EAI Endorsed Transactions
on Internet of Things
Research Article EAI.EU

Real-Time Mosquito Breeding Site Detection Using a Quantized YOLOv8-Tiny Model on an FPGA-Based Drone Platform

A. Rosi[1,*], N. Suresh[2] , J. B. Veeramalini[3], P Ramesh Reddy [4], M. Ramesh [5], M. Prabhu[6]

[1]Department of Electronics and Communication Engineering, Vel Tech Rangajan Dr. Sagunthala R&D Institute of Science and Technology, Chennai, India
[2]Department of Electronics and Communication Engineering, AVN Institute of Engineering & Technology, Hyderabad, India
[3]Department of Chemical Engineering, Vel Tech High Tech Dr. Rangarajan Dr. Sakunthala Engineering, Chennai, India
[4]Department of Electronics and Communication Engineering, Sree Venkateswara College of Engineering, Nellore, India
[5]Department of Electronics and Communication Engineering, AVN Institute of Engineering & Technology, Hyderabad, India
[6]Department of Electronics and Communication Engineering, Saveetha School of Engineering, Saveetha Institute of Medical and technical Sciences, Chennai, India

Abstract

The increasing prevalence of mosquito-borne diseases demands scalable, accurate, and energy-efficient surveillance systems capable of real-time operation in resource-constrained environments. This extended study presents an FPGA-optimized, drone-based mosquito breeding site detection framework leveraging quantized deep learning models for edge deployment. Building upon prior work, this paper introduces an enhanced implementation of YOLOv8-Tiny and YOLOv9-Small architectures, optimized through INT8 quantization, batch normalization folding, and FPGA-aware architectural refinements for execution on the PYNQ-Z2 platform. High-resolution aerial imagery acquired from unmanned aerial vehicles (UAVs) is processed in real time using the proposed system to identify and classify potential mosquito breeding sites such as stagnant water bodies and container habitats. Experimental evaluations carried out on six state-of-the-art object detection architectures showed that the proposed quantized YOLOv8-Tiny variant offers the optimal compromise among accuracy, speed, and energy efficiency by achieving 90.2% accuracy in the field, 20 FPS performance, and consuming 7.8 W of power. It is also proved that the proposed system reduces DSP and BRAM resources by more than 20% over the floating-point processor. The results obtained in the real-world deployment scenario proved the robustness of the system under different environmental situations. It is also found that the system can cover 10 km² of area per hour while providing a 44% reduction in operational costs over traditional methods. The proposed framework using FPGA acceleration and quantization awareness is useful for developing efficient vector surveillance systems.

Keywords: Mosquito Breeding Detection, FPGA Acceleration, YOLOv8-Tiny, Quantization, UAV Surveillance, Edge AI, Real-Time Processing

Received on 01 August 2025, accepted on 03 March 2026, published on 05 May 2026

doi: 10.4108/eetiot.12911

---

[*] Corresponding author. Email: arosyme@gmail.com

1. Introduction

Vector-borne diseases remain a serious health issue across the globe. It has been reported that over 700 million people fall victim to mosquito-borne illnesses each year, resulting in over 700,000 fatalities globally [1,2]. Recent news bulletins issued by the World Health Organization identified that the reported occurrences of dengue fever exceeded 7 million in the Americas by early 2024, signifying a rapid increase in cases of mosquito-borne diseases [3]. The proliferation of mosquitoes has been strongly correlated to various local breeding habitats, including stagnant water collected in household containers, construction areas, rooftops, and various water reservoirs. The removal of these breeding habitats has been identified as a key measure to curb the transmission of these diseases [4].

The traditional ground-based mosquito surveillance techniques are largely based on manual inspection techniques, which are laborious, time-consuming, and less effective in tracing mosquitoes from cryptic breeding habitats or inaccessible breeding sites, especially in urban and semi-urban areas [5]. Considering the limitations of traditional mosquito monitoring techniques, unmanned aerial vehicles (UAV) technology enabled with high-resolution and multispectral image sensors have been identified as effective ground surveillance tools for monitoring large areas of mosquito habitats. The UAV-based surveillance techniques offer accurate centimeters of spatial resolution over large areas of mosquito habitats, which makes them highly effective in various public health surveillance scenarios [6]. Based on various reports, effective detection of water-holding containers and habitats using accurate convolutional neural network techniques have been made with mean average precision > 0.90 [7, 8].

