International Journal of Creative and Open Research in Engineering and Management
ISSN: **3108-1754 (Online)**
Volume 02 Issue 05 May-2026 | Impact Factor: 3.5

Design and Implementation of An Autonomous AI-Based Wild Animal Detection and Repellent System for Crop Protection

Mrs. A. Asrin Mahmootha[1], Dr. B. Aysha Banu[2], Abdul Ahad L[3], Hari Krishnan M[4], Kaleel Rahman S[5], Mohamed Mubarak S[6]

[1]Assistant Professor, Department of Information Technology, Mohamed Sathak Engineering College, Kilakarai, Tamil Nadu, India
[2] Professor & Head, Department of Information Technology, Mohamed Sathak Engineering College, Kilakarai, Tamil Nadu, India
[3] B.Tech Student, Department of Information Technology, Mohamed Sathak Engineering College, Kilakarai, Tamil Nadu, India
[4] B.Tech Student, Department of Information Technology, Mohamed Sathak Engineering College, Kilakarai, Tamil Nadu, India
[5] B.Tech Student, Department of Information Technology, Mohamed Sathak Engineering College, Kilakarai, Tamil Nadu, India
[6] B.Tech Student, Department of Information Technology, Mohamed Sathak Engineering College, Kilakarai, Tamil Nadu, India asrinmahmootha@gmail.com

How to Cite this Article:

S, M. M., S, K. R., M, H. K., L, A. A. & Mahmootha, A. A. (2026). Design and Implementation of An Autonomous AI-Based Wild Animal Detection and Repellent System for Crop Protection. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(05). https://doi.org/10.55041/ijcope.v2i5.291

https://doi.org/10.55041/ijcope.v2i5.291

Abstract—

Human-wildlife conflict (HWC) causes an estimated 40% crop loss in border farming communities across South Asia, with annual damages exceeding ₹12,000 crore in India alone. Traditional deterrent methods including electric fencing, manual vigilance, and passive infrared (PIR) sensors suffer from high false-positive rates (up to 28%), lack of species specificity, and prohibitive maintenance costs. This paper presents an autonomous, solar-powered, edge-AI system that integrates YOLOv8 object detection, adaptive multi-modal repellent actuation, and real-time GSM-based remote monitoring for precision crop protection. The system employs a Raspberry Pi 4B edge processor running a quantized and pruned YOLOv8n model (12 MB, 0.93 mAP@0.5) that detects elephants, wild boars, and other crop-raiding species at 280 ms latency. Upon detection, species-specific repellent protocols are activated: infrasonic deterrence (10–20 Hz) for elephants, ultrasonic emission (18–25 kHz) for boars, and stroboscopic lights for nocturnal intrusions. Field trials conducted over a 10-acre farm in Tamil Nadu demonstrated a 92% animal repulsion rate, 99.7% system uptime, and a 4.2% false-positive rate—outperforming all baseline methods. A five-year cost-benefit analysis confirms a 340% return on investment versus traditional electric fencing. The proposed system demonstrates that edge AI can deliver sustainable, farmer-affordable, and ecologically non-lethal wildlife management at scale.

Keywords: YOLOv8, Edge AI, Human-Wildlife Conflict, Crop Protection, Solar-Powered IoT, Precision Agriculture, Wild Animal Deterrent

1. INTRODUCTION

Agriculture remains the economic backbone of rural South Asia, employing over 58% of the Indian workforce and contributing approximately 18% of national GDP. However, the productive coexistence of farming communities with encroaching wildlife has deteriorated sharply over the past two decades. Rapid deforestation, shrinking forest corridors, and climate-induced habitat disruption have displaced large mammal populations into agricultural peripheries, precipitating a crisis of human-wildlife conflict (HWC) of unprecedented scale.

In Tamil Nadu alone, the Forest Department documented 3,742 crop-raiding incidents in 2022–23, with elephants (Elephas maximus indicus), wild boars (Sus scrofa), and nilgai (Boselaphus tragocamelus) responsible for over 78% of reported losses [1]. National estimates place annual crop damage attributable to wildlife at ₹12,000–15,000 crore, while indirect costs—medical treatment of attack victims, psychological distress, and labour displacement—add another 30–40% to the burden [2]. The Food and Agriculture Organization (FAO) reports that smallholder farmers facing repeated HWC are 2.3 times more likely to abandon land, accelerating rural poverty and reducing regional food security [3].

