

Research Article

Fine-Tuned YOLOv8n for Automated Helmet and Seatbelt Compliance Detection in Resource-Constrained Road Safety Enforcement

Nnanna Ekedebe* 

Department of Information Technology, Federal University of Technology, Owerri, Nigeria

Abstract

Road traffic fatalities remain a significant global public health crisis, with the WHO identifying helmet and seatbelt non-compliance as among the leading contributing risk factors. According to the WHO Global Status Report on Road Safety 2023, Nigeria is part of WHO African region, which has the among the highest road traffic death rates in the world at 19.4 deaths per 100,000 people, with enforcement of helmet and seatbelt compliance remaining persistently limited by personnel constraints and the scale of the road network. With over 195,000 kilometers of road network to monitor, manual enforcement by Federal Road Safety Corps officers alone is not scalable, motivating the need for an automated, resource-efficient screening approach. This study fine-tunes a pre-trained YOLOv8n object detection model for helmet and seatbelt compliance detection using the publicly available SeatbeltHelmet dataset comprising 9,166 images across 9 annotated classes. The dataset was split into 8,020 training, 764 validation and 382 test images, and the model was trained for 50 epochs with a batch size of 16 and an image size of 640×640 pixels using the AdamW optimizer on a free-tier Google Colab GPU. The model was evaluated on a held-out test set of 382 images using mAP50, mAP50-95, Precision and Recall. The fine-tuned model achieved an overall mAP50 of 0.599, with Person-Seatbelt achieving the strongest per-class performance at mAP50 of 0.912 and Person-NoSeatbelt at 0.847. Helmet and No-Helmet detection performed comparatively weaker, at mAP50 of 0.537 and 0.380 respectively, reflecting the smaller object size and visual similarity of helmets at typical traffic camera distances. These results demonstrate that a lightweight fine-tuned object detection model is viable for automated road safety compliance screening in resource-constrained African settings where manual monitoring alone is insufficient to address the scale of the problem.

Keywords

Helmet and Seatbelt Compliance Detection, YOLOv8, Road Safety Enforcement, Object Detection, Resource-Constrained Deployment

1. Introduction

For a long time, road traffic fatalities have represented a persistent global public health crisis, with the non-use of seatbelts and helmets being among the leading contributing

factors. According to a Lancet study drawing on data from 185 countries, enforcing seatbelt use could save an estimated 121,083 lives globally per year, while enforcing motorcycle

*Correspondence: Nnanna Ekedebe (nnanaekedeb@gmail.com)

Received: 13 July 2026; **Accepted:** 24 July 2026; **Published:** 18 September 2026



Copyright: © The Author(s), 2026. Published by Science Publishing Group. This is an **Open Access** article, distributed under the terms of the Creative Commons Attribution 4.0 License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

helmet use could save a further 51,698 lives annually [1]. The WHO further notes that seatbelt use by front-seat passengers reduces the risk of fatal injury by 45 to 50 percent, and reduces the risk of death and serious injury among rear-seat occupants by 25 percent [2].

According to the WHO Global Status Report on Road Safety 2023, Nigeria ranks amongst the highest road traffic death rates in Africa at 21.4 deaths per 100,000 people [2]. With motorcycles, locally known as "okada," being one of the most popular modes of transport, this statistic makes helmet compliance particularly critical. In the case of seatbelts, which have been mandated in Nigeria since 2003 under the FRSC Act, compliance has remained persistently low especially amongst commercial vehicle drivers.

Nigeria has over 195,000 kilometers of road networks, making the current manual method of enforcement, where FRSC officers' conduct roadside checks and issue fines for non-compliance not scalable across the full road network. This approach is further limited by the availability of personnel, and the fines currently imposed are insufficient to deter violations consistently.

Computer vision and deep learning present a solution to the scale and consistency problem. Unlike manual checks, an automated system can process images and video frames continuously and consistently without personnel constraints. Object detection models like YOLO can identify whether riders are wearing helmets or whether drivers are wearing seatbelts directly from images, making automated compliance screening possible.

This study aims to fine-tune a pre-trained YOLOv8n object detection model for helmet and seatbelt compliance detection, evaluate it using standard object detection metrics including mAP, precision, and recall, and demonstrate its relevance to road safety enforcement in resource-constrained settings such as Nigeria and Africa as a whole. The paper follows the following structure: section two reviews the related work, section three describes the methodology, section four presents the results, section five the discussion, and section six the conclusion.

