

Detecting, Tracking and Counting Motorcycle Rider Traffic Violations on Unconstrained Roads

Aman Goyal* Dev Agarwal*
Anbumani Subramanian C.V. Jawahar Ravi Kiran Sarvadevabhatla Rohit Saluja

Centre For Visual Information Technology (CVIT), IIIT Hyderabad, INDIA

<https://github.com/iHubData-Mobility/public-motorcycle-violations>

Abstract

In many Asian countries with unconstrained road traffic conditions, driving violations such as not wearing helmets and triple-riding are a significant source of fatalities involving motorcycles. Identifying and penalizing such riders is vital in curbing road accidents and improving citizens' safety. With this motivation, we propose an approach for detecting, tracking, and counting motorcycle riding violations in videos taken from a vehicle-mounted dashboard camera. We employ a curriculum learning-based object detector to better tackle challenging scenarios such as occlusions. We introduce a novel trapezium-shaped object boundary representation to increase robustness and tackle the rider-motorcycle association. We also introduce an amodal regressor that generates bounding boxes for the occluded riders. Experimental results on a large-scale unconstrained driving dataset demonstrate the superiority of our approach compared to existing approaches and other ablatable variants.

1. Introduction

Automated road surveillance has become increasingly crucial as road crashes have become the 8th leading cause of death worldwide. A World Health Organization study on road safety [18, 25] claims that violations lead to 1.35 million fatalities and affect 50 million people yearly. Another recent report by World Bank [2] mentions that more than 50% of road fatalities involve two-wheeler vehicles, also showing that 'no helmet' and triple-riding (more than two riders) violations are common causes. Studies carried out in Asian countries also account for two-wheeler vehicles among the significant share for road fatalities [19, 25, 31].

*Authors contributed equally to this research.

{aman.goyal1099, devagarwal.iitkpg}@gmail.com,
{anbumani, jawahar, ravi.kiran}@iiit.ac.in,
rohit.saluja@research.iiit.ac.in



Figure 1. (a) Frame level rider and motorcycle predictions from our curriculum learning-based detector, (b) Outputs from violation detection and counting module, (c) Violation detections and counts overlaid on the frame; triple-riding and helmet violations in orange trapezium (note rider/head occlusions) and red boxes. Refer Fig. 2 for remaining color schemes.

Often, static cameras may not be present on the majority of the streets. Installing cameras everywhere may not be an economically viable and sustainable solution. Therefore, we adopt an approach that operates on videos from a dashboard camera mounted within a police vehicle.

Existing methods for triple-riding violations [15, 21] involve heuristic rules applied to outputs of rider detectors

and show qualitative results for a few samples. In addition, these methods are not robust to occlusions and crowded scenes. In contrast, we propose a novel trapezium representation for identifying triple riding violations. The representation largely reduces false positives and is robust as it does not require any heuristic rules. For helmet violations, a majority of the existing works [4, 6, 11, 22, 23] perform classification over upper portions of predicted riders. Due to the unavailability of context in the Regions Of Interest (ROIs), such works are not robust in crowded road scenarios. Our proposed technique performs well even in crowded scenarios, as we use the ROIs of motorcycle-driving instances containing sufficient context for *helmet/no-helmet* detector to provide accurate results. As an integral part of our pipeline, we use a curriculum learning-based detector to predict motorcycle and rider boxes to overcome inter-class confusion due to overlapping regions among the two classes. Recent methods in curriculum learning mainly focus on object classification [9, 16, 29] with very few works on detection [24, 28] which involve handling intra-class scale variations or weakly and semi-supervised training.

To the best of our knowledge, no existing solution jointly tackles helmet and triple-riding violations for crowded scenes and provides robustness in occluded rider scenarios. Fig. 1 shows outputs of various components of the proposed architecture wherein we train a curriculum learning-based model to predict motorcycles and riders over the video frame followed by a violation detection and counting module. Our contributions in this paper are:

1. A novel dataset pre-processing pipeline involving an amodal regressor to generate boundary representations even for occluded riders.
2. An innovative trapezium-shaped box regressor for associating a motorcycle with its corresponding riders.
3. A curriculum learning-based architecture to jointly detect, track and count triple-riding and helmet violations on unconstrained road scenes.

2. Related Works

Triple-riding Violations: There are very few existing methods in the space of triple-riding violations. The method proposed by [15] detects violation on the lower portion of the motorcycle, is not robust to false positives as no counting algorithm is involved, and all the images used for qualitative analysis have similar camera angles and lighting conditions. The work by [21] focuses on using an association algorithm based on Euclidean distance after detecting motorcycle and rider boxes to count the number of riders on a motorcycle. It fails in fairly crowded road scenarios and lacks qualitative analysis on a large-scale dataset. However, our proposed approach based on data pre-processing

to handle occlusions and trapezium-shaped bounding boxes can robustly detect and track triple-riding violations in each video frame.

Helmet Violations: Existing works on helmet violations by [4, 6, 11, 22, 23] classify the ROIs generated from the upper portion of the detected riders. Such approaches are not robust to false negatives and false positives in the form of truncated heads in the cropped ROIs and riders of other motorcycles, respectively. The work by [27] uses conventional techniques utilizing hand-engineered features such as scale-invariant feature transform, histogram of gradients etc. [5, 7, 8, 14] do not perform well on large-scale unconstrained and crowded road scenarios. Works of [12] focus on detecting helmet violations directly from complete scene images. Such an approach leads to more false negatives and cannot detect violations for riders farther away from the camera. Apart from this, work by [32] requires foreground segmentation as a step before detecting *helmets/no-helmets*. We obtain robust performance by involving ROIs containing motorcycles and corresponding riders. Though works such as [4, 10, 13, 20], robustly detect helmet violations, they do not possess a triple-riding detection module. On the contrary, our proposed approach can jointly detect, track, and count triple-riding and helmet violations.