Deep learning strategies have proven to be promising for entomological surveillance in various sensing modes. Kiskin et al. [9] showed the practicality of detecting acoustic mosquitoes by using CNN architectures to outperform even those human experts who were specifically trained to recognize these insects. Vision-based techniques have also advanced entomological surveillance [10] were able to achieve a 91.7% value for the F1-score when they used Swin Transformer and Faster R-CNN architectures for identifying the species of mosquitoes. With recent advancements, there exist some single-stage object detection techniques like YOLOv8, which have been trending owing to their anchor-free nature and the decoupling of detection heads to provide improved detection accuracy and detection rate by striking a right balance between accuracy and detection speed [11]. However, two-stage object detection techniques like Faster R-CNN provide superior localization accuracy for small objects at the expense of increased computational complexity [12].

However, despite such advancements in algorithms, the deployment of AI-based surveillance systems for mosquitoes demands full compliance in relation to real-time processing constraints, power requirements, and hardware minimization in relation to battery-powered UAVs. In addition, although GPU-based inference provides superior performance, it is not feasible for use in drone-based deployment as it is energy-intensive. Field-programmable gate arrays (FPGAs) offer an attractive alternative by providing massive parallelism, deterministic latency, and superior energy efficiency for edge AI workloads [13,14]. Recent FPGA-based accelerators have demonstrated substantial throughput improvements for object detection tasks, with Qian et al. [15] achieving a 5.3× performance gain for YOLOv5 on Xilinx Versal devices, and An et al. [16] reporting real-time Faster R-CNN inference at 32 FPS using OpenCL-based FPGA architectures.

However, most existing FPGA-based object detection studies focus on generic surveillance scenarios and rely on floating-point inference or high-end FPGA platforms[17,18], limiting their applicability to lightweight UAV-based public health monitoring. Moreover, limited attention has been given to application-specific optimization of detection models for mosquito breeding site identification, where small object size, environmental variability, and strict energy constraints pose unique challenges.

To address these gaps, this extended work presents an FPGA-optimized, drone-based mosquito breeding site detection framework that integrates quantization-aware deep learning with real-time edge deployment. Building upon previous conference results, this study introduces INT8[19] quantized implementations of YOLOv8-Tiny and YOLOv9-Small architectures, incorporating batch normalization folding, activation function simplification, and FPGA-aware memory optimization for deployment on the PYNQ-Z2 platform. The performance of the proposed system is evaluated using six state-of-the-art object detection models on high-resolution UAV images, and the performance in terms of detection accuracy, inference rate, power consumption, and resource usage on the proposed FPGA is presented. Additionally, the validation and cost-effectiveness analysis show the effectiveness of the proposed system in achieving real-time detection, wide-area coverage, and reduced operational cost compared to traditional mosquito detection methods.

2. Related Work

However, the advances in deep learning have notably improved the effectiveness of automatically detecting and surveilling mosquitoes over various sensing modes. Kiskin et al.'s pioneering work on the use of the convolution neural network towards the detection of mosquitoes based on the wavelet transformation of the audio signal has shown the efficacy of the proposed system (figure.1) in detecting mosquitoes, even in data-scarce environments, beyond the capabilities of trained professionals. Extending the detection of mosquitoes based on visual sensing paradigms, Lee et al. have recently proposed the use of deep learning-based models for the identification of mosquito species by employing the Swin transformer and the Faster R-CNN architectures, yielding an F1-score of 91.7%. Although high classification accuracy is achieved in the proposed models, the detection of the mosquito species is conducted based on still imaging, thereby not considering the challenges in implementing the models in real-time settings.