2. Related Work

Research into automated helmet and seatbelt compliance detection has grown significantly in recent years, driven by the need for scalable enforcement systems. The field has consistently moved from traditional image processing methods toward CNN and YOLO based approaches. This section reviews the most relevant work and identifies where this study fits in.

Yung et al. conducted a key comparative study evaluating YOLOv5, YOLOv6, and YOLOv7 for safety helmet detection, demonstrating that newer YOLO versions consistently improved on their predecessors in both detection accuracy and speed [3]. The limitation however was that the study focused on construction site helmets rather than motorcycle helmet

compliance on public roads, and did not extend to YOLOv8. Similar architecture-focused improvements have since been explored: Chen et al. proposed a lightweight helmet detection algorithm built on an improved YOLOv4 [4], Zhao et al. introduced BDC-YOLOv5, a helmet detection model built on an improved YOLOv5 backbone [5], and Liu et al. presented a helmet-wearing detection algorithm based on improved YOLOv5 [6]. Like Yung et al., these studies remain centered on construction and industrial helmet detection rather than road traffic enforcement.

More recent work by Wei et al. applied deep learning specifically to helmet wearing detection using an improved YOLOv8 based approach, achieving strong detection performance and confirming YOLOv8 as a superior base model for helmet detection tasks [7]. The drawback was that the study focused on industrial and construction settings rather than road traffic enforcement and helmet compliance on public roads. Building on YOLOv8 as a base, Khalil et al. proposed a YOLOv8-driven smart helmet detection system aimed specifically at traffic safety monitoring [8], while Zhang et al. introduced a lightweight YOLOv8 variant for safety helmet detection with reduced computational overhead [9]. While Khalil et al.'s work extends helmet detection toward traffic contexts, neither this nor Zhang et al.'s work addresses seatbelt compliance or combines both detection tasks in a single model.

The IEEE SCEECs 2023 paper on real-time seatbelt detection using YOLO is more directly relevant to this study. It determined whether a driver's seatbelt was fastened from camera footage using a YOLO based model, demonstrating that YOLO is well suited for compliance detection tasks [10]. The limitation was that it focused on in-cabin camera setups rather than external traffic surveillance, limiting its applicability to large scale road monitoring.

Sutikno et al. represents the most directly relevant prior work, a 2023 study that used YOLOv8 specifically for automated detection of drivers and passengers without seatbelts, demonstrating that YOLOv8 could effectively classify seatbelt compliance from vehicle images [11]. The limitation was that the study was conducted in a controlled environment with limited image diversity and did not address the challenges of varied lighting, camera angles, and road conditions typical of real-world traffic surveillance. Related seatbelt-focused work includes Gu et al., who proposed a gated Bi-LSTM model with part-to-whole attention on diagonally sampled patches for seat belt detection [12], and a 2025 follow-up by Sutikno et al. that combined CLAHE preprocessing with YOLOv11 for seat belt compliance detection [13]. Together, these confirm the continued viability of the YOLO family for seatbelt-specific detection, though both still rely on relatively controlled or pre-processed imagery rather than raw external traffic surveillance footage.

Building on this body of work, this study fine-tunes YOLOv8n for combined helmet and seatbelt compliance detection in a single model. While existing studies have largely focused on Asian and Western road environments like India,

China, Thailand and similar contexts, the application of automated compliance detection to road safety enforcement in Nigeria and sub-Saharan Africa has not been directly addressed in the reviewed literature, bringing this approach to a region where manual enforcement capacity is limited and road fatality rates remain among the highest on the continent.

3. Methodology

This section outlines the steps taken to fine-tune a pre-trained YOLOv8n model for helmet and seatbelt compliance detection. It covers the following: dataset, data preprocessing, model selection, training configuration, and evaluation metrics.

3.1. Dataset

The dataset used in this study is the SeatbeltHelmet dataset by helmet-seatbelt-detection, a publicly available annotated dataset comprising 9,166 images across 9 annotated classes covering helmet and seatbelt compliance across varied road and traffic conditions. The dataset comes pre-annotated with bounding boxes in YOLO-compatible format, meaning no manual labeling was required. The dataset was split into 8,020 training images, 764 validation images and 382 test images [14].

3.2. Data Preprocessing

Before training, the images were resized to 640×640 pixels, the standard input size expected by YOLOv8. Data augmentation was applied, that is, artificially expanding the dataset by flipping, rotating, and adjusting the brightness of existing images so the model learns to handle varied real-world lighting and road conditions.