Curriculum Learning: The researchers have widely focused on applying curriculum learning-based techniques to object classification [9, 16, 29]. Existing curriculum-learning-based detection works [24, 28] either concentrate on improving the performance by handling intra-class scale variations or involve weakly and semi-supervised training. Our method utilizes the curriculum learning-based object detector to obtain robust detections of motorcycles and riders and avoid inter-class confusion due to high overlap between the bounding boxes of the two classes. The object detector further helps in robustly predicting triple riding and helmet violations.

3. Dataset

We train various models for identifying motorcycle violations. We use a subset of India Driving Dataset (IDD) [26] extracted by selecting images with motorcycles and riders followed by filtering out small bounding boxes based on an area threshold of 900 squared pixels, as low-resolution boxes add noisy samples to the training data. We annotate *helmet class*, *no-helmet class*, and *trapezium-shaped driving instance class* (see Fig. 2 b-d, i-k). We use conventional rectangular boxes for helmet violations and propose trapezium bounding boxes to detect triple-riding violations. We now discuss how we pre-process the dataset.

Data Pre-Processing: We note two crucial issues in the existing IDD dataset annotations, which pose a problem in accurately detecting the riders and subsequent triple-riding and helmet violations. i) Some IDD samples had multiple

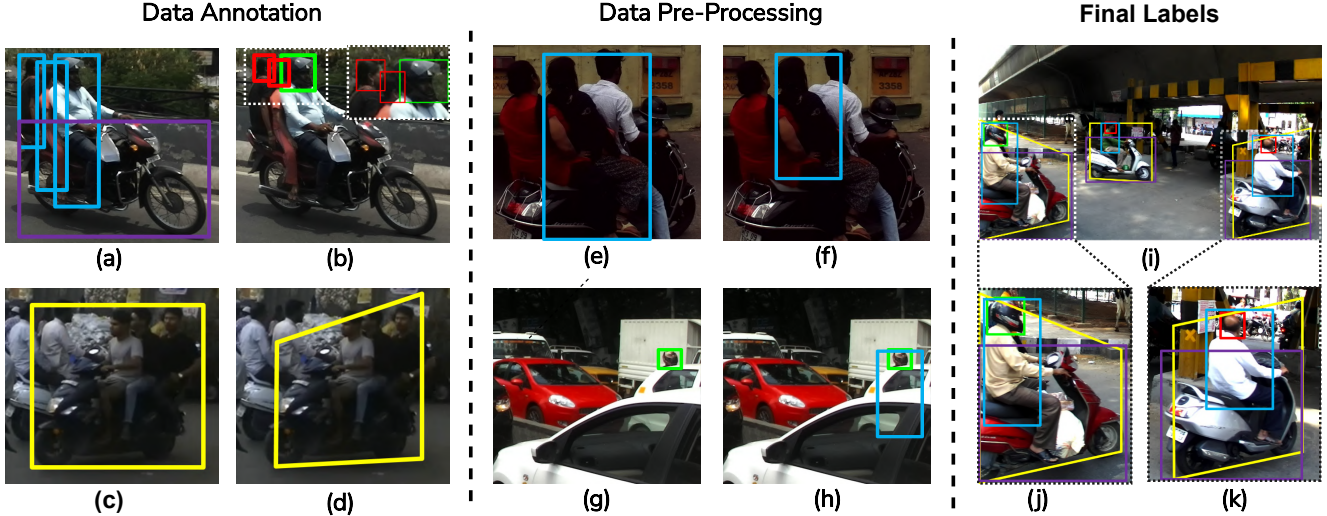


Figure 2. (a) Existing annotations of rider/motorcycle in IDD dataset, (b) *helmet/no-helmet* annotations additionally annotated by us, (c) Conventional *rider-motorcycle* instance representation. (d) Novel trapezium box representation for the *rider-motorcycle* instance in contrast to (c); reduces false-positives from nearby riders, (e) Rider box annotation present in IDD dataset, (f) Manually corrected rider boxes, (g) Annotated helmet box on an occluded rider, (h) Generating boxes for occluded riders with amodal regressor using helmet boxes (refer Sec. 3), (i) Sample frame with final labels & its crops (j), (k). [color scheme: red-helmet violation, green-helmet, orange-triple-riding violation, yellow-driving instance, purple-motorcycle, blue-rider]

riders in a single bounding box (due to inclusion of riders’ legs), as shown in Fig. 2 (e). ii) There were frequent rider occlusions, to the extent that only the head portion is visible, as shown in Fig. 2 (g). We overcome such issues by:

(i) Processing the Rider Boxes: Rider bounding boxes are manually processed such that each bounding box contains only one rider. Fig. 2 (f) shows the output of such processing for a rider in Fig. 2 (e).

(ii) Generating Boxes for Rider Occlusions with Amodal Regressor: To obtain full bounding boxes for occluded riders, we train a two-layered deep network with fully connected layers, which we refer to as amodal regressor. With *helmet/no-helmet* bounding box as input and corresponding *rider* bounding box as output, the amodal regressor has 16 and 64 nodes in the respective hidden layers. We train it on non-occluded riders in our data with *relu* activation and *learning rate* of 0.001. We use the trained model to generate boxes for occluded riders in our data. Fig. 2 (h) depicts the example of a bounding box generated for an occluded rider.