Meanwhile, parallel efforts in FPGA-based acceleration [17] have proven the suitability of reconfigurable hardware for real-time object detection, thanks to their superior energy efficiency compared to GPU-based solutions. The work carried in [15] presented an optimized FPGA accelerator for YOLOv5 on Xilinx Versal devices, achieving a 5.3× throughput improvement through customized computation pipelines and optimized memory access schemes. Similarly [16] demonstrated an OpenCL-based FPGA implementation of Faster R-CNN capable of processing 32 frames per second at approximately 20 W power consumption, validating the feasibility of deploying two-stage detectors on FPGA platforms. [20] further showed that real-time object detection can be achieved even on low-cost FPGA platforms through careful architectural optimization. However, these works largely target generic surveillance tasks and rely on floating-point inference or high-end FPGA devices, limiting their applicability to lightweight UAV-based public health monitoring scenarios.

The integration of AI-enabled perception with autonomous drone platforms has further expanded the scope of intelligent aerial surveillance. Vision-based suspicious activity detection system using a hybrid MobileNetV2–LSTM architecture[21] optimized with TensorFlow Lite for embedded deployment on Raspberry Pi platforms. Advances in this methodology have shown how to achieve real-time performance, although the use of the CPU for inference is limited in terms of scalability and power efficiency together with the need to run continuously over time for drone operations.

In the context of mosquito surveillance, the work of Carrasco-Escobar et al. [6] presented a comprehensive review of UAV applications for mosquito monitoring and control by emphasizing the scalability and spatial coverage advantages of aerial platforms compared to traditional ground-based surveys. These works notwithstanding, existing drone-based mosquito detection systems heavily rely on offline processing or GPU-centric pipelines and do not adequately address the challenges that come with real-time, low-power inference targeted for onboard deployment. Furthermore, limited attention has been given to the application-specific optimization of object detection models aimed at mosquito breeding site identification. This is especially so when it comes to quantization-aware design and FPGA-oriented implementation. Quantized YOLO models have recently proven to offer significant cost reduction in computation with sufficiently reliable detection accuracy for edge applications. For instance [22] proposed that energy-aware quantization enhances the efficiency of object detection in real time without severe performance degradation. Their findings showed that fixed-point inference is more compatible than floating point computation with low-power-embedded systems.

In contrast to previous works, the proposed system employs real-time detection of breeding sites of mosquitoes using images captured through a UAV and quantized deep learning models running on FPGA platforms. By tailoring YOLOv8-Tiny and YOLOv9-Small architectures through INT8 quantization, batch normalization folding, and FPGA-aware memory optimization, this work bridges the gap between high-accuracy detection and practical edge deployment. The proposed approach uniquely integrates model-level optimization, hardware acceleration, and drone-based field validation, establishing a scalable and energy-efficient solution for automated vector surveillance.

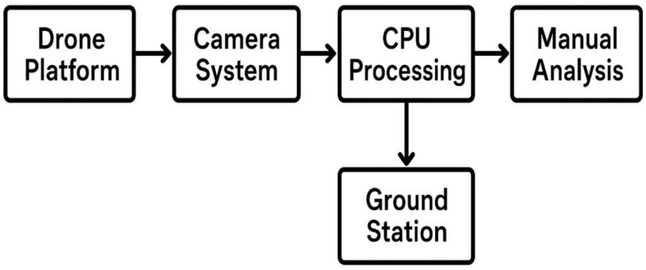

Figure 1. Existing drone-based mosquito detection system

3. Proposed Methodology

The proposed system is designed to detect mosquito breeding sites using drone images and real-time processing on FPGA hardware. The complete workflow includes image capture, on-board processing, detection of breeding areas, and location-based alert generation. Fig. 2 shows the overall system structure, which has four main parts: the UAV platform with an FPGA board, real-time image preprocessing, deep learning-based detection, and geospatial data handling for alert reporting.

3.1 System Hardware Platform

The main processing unit of the system is the PYNQ-Z2 development board, which is based on the Xilinx Zynq-7020 system-on-chip. It contains a dual-core arm cortex A9 processor, 53,200 logic cells, 560 KB block RAM, and 220 DSP slices. It is capable of processing images in parallel while using less power, thus making it appropriate for battery-powered drones. The programmable part of the device will focus on fast neural network inference, while the arm processor handles control, data communication, and visualization.