Existing countermeasures fall into three broad categories: physical barriers, chemical/acoustic deterrents, and human vigilance. Electric fences, the most widely deployed physical barrier, achieve only 65% containment effectiveness due to maintenance failures, voltage drops during monsoon seasons, and the learning behaviour of elephants that test fence integrity [4]. Chemical repellents such as chilli-smoke generators require constant replenishment and offer no real-time intelligence. Manual night guarding—still practised by 62% of border farmers—exposes community members to physical danger and is unsustainable given rural labour shortages [5]. Passive infrared (PIR)-based alarm systems suffer from 28% false-positive rates triggered by vegetation movement and thermal gradients, causing alert fatigue.

The convergence of three enabling technologies now makes a superior solution feasible: (i) lightweight deep neural networks capable of real-time object detection on resource-constrained hardware, (ii) affordable solar-powered edge computing platforms, and (iii) species-specific acoustic deterrence validated by animal behaviour research. YOLOv8 [6], the latest evolution of the You Only Look Once detection paradigm, achieves state-of-the-art accuracy-latency tradeoffs on embedded processors. Simultaneously, the declining cost of single-board computers (SBCs) and long-range GSM modules enables deployment in remote, grid-independent agricultural settings.

1.1 Research Contributions

This paper makes the following primary contributions:

- A complete end-to-end autonomous wildlife detection and repulsion system architecture optimized for edge deployment on solar power.
- A custom annotated dataset of 20,000 field images (60,000 after augmentation) covering elephants, wild boars, and other species under diverse lighting, weather, and occlusion conditions.
- A YOLOv8n model quantized and pruned to 12 MB achieving 0.93 mAP@0.5 at 280 ms inference on Raspberry Pi 4B, a 3.1× size reduction over baseline with less than 1% accuracy loss.
- An adaptive multi-modal repellent controller that selects species-specific deterrent frequencies and validates efficacy, achieving 92% repulsion in field trials across 10 acres over six months.
- A rigorous five-year economic analysis demonstrating 340% ROI and a total system cost of ₹16,000 per installation—accessible to smallholder farmers.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 describes the overall system architecture. Section 4 details the methodology including dataset preparation, model training, and edge optimization. Section 5 covers hardware implementation. Section 6 presents results and discussion. Section 7 concludes with future directions.

2. RELATED WORK

Research on automated wildlife deterrence and agricultural surveillance spans three converging domains: computer vision-based monitoring, smart repellent actuation systems, and edge AI deployment for agricultural applications.

2.1 Vision-Based Wildlife Monitoring

Early camera-trap deployments for wildlife monitoring relied on motion-triggered still cameras with manual review, achieving practical identification rates only 60–70% of the time owing to poor image quality and human resource constraints [7]. Beery et al. [8] introduced Context R-CNN, a cloud-based architecture fusing temporal context across frames for wildlife identification, achieving 0.87 mAP but requiring sustained internet connectivity and cloud compute—prohibitive for rural deployment. Norouzzadeh et al. [9] demonstrated that deep CNNs trained on Snapshot Serengeti data could match expert-level species classification at scale, yet inference latency exceeded 500 ms per frame on GPU hardware, far exceeding real-time requirements.

More recent work has explored lightweight architectures. MobileNet-SSD variants deployed on NVIDIA Jetson Nano by Kholiavchenko et al. [10] achieved 0.81 mAP at 312 ms latency for large mammal detection in savannah environments. Parham et al. [11] benchmarked six YOLO variants for field deployment, finding YOLOv5s optimal for constrained hardware, though their study predates the YOLOv8 release. Our work extends this line by adopting YOLOv8n with post-training quantization and structured pruning, achieving superior accuracy-latency performance on a lower-cost platform.

2.2 Automated Repellent Systems

Wildlife deterrence technology has evolved from simple acoustic alarms to species-adaptive multi-modal systems. Hedwig et al. [12] validated that playback of bee swarm sounds at frequencies between 300–8,000 Hz caused statistically significant African elephant avoidance in 98% of trials, confirming acoustic sensitivity as a viable deterrence vector. Shivik et al. [13] conducted a systematic review of non-lethal carnivore deterrence, identifying the critical principle of stimulus unpredictability: animals rapidly habituate to fixed-frequency deterrents, necessitating randomized frequency selection—a design principle incorporated into our adaptive controller.

For wild boar deterrence, Mayer [14] established the ultrasonic sensitivity range at 18–40 kHz, with optimal aversive response at 20–25 kHz. Combined visual-acoustic deterrence (strobe + ultrasonic) produced 85% longer avoidance persistence versus acoustic-only systems in controlled trials [15]. Our system implements this multi-modal principle with species-specific protocol selection driven by real-time AI classification.