4. The Pipeline

This section discusses the proposed pipeline, shown in Fig. 3, for identifying triple-riding and helmet violations. The videos from a camera mounted on a moving vehicle are input to a Curriculum Learning (CL) based object detection module, predicting riders and motorcycles (refer to Fig. 3 a). Thereafter, we send each motorcycle prediction and riders’ predictions intersecting with it to the next stage

(see Fig. 3 b-c). A novel trapezium box regressor, shown in Fig. 3 (d), then generates bounding boxes over each driving instance. Next, a triple-riding detection module calculates the number of riders within each trapezium box to identify the possibility of a triple-riding violation. A helmet violation detector also operates parallelly (see Fig. 3 d), which works on the Regions Of Interest (ROIs) extracted from the previous detector’s predictions (motorcycle and riders boxes). Finally, we modify the predictions from DeepSORT [30] to jointly track the rectangular *helmet/no-helmet* boxes and trapezium instance boxes as shown in Fig. 3 (f-g). We now elaborate on each step of the proposed pipeline.

4.1. Curriculum Learning-based Detector

A generic two-class *rider-motorcycle* detector performs poorly due to overlapping regions between a motorcycle and its riders. The overlapping regions lead to inter-class confusion and, in turn, lead to false negatives due to lower confidence of one or more overlapping predictions and Non-Maximum Suppression (NMS). As shown in Fig. 3 (a), we propose a CL based detector to avoid such issues. The model is first trained on only motorcycle instances and then retrained after adding rider instances. As we will see in Sec. 5, the recall improves significantly compared to the generic model, as an outcome of a reduction in false positives.

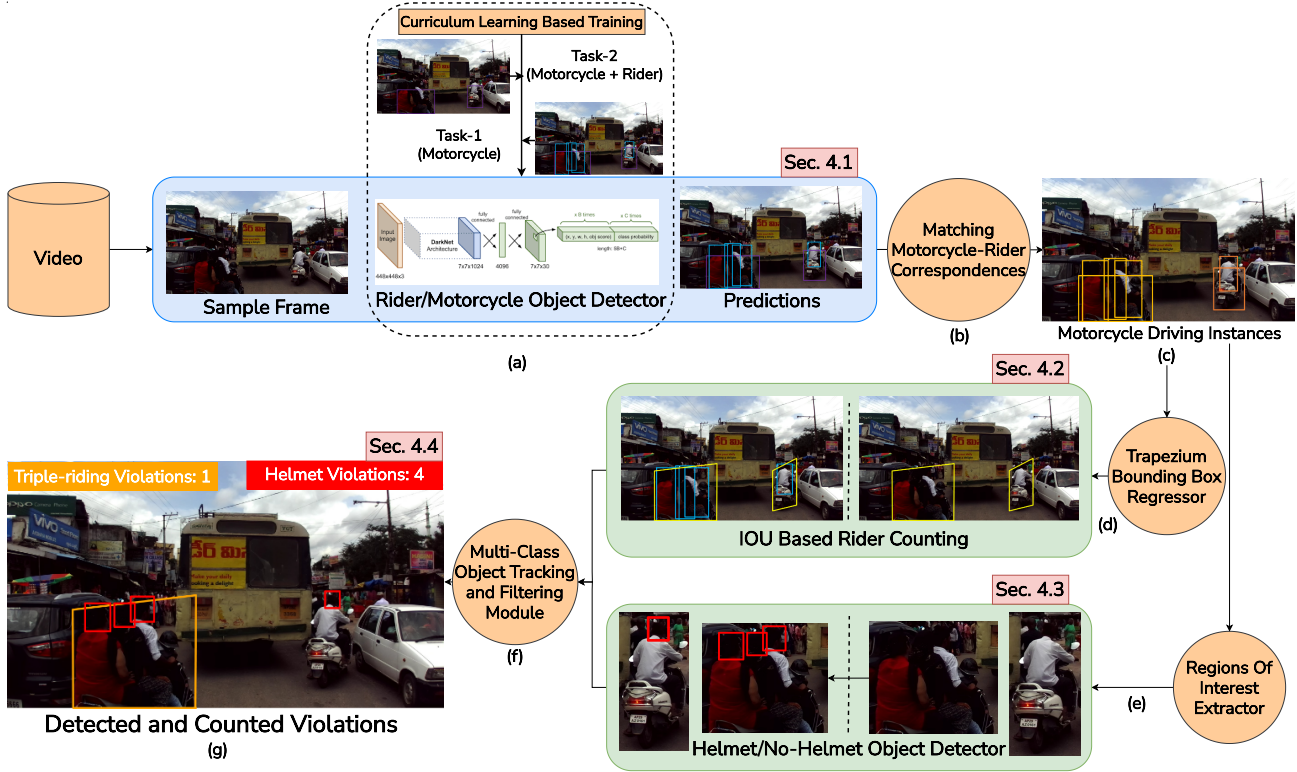


Figure 3. Our pipeline: (a) Video frames are input to Curriculum Learning-based rider/motorcycle detector, (b) The intersecting motorcycles and riders are matched based on IOUs, (c) Each Rider-Motorcycle instance is shown in a different color, (d) Trapezium boxes from a regressor are passed to Triple-riding violation module, (e) Extracted ROIs are passed to a *helmet/no-helmet* detector, (f) Violations are tracked and counted over the video, (g) Violations/counts are shown in red and orange colors.