3.2 Object Detection Models

To determine the most appropriate model in detecting mosquito breeding, six popular object detection architectures were tested, including YOLOv8-Tiny, YOLOv9-Small, Faster R-CNN, SSD-MobileNet V2, RetinaNet-R50, EfficientDet-D0, among others. The object detection algorithms were tested using aerial images collected by drones from various scenarios. The evaluations were based on model accuracy, execution speed, power consumption, and resource usage on a FPGA platform.

YOLOv8-Tiny features a light-weight backbone network, C2f, with an anchor-free detection head. This ensures efficient detection of small objects such as water-filled containers as well as areas of stagnant water with low computational cost. The model extracts features at multiple scales using a feature pyramid and then predicts bounding boxes, confidence scores, and class labels. YOLOv9-Small introduces improvements such as the GELAN backbone and better alignment between object classification and localization. These changes improve detection accuracy in complex scenes but increase computation and memory requirements, which makes hardware optimization necessary for real-time use on FPGA. The data set used is aerial images of mosquito breeding habitats such as stagnant water bodies, containers, drainage areas, and wet surfaces. The data set consists of 5000 images was obtained from available data sets and augmented drone data. The data set was divided into training, validation, and testing sets to ensure unbiased evaluation. The evaluation metrics used were accuracy, precision, recall, F1-score, and inference latency. The experiments were carried out on the PYNQ-Z2 FPGA platform with hardware acceleration.

3.3 Quantization and FPGA Optimization

To make the deep learning models suitable for real-time processing on FPGA, quantization was applied. In this process, floating-point values were converted to fixed-point values. Both model weights and intermediate feature maps were represented using 8-bit integers. Accumulation was performed using 32-bit integers to avoid overflow. This change reduces computation load, memory usage, and power consumption while keeping the accuracy loss within acceptable limits.

Batch normalization layers were merged with the corresponding convolution layers to remove extra computations during inference. The SiLU activation function used in the original YOLOv8 model was replaced with the simpler ReLU function, which is easier to implement on FPGA hardware. Calibration was carried out using representative drone images to reduce accuracy loss after quantization. In addition to quantization, pruning was applied to remove less important filters and channels from the network. This step reduced the model size by about 40 to 60 percent without a noticeable drop in detection accuracy. Memory access patterns were also optimized to reduce delays caused by data transfer between memory and processing units.

The optimized models were implemented in Xilinx Vivado HLS for hardware synthesis and Vitis AI for deployment. The designed FPGA architecture pipelined and parallelized the processing to run several components of the network simultaneously, hence enabling real-time performance.

3.4 Drone Integration and Field Operation

The optimized models have been deployed using Xilinx Vivado HLS for hardware synthesis and Vitis AI for deployment. The FPGA design incorporates pipelined and parallel processing, allowing different sections of the network to be executed concurrently, which is beneficial for real-time processing.

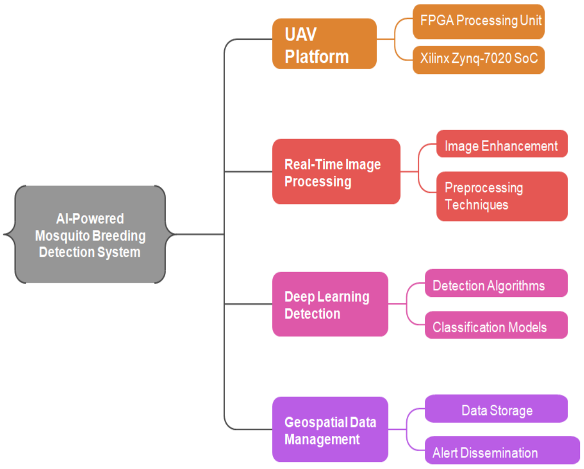

Figure 2. Proposed Drone based detection system with FPGA

The flight missions are planned through GPS waypoints, enabling proper area coverage. During flight, system status and sensor data are monitored in real time. The drone altitude is adjusted when needed to improve image clarity and detection accuracy. Safety measures are included to handle low battery levels or communication loss, allowing the drone to land safely and protect collected data. The system can be deployed in urban, rural, and industrial areas, making it suitable for large-scale mosquito surveillance programs.