3.3. Model Selection and Architecture

YOLOv8n was selected as the object detection model for this study, as it represents the most current and optimized lightweight variant of the YOLO family, offering strong detection performance without requiring large computational resources [15]. The model is fine-tuned rather than trained from scratch, it begins with weights pre-trained on the COCO dataset and is then adapted to the helmet and seatbelt compliance detection task. The model handles 9 detection classes across the dataset, with the five compliance-relevant classes being Helmet, No-Helmet, Person-Seatbelt, Person-NoSeatbelt and Seatbelt.

3.4. Training Configuration

The model was trained using the following configuration: a batch size of 16 images per step, an image input size of 640×640 pixels, and training was run for 50 epochs. The

optimizer was automatically selected by YOLOv8 as AdamW. All training was performed on Google Colab using a free-tier GPU instance, which provided sufficient computational resources for fine-tuning a lightweight YOLOv8n model on this dataset.

3.5. Evaluation Metrics

The model's performance was evaluated using standard object detection metrics mAP@0.5, mAP@0.5-0.95, Precision, Recall, and IoU with a threshold of 0.5. Since the model handles 9 detection classes, results are reported both overall and per class to give a complete picture of detection performance across all compliance-relevant classes.

4. Results

This section presents the results obtained from fine-tuning YOLOv8n on the SeatbeltHelmet dataset, evaluated on a held-out test set of 382 images across 9 classes. Performance is reported using mAP50, mAP50-95, Precision and Recall both overall and per class. The analysis and commentary focus specifically on the five compliance-relevant classes Helmet, No-Helmet, Person-Seatbelt, Person-NoSeatbelt and Seatbelt.

4.1. Overall Performance

The fine-tuned model achieved an overall Precision of 0.599, Recall of 0.583, mAP50 of 0.599 and mAP50-95 of 0.370, demonstrating the model's general ability to detect compliance-relevant objects across varied traffic and road conditions. It is important to note that the gap between mAP50 and mAP50-95 is expected for a lightweight model like YOLOv8n and is consistent with results reported in similar fine-tuning studies. For a road safety enforcement context such as Nigeria's, where the primary goal is reliable flagging of violations rather than pixel-perfect localization, mAP50 is the more operationally relevant metric and 0.599 represents a functional detection baseline.

4.2. Per-Class Performance

It was observed from the results that Person-Seatbelt achieved the highest mAP50 at 0.912 and Person-NoSeatbelt at 0.847, indicating that the model detects seatbelt compliance and violations with high confidence. The Seatbelt class also performed strongly at mAP50 0.780, further reinforcing that seatbelt detection is the strongest aspect of the fine-tuned model. Helmet detection achieved mAP50 of 0.537 and No-Helmet achieved 0.380 lower but still functional, most likely attributable to the smaller apparent object size and the visual similarity between helmeted and unhelmeted riders at typical traffic camera distances. Among the auxiliary classes, Bike achieved mAP50 of 0.883 and Car achieved 0.847, confirming that the model handles vehicle context detection well. Table 1

presents the precision, recall, mAP50 and mAP50-95 achieved by the fine-tuned model across all 9 classes on the 382-image held-out test set.

Table 1. Per-Class Test Set Performance.

Class	Images	Instances	Precision	Recall	mAP50	mAP50-95
All	382	1453	0.599	0.583	0.599	0.370
Bike	158	338	0.729	0.953	0.883	0.631
Car	87	146	0.689	0.870	0.847	0.683
Helmet	93	126	0.538	0.563	0.537	0.204
Mobile	8	9	0.000	0.000	0.0106	0.0025
No-Helmet	119	222	0.445	0.407	0.380	0.130
NoPlate	72	142	0.433	0.148	0.195	0.070
Person-NoSeatbelt	98	116	0.842	0.735	0.847	0.577
Person-Seatbelt	136	181	0.894	0.843	0.912	0.632
Seatbelt	129	173	0.824	0.731	0.780	0.398

4.3. Training Curves

The mAP50 improved steadily throughout the training process, with validation performance closely tracking training performance, indicating that the model generalised well rather than overfitting. Training and validation loss decreased

consistently across all 50 epochs, confirming stable convergence throughout. The chosen training configuration was, batch size 16, image size 640 and AdamW optimizer which proved appropriate for this dataset and task. Figure 1 presents the training and validation curves, including box loss, classification loss, precision, recall and mAP recorded across all 50 epochs.




Figure 1. YOLOv8n Training Performance Curves over 50 Epochs.