Figure 4. Different IOU-based correlation approaches to associate riders to their motorcycle. **Note:** (a) demonstrates the use of motorcycle boxes (purple) for IOU based association with rider boxes (blue), but the former have low IOU with all the latter, (b) demonstrates the use of instance boxes (yellow), but the box closer to the camera has high IOU with all the three riders, (c) **[Proposed]** approach uses a trapezium-shaped instance box (yellow) has high IOU with its two true riders and low IOU with the rider from another motorcycle.

4.2. Identifying Triple-riding Violations

This section discusses the trapezium box regressor (see Fig. 3 d) and the rider counting step for identifying triple-

riding violations in a video frame. Each predicted motorcycle box and rider bounding boxes that intersect with it form the input to the trapezium regressor. The regressor detects driving instances for counting the rider boxes having the highest IOU with the obtained trapeziums; to identify the triple-riding violations. The resulting reduction in false-positive rider correspondences also leads to a massive 55.44% precision gain for triple-riding violations in Tab. 2.

Trapezium Bounding Box Regressor: Our data contains highly unstructured and crowded scenarios. Thus, as is illustrated in Fig. 4 (a), relying on derived IOU-based correlations between riders and their motorcycles become erroneous for accurate identification of triple-riding violations. Similarly, drawing rectangular boxes over such correlated riders and motorcycles also becomes an issue, as in many cases, a significant portion of such boxes cover the background area and intersect with riders of other motorcycles, as shown in Fig. 4 (b). Hence, we utilize novel trapezium bounding boxes which circumscribe the correlated riders and motorcycles more tightly than rectangles, as shown in Fig. 4 (c). We train a two-layered deep network (see Fig. 5 b) having 512 and 256 nodes in the respective hidden layers and use *tanh* activation with *learning rate* of 0.001.

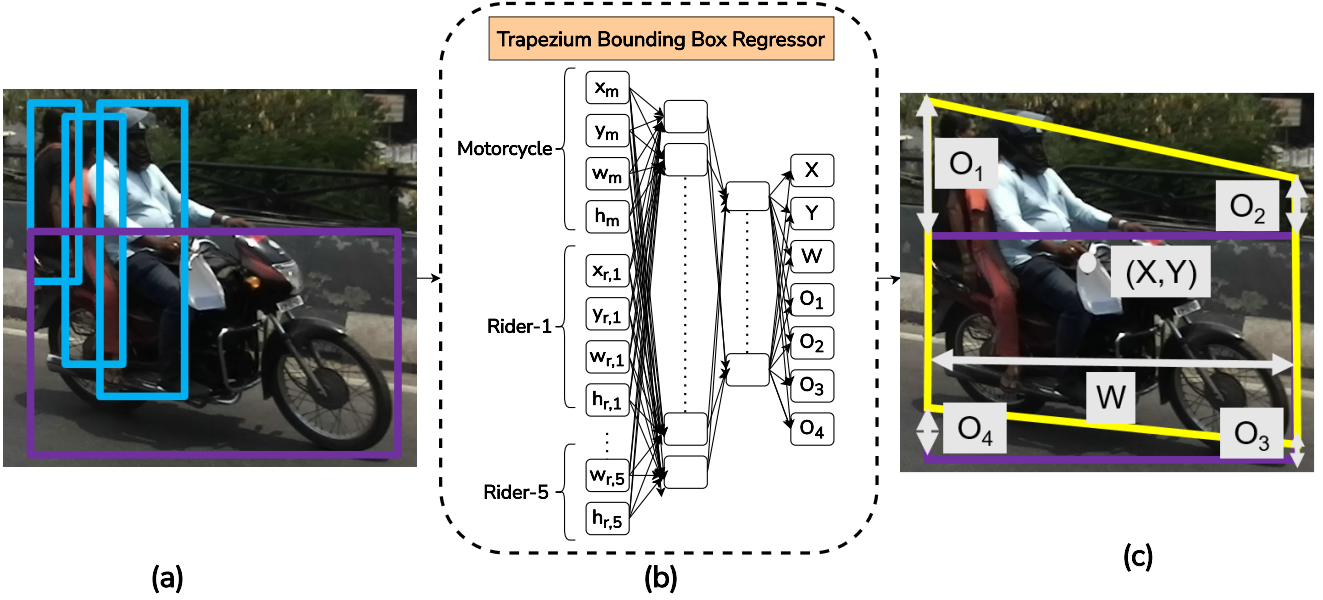


Figure 5. Input (a) and output (c) of the trapezium box regressor (b). Refer to Sec. 4.2 for additional details.

Mean squared error [1] was used as loss function in regression.

Input: As shown in Fig. 5 (a), corresponding motorcycle and riders' boxes are concatenated to form the input of the trapezium regressor. The number of inputs is fixed to 24; the first four values correspond to the motorcycle box and the rest to its riders. If the number of riders is lower than 5 (maximum number of rider space in input as also advocated by [13]), then we initialize the remaining values to 0.

Output: The center of gravity [17] of the trapezium is taken as its central position (X, Y) , the perpendicular distance between the parallel sides (vertical in our case) is taken as its width (W) . For slopes of non-parallel sides, we take the offset of each corner of a trapezium from its corresponding motorcycle box corner as illustrated in Fig. 5 (c).