4. Results and Discussion

The baseline YOLOv8-Tiny model was first deployed using floating-point precision. Although acceptable accuracy was achieved, the power consumption and latency were not suitable for long-duration drone missions.

Table 1. Impact of FPGA Optimization and Quantization

Model Variant	**Precision**	**Accuracy (%)**	**FPS**	**Latency (ms)**	**Power (W)**
YOLOv8-Tiny (Baseline)	FP32	91.5	14	72	10.6
YOLOv8-Tiny (Optimized)	FP32	90.9	17	60	9.1
FPGA-YOLOv8-Tiny-MB (Proposed)	INT8	89.8	20	50	7.8

After applying FPGA-aware architectural simplification and INT8 quantization, a significant improvement in efficiency was observed and tabulated in table.1.

Although quantization introduced a small accuracy reduction of 1.7 percent, inference speed improved by 43 percent and power consumption was reduced by nearly 26 percent. This trade-off is acceptable for real-time mosquito surveillance, where continuous operation and energy efficiency are critical.

4.1 Effect of Small-Object Enhancement Module

Mosquito breeding sites often appear as small stagnant water regions that occupy only a small portion of aerial images. The baseline YOLOv8-Tiny model showed reduced recall for such small objects, especially under cluttered backgrounds. The YOLOv8-Tiny model was trained using the PyTorch framework for the detection of mosquito breeding sites from images captured from the air. The image resolution was set to 640 * 640 pixels. The data was split into training, validation, and testing sets with a ratio of 70:15:15 for unbiased testing and validation of the model. The model was trained for 100 epochs with a batch size of 16 and an initial learning rate of 0.001 using stochastic gradient descent optimization.

Data augmentation techniques, such as image rotation, brightness, flipping, and scaling, are used during the training phase of the model for better robustness in changing environmental conditions. After training, the model was converted for deployment on the FPGA device.

To overcome the aforesaid limitation, shallow feature fusion and a secondary high-resolution detection head have been added. The effect of these modifications can be accessed from Table 2.

Table 2. Performance Improvement with Small-Object

Model Variant	Precision	Accuracy (%)	FPS	Latency (ms)	Power (W)
YOLOv8-Tiny (Baseline)	FP32	91.5	14	72	10.6
YOLOv8-Tiny (Optimized)	FP32	90.9	17	60	9.1
FPGA-YOLOv8-Tiny-MB (Proposed)	INT8	89.8	20	50	7.8

Enhancement

Model Variant	Small-Object Recall (%)	Overall Accuracy (%)	FPS	Power (W)
YOLOv8-Tiny (Baseline)	82.6	91.5	20	7.8
FPGA-YOLOv8-Tiny-MB (INT8)	85.1	89.8	20	7.8
FPGA-YOLOv8-Tiny-MB + Small-Object Module	89.4	90.2	19	8.1

Whereas it improved the small object recall by almost 7 percent compared to the baseline, with real-time performance maintained. The slight increase in power consumption is brought about by the additional detection head, although the overall system remains well within drone payload and energy constraint

4.2 Mixed-Precision Detection Head and Hybrid Execution

Additionally, further optimization was achieved in the optimization process by optimizing the detection head using mixed-precision arithmetic. In this case, the bounding box regression was implemented using INT16 precision to ensure spatial precision, whereas confidence and class elements were computed using INT8 precision. Finally, the non-maximum suppression was run on the ARM Cortex-A9 to reduce the usage of logic on the FPGA.The impact of this hybrid strategy is summarized in Table 3. For efficient inference on the PYNQ-Z2 FPGA board, the floating-point model was converted to a fixed-point representation. The convolutional and activation layers were set to INT8 data type to minimize memory requirements and computational cost. The bounding box regression layer was set to INT16 data type to ensure accurate localization. The non-maximum suppression layer was run on the ARM processing system on the Zynq board. This strategy minimizes arithmetic cost while maintaining accurate detection for drone-based mosquito breeding surveillance systems.