4.4. Confusion Matrix

It was observed that Person-Seatbelt and Person-NoSeatbelt

showed strong diagonal values with minimal confusion between them, the most operationally critical distinction for compliance enforcement. Confusion was more pronounced between the Helmet and No-Helmet classes, consistent with

their lower mAP50 scores, and is most likely attributable to the smaller object size and visual similarity between helmeted

and unhelmeted riders at distance. Figure 2 presents the normalized confusion matrix obtained on the 382-image test set.




Figure 2. Normalized Confusion Matrix for Helmet and Seatbelt Compliance Detection.

5. Discussion

The overall results of this study demonstrate that YOLOv8n can be effectively fine-tuned for helmet and seatbelt compliance detection even in a multi-class setting. Achieving an overall mAP50 of 0.599 on a 9-class dataset using a light-weight fine-tuned model represents a practically viable result for resource-constrained deployment. While the model is not perfect, it is a functional one for an enforcement support tool in a setting like Nigeria where no automated compliance detection system currently exists, a functional baseline is a meaningful contribution.

The most notable finding requiring explanation is the clear performance gap between seatbelt detection and helmet detection. Person-Seatbelt achieved mAP50 of 0.912 and Person-NoSeatbelt achieved 0.847, significantly outperforming Helmet at 0.537 and No-Helmet at 0.380. This gap can be attributed to the fact that seatbelts are detected in the context of the full upper body of a person inside a vehicle, providing a much larger and more consistent visual region for the model to work with, compared to helmets which are smaller objects that vary significantly in color, shape and size across different brands and types.

The seatbelt detection results of this study are comparable to and in some cases stronger than results reported in prior work. Sutikno et al. reported strong seatbelt classification

performance in a controlled environment, whereas this study achieves comparable results on a more diverse real-world dataset. In the case of Wei et al., their improved YOLOv8 approach focused on industrial settings where helmets are more uniform and larger relative to the frame, which partly explains why their helmet detection performance exceeded that of this study. It is important to note that these performance differences are largely attributable to dataset characteristics and deployment context rather than model architecture limitations.

Relative to the broader related work, this study's contribution is distinct in scope rather than raw performance. Yung et al. compared successive YOLO versions for helmet detection but did not extend to YOLOv8 or address seatbelt compliance [3], while the SCEECs seatbelt study relied on in-cabin camera footage rather than external traffic surveillance [10], limiting its applicability to large-scale road monitoring. This study's use of a single fine-tuned model to detect both helmet and seatbelt compliance from external traffic imagery therefore addresses a combination of tasks, deployment setting, and geographic context not covered together in the reviewed literature.

Regarding limitations, the most significant is the class imbalance in the dataset, particularly the No-Helmet class which underperformed relative to the Helmet class, which raises a concern for enforcement since detecting the violation is operationally more important than detecting compliance, in hindsight, training on a dedicated dataset containing only helmet

and seatbelt classes rather than a 9-class dataset would likely have yielded stronger per-class performance on the compliance-relevant classes, as the model's learning capacity would not have been distributed across auxiliary classes. It is also worth noting that the dataset was not collected specifically in Nigerian or African road conditions, the images are predominantly from Asian traffic environments, which may limit direct transferability to Nigerian road settings where lighting conditions, camera angles and vehicle types differ. Despite these limitations, the model's strong seatbelt detection performance has immediate practical relevance. Seatbelt non-compliance remains persistently high among commercial vehicle drivers in Nigeria, and the model's ability to reliably flag both compliant and non-compliant cases means it could be deployed on existing CCTV infrastructure at major road checkpoints to automate that screening, reducing the personnel burden on FRSC officers and improving the consistency of enforcement across the road network.

6. Conclusion

This study fine-tuned a pre-trained YOLOv8n object detection model for helmet and seatbelt compliance detection using the publicly available SeatbeltHelmet dataset comprising 9,166 images across 9 classes. The model was evaluated on a held-out test set of 382 images using standard object detection metrics such as mAP50, mAP50-95, Precision and Recall.

The fine-tuned model demonstrated strong seatbelt compliance detection performance, confirming that YOLOv8n is a viable lightweight model for automated compliance screening in resource-constrained settings. This study also demonstrated the feasibility of fine-tuning a single model for combined helmet and seatbelt detection simultaneously, rather than requiring two separate dedicated models for each compliance type. Overall, this study brings automated compliance detection to a regional context where manual enforcement capacity is limited and road fatality rates remain among the highest on the continent.

Future work should focus on collecting or curating a dataset specifically from Nigerian and African road conditions to improve direct transferability and real-world performance in the target deployment environment. A dedicated helmet and seatbelt only dataset without auxiliary classes like Mobile and No-Plate, would likely yield stronger per-class performance on the compliance-relevant classes specifically.