Equation of Trapezium's Center: Trapezium's centre, (X, Y) which is defined by its n ($= 4$) vertices $(x_0, y_0), (x_1, y_1), \dots, (x_{n-1}, y_{n-1})$ can be calculated by using the following equations. The vertices should be provided as input in clockwise or anti-clockwise order and the polygon should be closed such that the vertex (x_0, y_0) is the same as the vertex (x_n, y_n) :

$$X = \left[\sum_{i=0}^{n-1} (x_i + x_{i+1})(x_i y_{i+1} - x_{i+1} y_i) \right] / 6A \quad (1)$$

$$Y = \left[\sum_{i=0}^{n-1} (y_i + y_{i+1})(x_i y_{i+1} - x_{i+1} y_i) \right] / 6A \quad (2)$$

Where A is a polygon's signed area:

$$A = \left[\sum_{i=0}^{n-1} (x_i y_{i+1} - x_{i+1} y_i) \right] / 2 \quad (3)$$

Counting Riders: Counting the riders with boxes having the highest IOU with each trapezium bounding box are considered to be riding the motorcycle and therefore help identify violations.

One of the advantages of the trapezium-shaped boxes is that they are generic and can represent many other objects such as lane-markings and railway tracks recorded from a front view camera on a vehicle or train. It can benefit from the speed of detection methods and the accuracy of segmentation methods. We note that rider and motorcycle predictions obtained using semantic information are sufficient for the trapezium regressor. The trapeziums are robust and avoid bounding box area and region-based rules [15,21] for different conditions across images.

4.3. Identifying Helmet Violations

As an input to the helmet violation model, the width of the predicted motorcycle and maximum height of predicted rider boxes of a particular motorcycle driving instance are used to extract the relevant ROIs. The cropped ROIs are extended by 10% and then passed to the trained detector to predict the violations at the inference stage. Unlike motorcycle and rider classes, helmet and no-helmet classes have much lesser overlapping due to which a two-class YOLOv4 detector is trained (refer Fig. 3 e and Fig. 2 b).

For all detectors mentioned above, we use *binary cross entropy* and *CIoU (Complete Intersection over Union)* for classification and regression, respectively [3].

Table 1. Evaluation of Triple-riding Violation Detection (rider and motorcycle) & Identification (counting riders on a bike). Note: CL based YOLOv4 (v4) model performs better than the conventional two-class (non-CL) v4 & YOLOv3 (v3) models due to high IOU of riders with motorcycles. Trapezium-shaped boxes further boost the F-score.

S. No.	Method	Rider-Motorcycle Detection		Rider-Motorcycle Association Approach	Violation Identification Scores		
		Base Model	mAP		Precision	Recall	F-score
1	Saumya et al. [21]	v3_non-CL	70.13%	Euclidean Distance	29.00%	50.00%	36.70%
2	Ours	v4_non-CL	73.62%	Euclidean Distance	30.47%	51.26%	38.22%
3				Rider-Motorcycle Box	33.73%	53.84%	41.47%
4				Rectangular-shaped Instance Box	41.66%	57.69%	48.38%
5				Trapezium-shaped Instance Box	73.80%	59.61%	65.95%
6	Proposed	v4_CL	82.61%	Trapezium-shaped Instance Box	84.44%	73.07%	78.34%

Table 2. Helmet Violation Detection and Identification Performance at Rider-level (rows 1-5) and Instance-level (rows 6-7). Note: CL based YOLOv4 (v4) model trained on *helmet/no-helmet* classes performs poorer than the conventional two-class (non-CL) v4 model due to low IOU between helmets and no-helmets. Refer Tab. 1 caption and detection performance for acronyms and *rider-motorcycle* mAP scores. N/A: Not Applicable.

S. No.	Method	Helmet/No-Helmet Detection		ROI Extraction Approach	Violation Identification Scores		
		Base Model	mAP		Precision	Recall	F-score
1	Rithish et al. [20]	v4_non-CL	90.00%	Rider Instance Crop	53.21%	81.02%	64.23%
2	Ours			Upper Half Instance Crop	49.9%	76.40%	60.36%
3	Ours			Full Resolution Image	71.36%	90.13%	79.65%
4	Ours	v4_CL	83.5%	Rider-Motorcycle Instance Crop	68.56%	83.25%	75.19%
5	Proposed	v4_non-CL	90.00%	Rider-Motorcycle Instance Crop	77.86%	92.23%	84.43%
6	Chairat et al. [4]	v3_GoogleNet	N/A	Upper Half Instance Crop	84.61%	54.45%	66.25%
7	Proposed	v4_non-CL	90.00%	Rider-Motorcycle Instance Crop	99.01%	95.23%	97.08%

4.4. Tracking

As shown in Fig. 3 (f-g), we modify DeepSORT [30] predictions to jointly track *helmet/no-helmet* boxes and trapezium boxes in traffic videos. DeepSORT tracks only rectangular-shaped boxes, so the trapezium box is tracked via its corresponding motorcycle bounding box.

5. Experiments and Results

We use a dataset of 1281 images containing 907 riders without helmets and 98 instances of triple-riding violations amongst 3313 riders and 4573 motorcycles. We use 70:30 as a train:test split for our data. Our test set consists of 310 images, where 260 are from the filtered IDD dataset (refer to Sec. 3), and 50 are randomly chosen images from web search to enrich the testing for rare triple-riding violations. It has 234 riders with helmet violations and 42 instances of triple-riding violations among 758 riders and 685 motorcycles. For the curriculum learning-based model, the initial learning rate is set to 0.001 every time a class is added with a decay of 10. For the helmet violation detector as well, the learning rate is initialized with 0.001 with the decay of 10. We now present the results of our models on triple-riding and helmet violations, along with the previous works trained and tested on our data. We also present ablation studies, which show the performance boost by utilizing the proposed pre-processing and curriculum

learning-based object detector.

