

ABSTRACT

Real-time object detection in dynamic outdoor environments is affected by variations in illumination, weather, occlusion, object scale, background clutter, and computational constraints. This study developed and evaluated adaptive, hybrid, and ensemble-based approaches for improving object-detection accuracy, efficiency, and adaptability under such conditions. A dataset comprising 1,000 daytime outdoor images and 6,420 annotated object instances was collected from Kampala's five divisions, Uganda. The dataset contained 11 object classes and represented morning, midday, and late-afternoon illumination; clear, overcast, and rainy weather; and urban and peri-urban environments. Images were preprocessed, augmented, annotated using bounding boxes, and divided into training, validation, and testing subsets in a 70:20:10 ratio. Baseline one-stage and two-stage detectors were evaluated using precision, recall, F1-score, mean Average Precision (mAP), inference latency, and frames per second (FPS). A hybrid YOLOv5s–Faster R-CNN model was subsequently developed by using YOLOv5s for rapid proposal generation and Faster R-CNN components for region-level classification and bounding-box refinement. Kullback–Leibler (KL), Jensen–Shannon, Hellinger, and squared Euclidean divergence functions were compared for probability-distribution alignment. A separate YOLOv11m continuous-learning framework combined experience replay with KL-regularised incremental updating. The initial hybrid model achieved an mAP@0.50 of 0.89, precision of 0.93, recall of 0.90, and 45 FPS. Following KL-regularised fine-tuning, its mAP@0.50 increased to 0.93, precision to 0.96, and recall to 0.94, while maintaining 43 FPS. KL divergence provided the strongest detection performance and convergence behaviour among the evaluated divergence functions. Continuous learning increased YOLOv11m's mAP@0.50 from 0.83 to 0.87 and recall from 0.77 to 0.83, while maintaining real-time inference at 65 FPS. The findings demonstrate that selective hybrid refinement improves localisation and detection accuracy, while replay-supported continuous learning provides a favourable balance between incremental adaptability and inference efficiency. Further validation using larger multi-location datasets, extended update sequences, and target edge hardware is required before operational deployment.

Keywords: *Adaptive learning, continuous learning, hybrid object detection, Kullback–Leibler divergence, real-time object detection, YOLO.*