2.3 Edge AI for Precision Agriculture

The deployment of neural inference at the edge of agricultural networks has gained substantial momentum since 2021, driven by the need for real-time, low-latency decision making in remote farm environments where continuous cloud connectivity is unreliable. Kamilaris and Prenafeta-Boldú [16] reviewed more than 40 deep learning applications in agriculture and reported accuracy levels ranging from 89–99% for plant disease detection and crop classification. Despite this progress, their survey highlighted wildlife monitoring and crop-raiding prevention as a critical and underexplored research gap, particularly in smallholder farming scenarios where infrastructure and budget constraints are significant.

Koirala et al. [17] demonstrated the feasibility of edge-based inference by deploying YOLOv3 on the NVIDIA Jetson TX2 platform for mango detection, achieving 95% precision at 18 FPS in field conditions. While this work proved the practicality of real-time agricultural computer vision, the high cost of Jetson-class hardware—typically 4–10× more expensive than Raspberry Pi devices—limits large-scale adoption in developing regions. This cost barrier has slowed the translation of edge AI research into affordable solutions for small and marginal farmers.

Solar-powered IoT agricultural systems have been widely validated for environmental and soil monitoring applications [18], particularly for sensing temperature, humidity, and moisture in off-grid farms. However, these systems generally support low-power sensing tasks and rarely address the higher energy demands of continuous image capture, edge inference, and real-time wildlife deterrence. Active wildlife management requires sustained computation, reliable communication, and uninterrupted power availability—factors that remain insufficiently explored in prior research.

To address these limitations, our system introduces a solar-powered edge-AI architecture specifically designed for wildlife detection and crop protection. The proposed platform integrates a 40W photovoltaic array with a lithium iron phosphate (LiFePO4) battery system optimized for high cycle life, thermal stability, and deep-discharge tolerance. Field deployment over a six-month period demonstrated 99.7% system uptime, confirming the feasibility of sustained, autonomous operation in remote agricultural environments. This work therefore bridges the gap between low-power agricultural IoT sensing and high-compute edge AI applications for real-time wildlife deterrence.

Table 1: Comparative Analysis of Related Work

Study	Method	mAP	Latency	Power Source	Deployment
Beery et al. [8]	Context R-CNN	0.87	500 ms	Cloud GPU	Monitoring Only
Norouzzadeh et al. [9]	CNN Ensemble	0.91	520 ms	GPU Server	Monitoring Only
Kholiavchenko et al. [10]	MobileNet-SSD	0.81	312 ms	Jetson Nano	Monitoring Only
Parham et al. [11]	YOLOv5s	0.88	340 ms	Jetson TX2	Monitoring Only
Koirala et al. [17]	YOLOv3	0.95	55 ms	Jetson TX2	Crop Detection
Proposed System	YOLOv8n (Opt.)	0.93	280 ms	Solar/ LiFePO4	Active Repulsion

3. SYSTEM ARCHITECTURE

The proposed system is designed around a Sense–Think–Act (STA) paradigm, in which a multi-sensor perception layer continuously monitors the environment, an edge AI inference engine classifies detected entities, and an adaptive actuation module deploys species-specific deterrents. All processing occurs locally on-device, eliminating latency and connectivity dependence introduced by cloud architectures.

3.1 Hardware Architecture

The system integrates five hardware subsystems: (i) a Sony IMX219 wide-angle camera (8 MP, 160° FOV, IR cut filter) for daytime imaging, (ii) a 940 nm infrared array for passive night-time illumination, (iii) a Raspberry Pi 4B (4 GB LPDDR4) serving as the primary edge inference and control unit, (iv) a multi-modal deterrent assembly comprising a 25W infrasonic transducer (5–25 Hz), a 40 kHz piezoelectric ultrasonic emitter, and a 10W xenon strobe array, and (v) a SIM800L GSM module for remote telemetry and alert dispatch.

Power is supplied by a 40W monocrystalline photovoltaic panel with a PWM charge controller feeding a 50 Ah LiFePO4 battery pack (25.6V nominal). The system draws 6.8W during active inference and 1.2W in standby, yielding 72 hours of autonomous operation under zero-solar conditions. A waterproof ABS enclosure (IP67) with passive aluminium heat-sink maintains component temperatures below 65°C in ambient conditions up to 45°C.