5.1. Comparison with Existing Methods

Triple-riding Violations: As shown in the first four columns of Tab. 1, the mAP for the motorcycle and rider detections by the proposed curriculum learning (CL) based YOLOv4 model (see Sec. 4.1) is over 8.9% higher than the (non-CL) YOLOv4 model as well as Saumya et al. [21].

We now discuss the results on identifying triple-riding violations shown in the last four columns of Tab. 1. The Euclidean distance-based rider-motorcycle association by Saumya et al. [21], and the IOU-based approaches in Fig. 4 (a) and (b) of Sec. 4.2, are not robust to false-positive riders from other motorcycles. Hence, all these approaches have precision and F-Score lower than 50 (rows 1-4 of Tab. 1). However, using the proposed trapezium-shaped instance boxes as shown in Fig. 4 (c) improves the F-score by 17.57% as shown in Tab. 1 (row 5). Using the CL-based base model on our proposed violation identification approach further improves the F-score by 12.39%, leading to the highest precision, recall, and F-score of 84.44%, 73.07% and 78.34%.

Helmet Violations: Prior works identify helmet violations in two ways; i) Rider-level violations, which includes counting riders without a helmet, ii) Instance-level violations for motorcycles with at least one of their riders without




Figure 6. Qualitative results in complex scenarios: (a) Triple-riding violation with rider occlusions, (b) Helmet violations with head occlusions, (c) Triple-riding violation in a low light scene, (d) Helmet violation in a low light scene (top-left: crop enhanced for visualization).

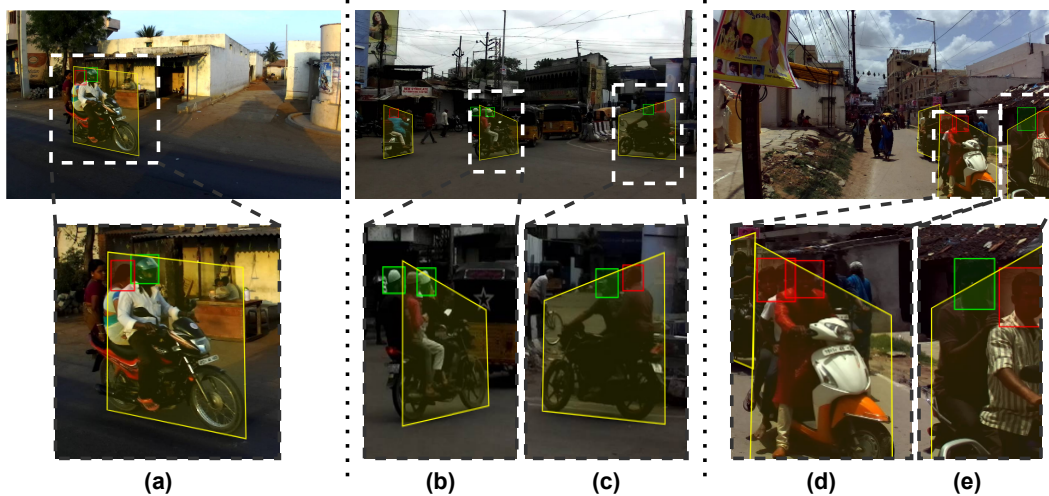


Figure 7. Failure cases. (a), (b), & (d): riders with non-distinguishable backgrounds leading to missed triple-riding violations, (b), (c) & (e): heads with caps and non-distinguishable backgrounds are sometimes falsely detected as helmets (green).

a helmet. We outperform previous works at both levels.

Unlike the rider and motorcycle boxes, the helmet and no-helmet boxes have comparatively lower IOU with each other. Thus, using a CL based detector to learn helmet and then no-helmet classes reduces the mAP from 90.00% to 83.5%, as shown in columns 1-4, rows 1-4 of Tab. 2. Therefore, as shown in rows 5 and 7, we propose to use the generic (non-CL) YOLOv4 to detect helmet violations.

As mentioned in Sec. 2, there are many ways to provide the ROIs to the violation detection model. The back-

ground context present in the ROI has a pronounced effect on the detection performance. As shown in the last four columns of Tab. 2, ROIs having complete *rider-motorcycle* instance (rows 5 and 7) provides enough background context when compared to rider-crops [20], upper body crops [4, 6, 11, 22, 23], or the full resolution images. We use ROIs obtained from predictions of the v4_CL model in Tab. 1. The proposed *rider-motorcycle* ROIs achieve the highest precision, recall, and F-score of 77.86%, 92.23% and 84.43%. For instance-level violations, our approach




Figure 8. Effect of using Amodal Regressor: Reduction in missing boxes corresponding to rider occlusions.

significantly improves the F-score by 30.83% compared to Chairat et al. [4] leading to precision, recall, and F-score of 99.01%, 95.23% and 97.08% as shown in the last two rows of Tab. 2.

Qualitative results are shown in Fig. 6, Fig. 7, Fig. 1 (c) and Fig. 3 (g). Fig. 6 shows violations identified by our model in complex scenarios at different scales and orientations. The scenarios include triple-riding and helmet violations with (a) rider occlusions and (b) head occlusions. Fig. 1 (c) and (d) also present qualitative results on both types of violations in the low light scenes. While analyzing the failure cases, we observe that our model sometimes fails in instances where (i) riders have non-distinguishable background as shown in Fig. 7 (a), (b) and (d), and (ii) heads with caps and non-distinguishable backgrounds, as shown in Fig. 7 (b) (c) and (e). We present more results in the demo video available on the github page.

