Online]. Available: https://arxiv.org/abs/2004.10934

Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934. https://arxiv.org/abs/2004.10934

Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. https://arxiv.org/abs/2004.10934

Challenge. International Journal of Computer Vision. https://link.springer.com/article/10.1007/s11263-009-0275-4

Everingham, M., et al. (2010). The PASCAL Visual Object Classes (VOC) Challenge. International Journal of Computer Vision. The PASCAL Visual Object Classes (VOC)

Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2010). The PASCAL Visual Object Classes (VOC) Challenge. International Journal of Computer Vision, 88(2), 303–338.

Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 580-587.

Glenn Jocher et al. (2023). YOLO by Ultralytics. [Online]. Available: https://github.com/ultralytics/yolov5

Glenn Jocher et al. (2023). YOLO by Ultralytics. GitHub Repository: https://github.com/ultralytics/yolov5

Jocher, G., et al. (2023). YOLO by Ultralytics. GitHub. https://github.com/ultralytics/yolov5

Liu, H., Chen, Y., Wang, Y., & Wang, J. (2019). Helmet detection based on improved YOLOv3 model. Proceedings of the 2019 IEEE 4th International Conference on Image, Vision and Computing (ICIVC), 168-172.

Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. [Online]. Available: https://arxiv.org/abs/1804.02767

Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767. https://arxiv.org/abs/1804.02767

Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. https://arxiv.org/abs/1804.02767

Ruder, S. (2017). An overview of multi-task learning in deep neural networks. ArXiv preprint arXiv:1706.05098. https://arxiv.org/abs/1706.05098

Xie, H., Zhang, J., Li, X., & Yin, B. (2020). Helmet wearing detection based on RCNN. IEEE Access, 8, 144517–144525.

Zhang, Y., Xu, Z., & Wang, H. (2021). Edge-based real-time traffic monitoring using deep learning and edge computing. Sensors, 21(10), 3354.

Zhao, Y., Wang, Z., & Liu, Y. (2017). License plate recognition based on deep learning. 2017 2nd International Conference on Image, Vision and Computing (ICIVC), 745-749.