Table 2: Hardware Specifications

Component	Specification	Purpose	Unit Cost (₹)
Raspberry Pi 4B	4 GB RAM, ARM Cortex-A72 @ 1.5 GHz	Edge inference & control	4,200
Sony IMX219 Camera	8 MP, 160° FOV, IR filter	Daytime capture	1,100
IR Illuminator Array	940 nm, 15 m range	Night-time capture	850
Infrasonic Transducer	25W, 5–25 Hz, 110 dB SPL	Elephant deterrence	2,400
Ultrasonic Emitter	40 kHz piezoelectric, 120 dB	Boar deterrence	1,200
Xenon Strobe	10W, 1–10 Hz adjustable	Nocturnal deterrence	950
SIM800L GSM Module	Quad-band 850/900/1800/1900 MHz	Remote telemetry	650
Solar Panel	40W monocrystalline	Primary power	1,800
LiFePO4 Battery	50 Ah, 25.6V, 1280 Wh	Energy storage	2,400
PWM Charge Controller	30A, 12/24V	Battery management	750
ABS Enclosure	IP67, 400×300×200 mm	Environmental protection	700
		Total System Cost	₹16,000

3.2 Software Stack

The software architecture is implemented in Python 3.11 with PyTorch 2.1 as the deep learning framework and OpenCV 4.8 for image pre-processing and camera interfacing. Inference is performed using the ONNX Runtime 1.16 execution engine with CPU optimization, selected over TensorFlow Lite for superior dynamic-shape support. The actuation control layer interfaces with GPIO peripherals through RPi.GPIO and drives the GSM module via the pyserial AT-command interface. A Flask 3.0 web server provides the farm-owner dashboard accessible over cellular data, with SQLite as the local telemetry database.

3.3 Sense–Think–Act Control Flow

The system operates in a continuous 500 ms polling loop. The perception pipeline captures and downsamples frames to 640×640 pixels, applies histogram equalization for low-light frames, and buffers the normalized tensor for inference. The YOLOv8n model produces bounding box predictions with class probabilities; detections with confidence above 0.65 are forwarded to the repellent controller. The actuation module selects the appropriate deterrent protocol, randomizes frequency within the species-effective band, activates the actuators for 30 seconds, and logs the event. A re-detection check at 60 seconds determines whether extended deterrence is required.

Algorithm 1: Adaptive Repellent Protocol Selection

```
INPUT: detection_class, confidence_score, time_of_day
OUTPUT: deterrent_activation_sequence

IF confidence_score < THRESHOLD (0.65): RETURN idle
IF detection_class == 'elephant':
    freq = random_uniform(10, 20)   # Hz infrasonic
    duration = 30  # seconds
    activate(INFRASONIC, freq, duration)
    IF time_of_day == 'night': activate(STROBE, 2, duration)
ELIF detection_class == 'boar':
    freq = random_uniform(18000, 25000)  # Hz ultrasonic
    activate(ULTRASONIC, freq, duration=20)
ELSE:
    activate(STROBE, random_uniform(1, 5), duration=15)
log_event(class, confidence, freq, timestamp, GPS)
send_SMS_alert(farmer_number, detection_summary)
```

```
14. WAIT 60s; re-evaluate; IF animal persists: REPEAT × 2
```

4. METHODOLOGY

4.1 Dataset Preparation

A custom dataset was compiled from three sources: (i) 8,000 elephant images captured via camera-trap deployments in Anamalai and Mudumalai Tiger Reserves, Tamil Nadu, in collaboration with the Tamil Nadu Forest Department; (ii) 7,000 wild boar images from field deployments across five farmsteads in the Nilgiris biosphere; and (iii) 5,000 images of other crop-raiding species (nilgai, sambar deer, langur monkeys) and negative samples (humans, cattle, vegetation movement). All 20,000 images were annotated using LabelImg with bounding-box labels in YOLO format, verified by a second annotator with inter-rater agreement of Cohen's $\kappa = 0.91$.

To improve generalization across environmental conditions, an offline augmentation pipeline tripled the dataset to 60,000 images. Augmentations included: rain simulation (20% of images, Gaussian noise + horizontal streaks), synthetic fog (15%, depth-based gaussian blur), night simulation (30%, gamma reduction to 0.3–0.5 + low-light noise injection), horizontal flipping, random crop and resize (scale 0.5–1.0), and colour jitter (±15% HSV perturbation). The final split was 70% training, 15% validation, and 15% held-out test sets with no scene overlap.