5.2. Ablation Studies

Data Pre-processing: As discussed in Sec. 3, we pre-process the IDD subset we use for training our model to include occlusions. More precisely, we introduced an amodal regressor that predicts a bounding box for an occluded rider using annotated *helmet/no-helmet* box (for the rider) as input. As shown in Fig. 8, the percentage of missing rider boxes is significantly reduced by using the proposed methodology. It is important to note that i) the amodal regressor detects only one rider box for each helmet box, and ii) we keep the threshold for IOU between predicted and labeled rider boxes low and avoid all the false positives. We also analyze the effect of our pre-processing on the triple-riding violations task in Fig. 9. We use a generic YOLOv4 model and trapezium boxes to associate the motorcycle with its riders for the analysis. Correcting the rider boxes has led to an increase in precision. False negatives or missing riders due to occlusion are also reduced, leading to improved recall. Overall, as shown in Fig. 9, we obtain 7% and 10% gains in mAP and F-scores respectively due to the proposed data pre-processing techniques.

Curriculum Learning: We utilize Curriculum Learning (CL) for training our object detector as discussed in Sec. 4.1

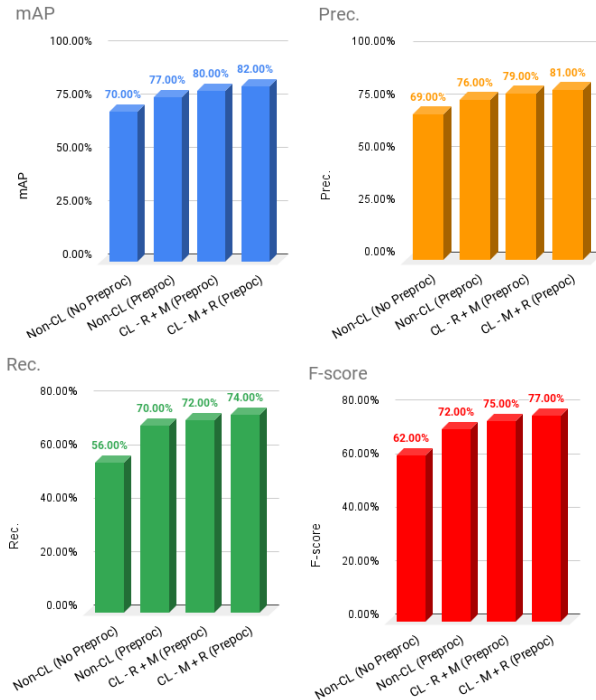


Figure 9. Detection Results: Effect of using proposed data pre-processing and curriculum learning.

and experiment with the class orders mentioned in the x-axis of Fig. 9. We find that training first on the motorcycle and then on the rider class (**M+R**) provides us the best results. A significant gain in recall compared to non-CL methods demonstrates the capability of the CL-based approach to robustly handle false negatives which arise due to overlapping predictions of the motorcycle and rider boxes.

6. Conclusion

This paper proposed a novel architecture capable of detecting, tracking, and counting triple-riding and helmet violations on crowded Asian streets using a camera mounted on a vehicle. Our approach jointly tackles both the violations. The proposed framework outperforms existing violation approaches because of its capability to handle rider occlusions, false negatives, and false-positive rides from other motorcycles in crowded scenarios. We demonstrate the effectiveness of using curriculum learning for motorcycle violations and trapezium bounding boxes instead of conventional rectangular boxes. Our work lays the foundation for utilizing such systems for increased road safety. It can be used for surveillance in any corner of the city without any considerable cost of a static camera network. In the future, we wish to explore the possibility of deploying our system to a distributed GPU setup for city-wide surveillance.

Acknowledgments: We thank iHubData, the Technology Innovation Hub (TIH) at IIIT-Hyderabad for supporting this project.