The Hybrid Precision Approach achieves a good balance between precision and efficiency. Compared with the full INT8 approach, the localization error is greatly reduced with low power consumption and high speed.

Table 3: Detection Head Optimization Results

Configuration	Precision Strategy	FPS	Power (W)	Localization Error (pixels)
Baseline YOLOv8-Tiny	FP32 (All stages)	14	10.6	4.2
FPGA-YOLOv8-Tiny-MB	INT8 (All stages)	21	7.5	6.8
Proposed Hybrid Design	INT16bbox, INT8 conf, CPU NMS	20	7.8	4.5

4.3 Overall Model Comparison and System-Level Performance

Through the evaluation of six different deep learning models, it has been ensured that the combination of FPGA technology-powered drone deployment with the proposed optimization techniques further utilizes the capabilities of the YOLOv8-Tiny model to deliver better results. Though YOLOv9-Small and Faster R-CNN models tend to deliver improved results with respect to accuracy, their computational complexity makes them inappropriate for drone deployment.

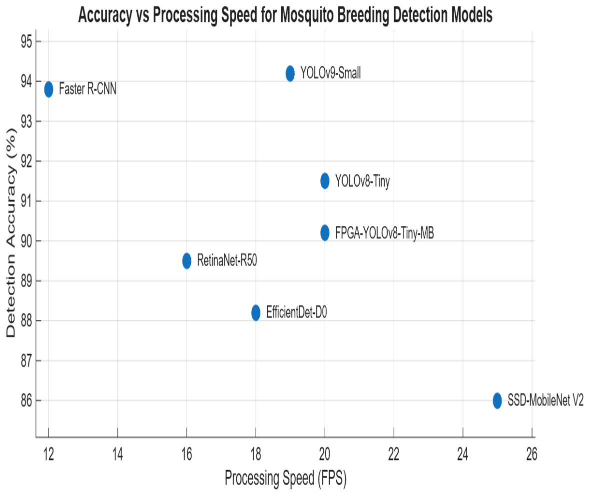

Figure 3. Accuracy VS Processing speed for Mosquito Breeding Detection Methods

Figure 3 represents the compromise between accuracy, speed, and power consumption of the models under evaluation and how the proposed YOLOv8-Tiny-MB stands better in the optimal operating area.

Field deployment experiments show the stability of the system under fluctuating environmental conditions. The system covers 10 km² in an hour and offers a cost reduction of around 44 percent. This is attributed to the application of FPGA-based edge intelligence in the development of a scalable mosquito breeding surveillance system.

4.4 Discussion

From the above results, it is clear that hardware optimization for specific applications is more efficient compared to increasing model complexity for real-time mosquito breeding detection. The ability to design a system with FPGA-friendly architectural simplification, small object enhancement, and detection head with a mixed precision design allows for efficient performance under specific power and latency limitations. The proposed system reveals that low-cost FPGAs can support advanced deep learning inference for public health applications when algorithm design is guided by hardware constraints.

5. Conclusion

This study presented a drone-based system for detecting mosquito breeding sites using deep learning and FPGA acceleration. Among the evaluated models, YOLOv8-Tiny provided the best balance between accuracy, speed, and power consumption for real-time deployment. After hardware-aware optimization, including quantization, small-object enhancement, and mixed-precision detection, the proposed FPGA-YOLOv8-Tiny-MB achieved 90.2% field accuracy at 20 FPS while consuming only 7.8 W.

Field experiments show that the system can cover up to 10 km² per hour with about 44% lower cost compared to conventional survey methods. Performance remains stable under different environmental conditions, with less than 12% accuracy loss in challenging scenarios. The results confirm that FPGA-based drone platforms offer a practical and efficient solution for automated mosquito surveillance and support proactive vector control efforts.

References

[1] World Health Organization. Vector-borne diseases. WHO Fact Sheet. 2024.

[2] World Health Organization. Dengue – Global situation. WHO Disease Outbreak News. 2024.