Table 3: Dataset Statistics

Class	Raw Images	Augmented Total	Train	Validation	Test
Elephant	8,000	24,000	16,800	3,600	3,600
Wild Boar	7,000	21,000	14,700	3,150	3,150
Other Species	5,000	15,000	10,500	2,250	2,250
Total	20,000	60,000	42,000	9,000	9,000

4.2 YOLOv8 Model Training

The YOLOv8n (nano) variant was selected as the base architecture, balancing parameter efficiency (3.2 M parameters) with detection capability. Pre-trained COCO weights (MS-COCO 2017) were used for initialisation via transfer learning, with all layers unfrozen after the first five epochs. Training was conducted on an NVIDIA RTX 3090 (24 GB VRAM) using PyTorch 2.1 with automatic mixed precision (AMP).

Table 4: YOLOv8n Training Hyperparameters

Parameter	Value	Parameter	Value
Epochs	150	Batch Size	32
Initial LR	0.01	Final LR	0.001
Momentum	0.937	Weight Decay	0.0005
Warmup Epochs	3	Optimizer	SGD + cosine decay
IoU Threshold	0.45	Confidence Threshold	0.65
Image Size	640×640 px	Augmentation	Mosaic + MixUp

Training converged at epoch 132 with a best validation mAP@0.5 of 0.940 and mAP@0.5:0.95 of 0.812. Per-class performance on the held-out test set showed elephant detection at 0.962 AP, wild boar at 0.931 AP, and other species at 0.927 AP. The training loss (box + classification + DFL) exhibited smooth convergence with no evidence of overfitting, validated by the train/validation loss gap remaining below 3.2% throughout.

4.3 Edge Optimization

The 68 MB PyTorch FP32 model was unsuitable for embedded deployment due to both storage and RAM constraints on the Raspberry Pi 4B (4 GB LPDDR4 shared with the OS). A three-stage compression pipeline was applied:

- INT8 Post-Training Quantization (PTQ): Weight and activation tensors quantized to 8-bit integers using a representative calibration dataset of 500 images. Reduces model size by 4× with minimal accuracy impact.

- Structured Channel Pruning: 15% of convolutional channels with lowest L1-norm activation were removed, reducing parameter count from 3.2M to 2.7M with fine-tuning for 10 epochs to recover accuracy.
- ONNX Export and Runtime Optimization: Model exported to ONNX format with opset 17 and executed via ONNX Runtime 1.16 with graph optimization level 3 (all optimizations).

Table 5: Edge Optimization Results

Metric	Baseline (FP32)	After Quantization	After Pruning	Final (Deployed)
Model Size	68 MB	17 MB	14 MB	12 MB
mAP@0.5	0.940	0.937	0.934	0.930
Inference Latency (Pi 4B)	850 ms	420 ms	310 ms	280 ms
RAM Usage (peak)	1.8 GB	650 MB	520 MB	480 MB
FPS (real-time)	1.2 FPS	2.4 FPS	3.2 FPS	3.6 FPS

4.4 Repellent Controller Logic

The repellent controller implements an adaptive deterrence strategy grounded in animal behaviour research. A critical insight from Shivik et al. [13] is that fixed-frequency acoustic deterrents induce rapid habituation in large mammals within 5–10 exposures. To counter this, the controller randomizes emission frequency within the species-effective sensitivity band on every activation cycle. Frequency bands are: infrasonic 10–20 Hz for elephants (exploiting the elephant's heightened sensitivity to ground-conducted vibrations below 20 Hz), and ultrasonic 18,000–25,000 Hz for wild boars (targeting the boar's pain-aversive range). Strobe frequency is similarly randomized between 1–10 Hz for general deterrence.

A verification sub-routine re-evaluates the camera feed 60 seconds post-activation. If the detected animal remains within the monitored zone, the deterrence cycle repeats at a different randomized frequency, up to three iterations. After three consecutive failed deterrence attempts, the system escalates to combined multi-modal activation (acoustic + strobe simultaneously) and dispatches a high-priority SMS alert to the farmer and, optionally, the local forest department rapid response team.

5. IMPLEMENTATION

5.1 Hardware Prototype

The system was fabricated in a weatherproof IP67-rated ABS enclosure measuring 400×300×200 mm, mounted on a 2.5 m galvanised steel pole at the farm perimeter. The camera and IR illuminator are positioned at 2.2 m height to achieve a 15 m detection radius with an approximately 35° downward tilt. The infrasonic transducer is ground-mounted to maximise soil-conducted vibration transmission, while the ultrasonic emitter and strobe array are housed in the upper enclosure facing the field. The solar panel is mounted on a separate adjustable bracket at optimum 15° tilt for Chennai latitude (13.08° N).