Abbreviations

YOLO	You Only Look Once
mAP	Mean Average Precision
IoU	Intersection over Union
CNN	Convolutional Neural Network
GPU	Graphics Processing Unit
WHO	World Health Organization

FRSC	Federal Road Safety Corps
CCTV	Closed-Circuit Television

Acknowledgments

I acknowledge the support provided by Federal University of Technology Owerri (FUTO), through the Information Communication Technology (ICT) research center.

Author Contributions

Nnanna Ekedebe: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing

Conflicts of Interest

The author declares no conflicts of interest.

References

- [1] Vecino-Ortiz AI, Nagarajan M, Elaraby S, Guzman-Tordecilla DN, Paichadze N, Hyder AA. "Saving lives through road safety risk factor interventions: global and national estimates." *Lancet*. 2022; 400(10347): 237–250. [https://doi.org/10.1016/S0140-6736\(22\)00918-7](https://doi.org/10.1016/S0140-6736(22)00918-7)
- [2] World Health Organization. "Global Status Report on Road Safety 2023." Geneva: WHO; 2023.
- [3] Yung, N. D. T., Wong, W. K., Juwono, F. H., Sim, Z. A. "Safety Helmet Detection Using Deep Learning: Implementation and Comparative Study Using YOLOv5, YOLOv6, and YOLOv7." *Proc. IEEE GECOST*. 2022: 164–170. <https://doi.org/10.1109/GECOST55694.2022.10010490>
- [4] Chen, J., Deng, S., Wang, P., Huang, X., Liu, Y. "Lightweight helmet detection algorithm using an improved YOLOv4." *Sensors*. 2023; 23(3): 1256. <https://doi.org/10.3390/s23031256>
- [5] Zhao, L., Tohti, T., Hamdulla, A. "BDC-YOLOv5: a helmet detection model employs improved YOLOv5." *Signal, Image and Video Processing*. 2023; 17(8): 4435–4445. <https://doi.org/10.1007/s11760-023-02674-8>
- [6] Liu, Y., et al. "Helmet wearing detection algorithm based on improved YOLOv5." *Scientific Reports*. 2024; 14: 8768. <https://doi.org/10.1038/s41598-024-59486-8>
- [7] Wei, L., Liu, P., Ren, H., Xiao, D. "Research on helmet wearing detection method based on deep learning." *Scientific Reports*. 2024. <https://doi.org/10.1038/s41598-024-57433-z>
- [8] Khalil, N., Rashid, M. H., Rehman, S. A., Yasir, M., Hussain, M. "Improved Traffic Safety with YOLO-v8 Driven Smart Helmet Detection." *VFAST Transactions on Software Engineering*. 2025; 13(3): 29–42. <https://doi.org/10.21015/vtse.v13i3.2184>

- [9] Zhang, L., Ma, H., Huang, J., Zhang, C., Gao, X. "An Improved Lightweight Safety Helmet Detection Algorithm for YOLOv8." *Computers, Materials & Continua*. 2025; 83(2): 2245–2265. <https://doi.org/10.32604/cmc.2025.061519>
- [10] "Real-time seatbelt detection using YOLO deep learning model." *Proc. IEEE SCEECs*. 2023. <https://doi.org/10.1109/SCEECs57921.2023.10063114>
- [11] Sutikno, Sugiharto, A., Kusumaningrum, R. "Automated detection of driver and passenger without seat belt using YOLOv8." *Int. J. Adv. Comput. Sci. Appl.* 2023; 14(11): 806–813. <https://doi.org/10.14569/IJACSA.2023.0141181>
- [12] Gu, X., Lu, Z., Ren, J., Zhang, Q. "Seat belt detection using gated Bi-LSTM with part-to-whole attention on diagonally sampled patches." *Expert Systems with Applications*. 2024; 252: 123784. <https://doi.org/10.1016/j.eswa.2024.123784>
- [13] Sutikno, Sugiharto, A., Kusumaningrum, R. "Enhanced Automatic License Plate Detection and Recognition using CLAHE and YOLOv11 for Seat Belt Compliance Detection." *Engineering, Technology & Applied Science Research*. 2025; 15(1): 20271–20278. <https://doi.org/10.48084/etasr.9629>
- [14] Helmet + Seatbelt Detection. "SeatbeltHelmet Dataset." Roboflow Universe; 2023. Available: <https://universe.roboflow.com/helmet-seatbelt-detection/seatbelthelmet>
- [15] Ultralytics. "YOLOv8 Documentation." 2023. Available: <https://docs.ultralytics.com>