References

- [1] David M Allen. Mean Square Error of Prediction as a Criterion for Selecting Variables. *Technometrics*, 13(3):469–475, 1971. 5
- [2] World Bank. Traffic Crash Injuries and Disabilities: The Burden on Indian Society, 2021. 1
- [3] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal Speed and Accuracy of Object Detection. *arXiv preprint arXiv:2004.10934*, 2020. 5
- [4] Aphinya Chairat, Matthew Dailey, Somphop Limsoonthrakul, Mongkol Ekpanyapong, and Dharma Raj KC. Low Cost, High Performance Automatic Motorcycle Helmet Violation Detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3549–3557, March 2020. 2, 6, 7, 8
- [5] Corinna Cortes and Vladimir Vapnik. Support-Vector Networks. *Machine learning*, 20(3):273–297, 1995. 2
- [6] Kunal Dahiya, Dinesh Singh, and C. Krishna Mohan. Automatic Detection of Bike-riders without Helmet using Surveillance Videos in Real-time. In *International Joint Conference on Neural Networks (IJCNN)*, pages 3046–3051, 2016. 2, 7
- [7] Navneet Dalal and Bill Triggs. Histograms of Oriented Gradients for Human Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, volume 1, pages 886–893. Ieee, 2005. 2
- [8] Zhenhua Guo, Lei Zhang, and David Zhang. A Completed Modeling of Local Binary Pattern Operator for Texture Classification. *IEEE Transactions on Image Processing*, 19(6):1657–1663, 2010. 2
- [9] Guy Hacohen and Daphna Weinshall. On The Power of Curriculum Learning in Training Deep Networks. In *International Conference on Machine Learning*, pages 2535–2544. PMLR, 2019. 2
- [10] Wei Jia, Shiquan Xu, Zhen Liang, Yang Zhao, Hai Min, Shujie Li, and Ye Yu. Real-time Automatic Helmet Detection of Motorcyclists in Urban Traffic using Improved YOLOv5 Detector. *IET Image Processing*, 15(14):3623–3637, 2021. 2
- [11] Fahad Khan, Nitin Nagori, and Ameya Naik. Helmet Presence Detection on Motorcyclists Using Image Processing and Machine Vision Techniques. *International Engineering Journal For Research & Development*, 5(5):7–7, 2020. 2, 7
- [12] Shubham Kumar, Nisha Neware, Atul Jain, Debabrata Swain, and Puran Singh. Automatic Helmet Detection in Real-Time and Surveillance Video. In *Proceedings of ICM-LIP*, pages 51–60, 2020. 2
- [13] Hanhe Lin, Guangan Chen, and Felix Wilhelm Siebert. Positional Encoding: Improving Class-Imbalanced Motorcycle Helmet use Classification. In *IEEE International Conference on Image Processing (ICIP)*, pages 1194–1198. IEEE, 2021. 2, 5
- [14] David G Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 2
- [15] Nikhil Chakravarthy Mallela, Rohit Volety, RK Nadesh, et al. Detection of the triple riding and speed violation on two-wheelers using deep learning algorithms. *Multimedia Tools and Applications*, 80(6):8175–8187, 2021. 1, 2, 5
- [16] Hamidreza Mousavi, Maryam Imani, and Hassan Ghassemian. Deep Curriculum Learning for PolSAR Image Classification. *arXiv preprint arXiv:2112.13426*, 2021. 2
- [17] Robert Nürnberg. Calculating the area and centroid of a polygon in 2d. URL: <http://wwwf.imperial.ac.uk/rn/centroid.pdf>, 3, 2013. 5
- [18] World Health Organization et al. Global Status Report on Road Safety 2018: Summary. Technical report, World Health Organization, 2018. 1
- [19] Amjad Pervez, Jaeyoung Lee, and Helai Huang. Identifying Factors Contributing to the Motorcycle Crash Severity in Pakistan. *Journal of Advanced Transportation*, 2021:10, 2021. 1
- [20] Harish Rithish, Raghava Modhugu, Ranjith Reddy, Rohit Saluja, and C.V. Jawahar. Evaluating Computer Vision Techniques for Urban Mobility on Large-Scale, Unconstrained Roads. In *arXiv preprint arXiv:2109.05226*. IEEE, 2021. 2, 6, 7
- [21] Apoorva Saumya, V Gayathri, K Venkateswaran, Sarthak Kale, and N Sridhar. Machine Learning based Surveillance System for Detection of Bike Riders without Helmet and Triple Rides. In *International Conference on Smart Electronics and Communication (ICOSEC)*, pages 347–352. IEEE, 2020. 1, 2, 5, 6
- [22] Linu Shine and C. Jiji. Automated detection of helmet on motorcyclists from traffic surveillance videos: a comparative analysis using hand-crafted features and CNN. *Multimedia Tools and Applications*, 79, 05 2020. 2, 7
- [23] Dinesh Singh, C. Vishnu, and C. Krishna Mohan. Real-Time Detection of Motorcyclist without Helmet using Cascade of CNNs on Edge-device. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8, 2020. 2, 7
- [24] Deepak Kumar Singh, Shyam Nandan Rai, KJ Joseph, Rohit Saluja, Vineeth N Balasubramanian, Chetan Arora, Anbumani Subramanian, and CV Jawahar. ORDER: Open World Object Detection on Road Scenes. 2
- [25] Ministry Road Transport, Highway, et al. Road accidents in India. Technical report, Ministry of Road Transport and Highway, 2019. 1
- [26] Girish Varma, Anbumani Subramanian, Anoop Namboodiri, Manmohan Chandraker, and CV Jawahar. IDD: A Dataset for Exploring Problems of Autonomous Navigation in Unconstrained Environments. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1743–1751. IEEE, 2019. 2
- [27] C Vishnu, Dinesh Singh, C Krishna Mohan, and Sobhan Babu. Detection of Motorcyclists without Helmet in Videos using Convolutional Neural Network. In *International Joint Conference on Neural Networks (IJCNN)*, pages 3036–3041. IEEE, 2017. 2
- [28] Jiasi Wang, Xinggong Wang, and Wenyu Liu. Weakly-and Semi-supervised Faster R-CNN with Curriculum Learning. In *24th International Conference on Pattern Recognition (ICPR)*, pages 2416–2421. IEEE, 2018. 2

- [29] Yiru Wang, Weihao Gan, Jie Yang, Wei Wu, and Junjie Yan. Dynamic Curriculum Learning for Imbalanced Data Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5017–5026, 2019. [2](#)
- [30] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple Online and Realtime Tracking with a Deep Association Metric. In *IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE, 2017. [3](#), [6](#)
- [31] Lili Xiong, Yao Zhu, and Liping Li. Risk Factors for Motorcycle-related Severe Injuries in a Medium-sized City in China. *AIMS public health*, 3(4):907, 2016. [1](#)
- [32] B Yogameena, K Menaka, and S Saravana Perumaal. Deep learning-based helmet wear analysis of a motorcycle rider for intelligent surveillance system. *IET Intelligent Transport Systems*, 13(7):1190–1198, 2019. [2](#)