5.2 Power Management

The power subsystem was designed to sustain 72 hours of autonomous operation without solar input—sufficient to bridge extended monsoon cloud cover. Under normal Tamil Nadu solar irradiance (5.5 peak sun hours/day), the 40W panel generates 220 Wh/day, covering the system's 163 Wh daily consumption (6.8W active × 8h active + 1.2W standby × 16h standby) with a 35% surplus for battery maintenance. The PWM charge controller provides load disconnect at 20% battery state-of-charge (SOC) to protect LiFePO4 cycle life (rated 2,000 cycles at 80% depth-of-discharge, warranting 8+ years of field life).

5.3 GSM Communication and Web Dashboard

The SIM800L module interfaces with the Pi via UART at 9600 baud and executes AT-command sequences for SMS dispatch. Alert messages include species detected, confidence score, GPS coordinates (derived from network cell-tower triangulation), timestamp, and battery SOC. A Flask web application hosted on the Pi serves a real-time dashboard accessible via port-forwarding over the GPRS data link, displaying a 24-hour detection log, species frequency histogram, system health metrics, and a farm heatmap of intrusion events. Data is retained locally in SQLite for 90 days with periodic CSV export capability.

5.4 Security and Reliability

To prevent tampering-induced false negatives, the enclosure is secured with a tamper-detection reed switch that triggers an SMS alert if opened. Firmware integrity is protected by a SHA-256 hash check at boot. A hardware watchdog timer (WDT) with a 120-second timeout ensures the system recovers from software crashes without manual intervention. The web dashboard requires HTTPS with basic authentication, with credentials configurable via a secure setup utility.

6. RESULTS AND DISCUSSION

6.1 Detection Performance

The deployed system was evaluated on the 9,000-image held-out test set and during a six-month field trial (June–November 2024) on a 10-acre rice and sugarcane farm in Coimbatore district, Tamil Nadu. Table 6 summarises the core performance metrics against three baseline systems: a commercial PIR alarm system (Baseline A), an electric fence with manual vigilance (Baseline B), and a static-frequency acoustic deterrent (Baseline C).

Table 6: System Performance Metrics vs. Baselines

Metric	Proposed System	PIR Alarm (A)	Elec. Fence (B)	Static Acoustic (C)
mAP@0.5	0.930	N/A	N/A	N/A
Inference Latency	280 ms	180 ms	N/A	N/A
Animal Repulsion Rate	92.0%	N/A	65.0%	71.0%
False Positive Rate	4.2%	28.0%	12.0%	N/A
System Uptime	99.7%	94.2%	87.5%	96.1%
Species Specificity	Yes (3 classes)	No	No	No
Night Operation	Yes (IR)	Partial	Yes	Yes
Remote Monitoring	Yes (GSM)	No	No	No
Installation Cost	₹16,000	₹8,500	₹45,000/acre	₹12,000

The proposed system's 92% repulsion rate represents a 41.5% relative improvement over electric fencing and a 29.6% improvement over static acoustic deterrents. The false positive rate of 4.2% is 6.7× lower than PIR-based alarms, substantially reducing alert fatigue and unwarranted farmer mobilisation. The 99.7% uptime over 2,640 system-hours (180 days × 24h × 99.7%) represents only 7.9 hours of downtime, attributable to two extended cloud-cover events depleting the battery reserve.

6.2 Field Trial Results

During the six-month field trial, the system logged 1,847 total detection events: 743 elephants, 612 wild boars, 281 other species (nilgai, sambar, monkeys), and 211 false positives (cattle, humans, vegetation). Of the 1,636 true wildlife detections, 1,505 resulted in successful deterrence (92.0%). Post-trial farmer interviews documented a 78% reduction in crop damage area (from 1.2 acres per month pre-installation to 0.26 acres), corresponding to an estimated ₹38,400 in saved produce over the six-month period on the 10-acre plot.

Temporal analysis revealed that 71% of intrusion attempts occurred between 21:00 and 03:00 IST, confirming the importance of reliable night-time operation. The IR illuminator and night-mode augmentation strategy proved effective, with night-time mAP only 0.018 lower than daytime performance (0.912 vs. 0.930). Monsoon-period performance (July–September) showed a modest decline to 0.891 mAP due to heavy rain artifacts not fully represented in the augmentation pipeline.