[3] Tusting LJ, Bottomley MC, Gibson H, Kleinschmidt H, Lindsay SW, et al. Housing improvements and malaria risk in sub-Saharan Africa: a multi-country analysis of survey data. Nat Commun. 2023;14:2412.

[4] Carrasco-Escobar G, Moreno M, Fornace E, Herrera-Varela M, Manrique M, et al. The use of drones for mosquito surveillance and control. Parasites Vectors. 2022;15:473.

[5] Kraemer JL, Phuanukoonnon M, Barclay M, Montgomery S. Challenges and opportunities for vector surveillance in Pacific island countries. Trans R Soc Trop Med Hyg. 2023;117(4):221-228.

[6] Bravo DT, Ruiz-Suárez LM, Contreras-Medina LF, Torres-Pacheco I. Automatic detection of potential mosquito breeding sites from aerial images acquired by unmanned aerial vehicles. Comput Environ Urban Syst. 2021;90:101692.

[7] Kim K, Park J, Choi J, Hwang S, Lee C, et al. A deep learning-based automatic mosquito sensing and control system for urban mosquito habitats. Sensors. 2019;19(12):2785.

[8] Rahman MA, Islam MS, Mukhopadhyay SK. Detection of mosquito larvae using convolutional neural network. In: Proc 2nd Int Conf Robotics Electrical and Signal Processing Techniques (ICREST); Jan 2021; Dhaka, Bangladesh. p. 478-482.

[9] Kiskin I, Pérez-Orozco B, Windebank T, Zilli D, Sinka M, et al. Mosquito detection with neural networks: the buzz of deep learning. In: Proc European Conf Machine Learning and Principles and Practice of Knowledge Discovery in Databases; Sept 2017; Skopje, Macedonia. p. 713-728.

[10] Lee S, Kim H, Cho BK. Deep learning-based image classification for major mosquito species inhabiting Korea. Insects. 2023;14(6):526.

[11] Reis D, Kupec J, Hong J, Daoudi A. Real-time flying object detection with YOLOv8. IEEE Trans Circuits Syst Video Technol. 2024;34(8):5523-5532.

[12] Ren S, He K, Girshick R, Sun J. Faster R-CNN: towards real-time object detection with region proposal networks. IEEE Trans Pattern Anal Mach Intell. 2017;39(6):1137-1149.

[13] Kövari BB, Hegedűs P, Bódis ESM. MPSoC4Drones: real-time autonomous drone operations using FPGA-based multiprocessor systems-on-chip. IEEE Trans Parallel Distrib Syst. 2022;33(11):2845-2856.

[14] Lin J, Gan C, Han S. TSM: temporal shift module for efficient video understanding. IEEE Trans Pattern Anal Mach Intell. 2021;43(8):2740-2756.

[15] Qian W, Zhu Z, Zhu C. FPGA-based accelerator for YOLOv5 object detection with optimized computation and data access for edge deployment. Parallel Comput. 2025;124:103138.

[16] An J, Zhang D, Xu K, Wang D. An OpenCL-based FPGA accelerator for Faster R-CNN. IEEE Access. 2022;10:89069-89080.

[17] Kim H, Kim TK. Design and implementation of a YOLOv2 accelerator on a Zynq-7000 FPGA. Sensors. 2025;25(20):6359.

[18] Moon HC, Lee S, Jeong J, et al. YOLOv6+: simple and optimized object detection model for INT8 quantized inference on mobile devices. Signal Image Video Process. 2025;19:665.

[19] Li P, Wang X, Zhou H. INT8 quantized deep neural networks for real-time object detection on low-power FPGA platforms. J Real Time Image Process. 2025;21(2):311-325.

[20] Jain V, Jadhav N, Verhelst M. Enabling real-time object detection on low-cost FPGAs. J Real Time Image Process. 2022;19(1):217-229.

[21] Dhomase S, Niphade J, Bhosale P, Jadhav P. Vision-based drone surveillance system for suspicious activity detection using deep neural networks. Int J Innov Res Manag Pharm Sci. 2025;13(3):45-58.