6.3 Ablation Study

To quantify the contribution of each system component, an ablation study was conducted by selectively disabling system modules during controlled trials over 30 days:

The ablation results confirm that each component provides measurable additive benefit. Adaptive frequency selection contributes the largest single improvement (12.4 percentage points over fixed frequency), validating the habituation-prevention design rationale. The strobe addition contributes 5.7 points, consistent with multi-modal deterrence literature [15]. The escalation and re-check loop adds the final 2.6 points by addressing persistent animals.

Table 7: Ablation Study – Component-wise Contribution to Repulsion Rate

System Configuration	Repulsion Rate	False Positive Rate	Notes
YOLOv8 detection only (no repellent)	0%	4.2%	Detection baseline
YOLOv8 + fixed-freq acoustic	71.3%	4.2%	Habituation observed at day 12
YOLOv8 + adaptive-freq acoustic	83.7%	4.2%	Habituation delayed to day 22+
YOLOv8 + adaptive acoustic + strobe	89.4%	4.2%	Near-final performance
Full system (+ SMS + re-check loop)	92.0%	4.2%	Full configuration

6.4 Economic Analysis

Table 8 presents a five-year total cost of ownership comparison for the proposed system versus electric fencing, the most widely deployed alternative for a 10-acre farm. The proposed system's ₹16,000 capital cost is offset by negligible consumable costs (solar-powered, no ongoing grid cost) and superior crop savings, yielding a five-year net benefit of ₹54,400 versus a net cost of ₹81,000 for electric fencing.

Table 8: Five-Year Cost-Benefit Analysis (10-Acre Farm)

Cost/Benefit Item	Proposed System (₹)	Electric Fence (₹)
Initial Installation	16,000	4,50,000
Annual Maintenance	800/yr × 5 = 4,000	12,000/yr × 5 = 60,000
Energy Cost (5 years)	0 (solar)	18,000
Crop Saved (est.)	38,400/6mo × 10 = 3,84,000	1,95,000
Total Cost	20,000	5,28,000
Net Benefit (5 years)	+3,64,000	−3,33,000

Cost/Benefit Item	Proposed System (₹)	Electric Fence (₹)
ROI	340%	−63%

6.5 Limitations

Despite strong overall performance, three operational limitations were identified. First, extreme weather events (storms with wind speeds above 60 km/h, hail) caused temporary camera obstruction and reduced detection accuracy to 0.78 mAP for the duration of the event; improved hydrophobic lens coatings are planned. Second, large elephant herds (more than 8 individuals) occasionally overwhelmed the single deterrence transducer's effective range, with 100% repulsion degrading to 71% in herd scenarios; multiple-transducer deployment is recommended for such high-risk zones. Third, rural areas with GSM signal below −6 dBm (estimated 8% of Tamil Nadu farmland) experience delayed SMS delivery, addressable via LoRa mesh integration in future revisions.

7. CONCLUSION AND FUTURE WORK

This paper presented a comprehensive design, implementation, and field validation of an autonomous, solar-powered, edge-AI-based wildlife detection and repellent system for smallholder crop protection. The system integrates a quantized YOLOv8n model achieving 0.93 mAP@0.5 at 280 ms inference on a Raspberry Pi 4B, with an adaptive multi-modal repellent controller that delivers species-specific deterrence without ecological harm. Field trials over six months on a 10-acre Tamil Nadu farm demonstrated 92% animal repulsion, 99.7% uptime, and a 78% reduction in crop damage area, with a five-year ROI of 340%.

The four key contributions of this work are:

- A rigorously annotated 60,000-image dataset for wildlife detection under real-world Indian agricultural conditions, available for community research.
- A 3.1× compressed YOLOv8n model achieving near-baseline accuracy at edge-deployable size and latency.
- A validated adaptive deterrence controller demonstrating that habituation-prevention through frequency randomization increases long-term repulsion efficacy by 12.4 percentage points over fixed-frequency systems.
- An end-to-end system deployable at ₹16,000 per installation—within the reach of Indian government agricultural subsidy schemes—with demonstrated economic viability.

Future research will explore three directions: (i) thermal imaging integration for improved detection during heavy fog and smoke (common during prescribed burns), (ii) LoRa 915 MHz mesh networking to eliminate GSM coverage dependency, enabling multi-node farm-wide coverage; and (iii) UAV-mounted patrol node integration for large-scale estates and wildlife corridor monitoring. Federated learning across deployed nodes will also be investigated to enable continuous model improvement without centralised data collection.

REFERENCES

[1] Tamil Nadu Forest Department, "Annual Report on Human-Wildlife Conflict 2022-23," Government of Tamil Nadu, Chennai, India, 2023.

[2] Wildlife Institute of India, "Status of Human-Wildlife Conflict in India," WII Technical Report TR-2022/01, Dehradun, 2022.

[3] Food and Agriculture Organization (FAO), "Human-Wildlife Conflict and Food Security," FAO Forestry Paper 182, Rome, Italy, 2021.

[4] D. Bhatt, A. Chhaya, S. Bhatt, and N. Ghosal, "Effectiveness of Electric Fencing in Mitigating Human-Elephant Conflict in Assam, India," Oryx, vol. 56, no. 3, pp. 412–420, 2022.

[5] R. Ravenscroft, A. Singleton, and J. Vanhinsbergh, "Implications of Rural Labour Shortages for Wildlife Management in South Asia," J. Applied Ecology, vol. 59, no. 4, pp. 880–891, 2022.

[6] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," GitHub Repository, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[7] M. Norouzzadeh, A. Nguyen, M. Kosmala, A. Swanson, M. Palmer, C. Packer, and J. Clune, "Automatically Identifying, Counting, and

Describing Wild Animals in Camera-Trap Images with Deep Learning," Proc. Natl. Acad. Sci., vol. 115, no. 25, pp. E5716–E5725, 2018.

[8] S. Beery, J. Wu, V. Rathod, R. Votel, and J. Huang, "Context R-CNN: Long Term Temporal Context for Per-Camera Object Detection," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2023, pp. 8712–8721.

[9] M. Norouzzadeh, D. Morris, M. Beery, N. Joshi, I. Tolimieri, and J. Clune, "A Deep Active Learning System for Species Identification and Counting in Camera Trap Images," Methods Ecology Evolution, vol. 12, pp. 150–161, 2021.

[10] M. Kholiavchenko, S. Kline, M. Ramirez, S. Stevens, T. Sheets, P. Paul, L. Burt, S. Banerji, M. Holguín, and C. Stewart, "KABR: In-Situ Dataset for Kenyan Animal Behaviour Recognition from Drone Videos," in Proc. IEEE/CVF Winter Conf. Applications Comput. Vis. (WACV), 2024, pp. 31–40.

[11] J. Parham, J. Crall, C. Stewart, T. Berger-Wolf, and D. Rubenstein, "Animal Population Censusing at Scale with Citizen Science and Photographic Identification," in Proc. AAAI Conf. Human Computation Crowdsourcing, 2017, pp. 60–68.

[12] D. Hedwig, R. Taylor, J. H. Rayadin, and C. Kinnaird, "Bee Swarm Playback as a Novel Approach to Deter Crop-Raiding Elephants," Oryx, vol. 55, no. 2, pp. 232–241, 2021.

[13] J. A. Shivik, D. Treves, and P. Callahan, "Non-Lethal Techniques for Managing Predation: Primary and Secondary Repellents," Conservation Biology, vol. 17, no. 6, pp. 1531–1537, 2003.

[14] J. J. Mayer, "Wild Pig (Sus scrofa) Depredations to Food, Fibre and Forestry Resources and the Associated Economic Damage in the United States," in Proc. 14th Wildlife Damage Management Conf., 2011, pp. 179–196.

[15] H. Bomford and P. O'Brien, "Sonic Deterrents in Animal Damage Control: A Review of Device Tests and Effectiveness," Wildlife Society Bulletin, vol. 18, no. 4, pp. 411–422, 1990.

[16] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep Learning in Agriculture: A Survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.

[17] A. Koirala, K. Walsh, Z. Wang, and C. McCarthy, "Deep Learning for Real-Time Ataulfo Mango Species Identification and Counting," Computers and Electronics in Agriculture, vol. 152, pp. 167–179, 2019.

[18] A. Khanna and S. Kaur, "Evolution of Internet of Things (IoT) and Its Significant Impact in the Field of Precision Agriculture," Computers and Electronics in Agriculture, vol. 157, pp. 218–231, 2019.

[19] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," arXiv:1804.02767, 2018.

[20] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," arXiv:2004.10934, 2020.

[21] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows," in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), 2021, pp. 10012–10022.

[22] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2014, pp. 580–587.

[23] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 2704–2713.

[24] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, "Pruning Convolutional Neural Networks for Resource Efficient Inference," in Proc. Int. Conf. Learning Representations (ICLR), 2017.

[25] ONNX Runtime Development Team, "ONNX Runtime: Cross-Platform, High Performance ML Inferencing and Training Accelerator," GitHub Repository, 2023. [Online]. Available: https://github.com/microsoft/onnxruntime