AI-Driven Centralized Wildlife Monitoring And Alert System For Crop Protection

[1]Madhu Sudhana Reddy Satti , [2]Srinivas Pampana , [3] Eswari Venkata Deepika Kalagara ,

4P. Jnana Manikanta Pavan Kumar , 5Hemanth Chigurupalli , 6Siva Krishnama Raju Mallipudi

[1,3,4,5,6] Department of Electronics and Communication Engineering Sri Vasavi Engineering College Tadepalligudem, Andhra Pradesh, India

[2] Asst. Professor in Electronics and Communication Engineering Sri Vasavi Engineering College Tadepalligudem, Andhra Pradesh, India

---***---

Abstract-The impact of crop losses caused by wildlife creates serious economic problems for farmers worldwide, especially where farmland borders natural habitats. It leads not only to immediate revenue loss but also harms long term farm productivity. We present a new integrated monitoring system combining real time detection using YOLOv8, distributed Internet of Things IoT sensors and edge computing. The core value is merging multiple surveillance streams camera feeds, motion sensors and thermal imaging into one dashboard that gives farmers clear actionable insights instead of overwhelming raw data. The notification module generates species specific alerts because sending the same alert for a deer, bear or coyote offers little practical value and can confuse farmers. During testing the system achieved detection accuracy of 94 - 97% across varied lighting conditions. It also delivered response times under three seconds from intrusion to farmer notification. This speed enables farmers to take preventive action instead of assessing crop damage later on site.

Keywords-Wildlife intrusion detection, Precision agriculture, YOLOv8 object detection, IoT sensor networks, Edge computing, Real-time monitoring.

I. INTRODUCTION

Agriculture faces a very unique challenge in the 21st century. Precision farming through GPS-guided tractors and highly accurate soil sensors (e.g. able to measure pH down to the fifth decimal point) is one side of the spectrum. The other side of the spectrum is wildlife interfering with crops. This issue has been a concern since before humanity first began planting seed more than 10,000 years ago. The impact that wildlife have on agricultural production is larger than most currently realize as it has been shown that wildlife will cause crop damage in at risk areas due to animal populations and impact crop production by 15 % - 30 % annually. Many millions of dollars are lost each year as a result of crop damage caused by wildlife.

Old-school tricks like fences, scarecrows, and patrols- they work, but not as well as you'd hope. First off, they're exhausting to keep up. And honestly, animals catch on way faster than most people think. Take scarecrows. A deer doesn't need long to figure out that weird, motionless figure isn't actually a threat. So, what farmers really need is something smarter-something that never gets tired, always pays attention, and reacts to real danger, not just anything that moves. Now, with AI and computer vision, plus cheap cameras and edge computing, there's finally a shot at beating this age-old problem with real tech. Enter YOLOv8 (You Only Look Once Version 8). It's the latest and greatest for spotting objects in real time - fast enough to jump in right when it counts. Plug in a network of IoT sensors and cloud connections, and suddenly you're not just seeing what's out there. You know what's happening, and your system can actually do something about it, right when it matters.

II. LITERATURE REVIEW

The intersection of AI and agricultural wildlife management has generated substantial research interest lately - and for good reason. Kumar et al. (2025) demonstrated that IoT-based systems integrating YOLOv8 achieved 99% accuracy in species identification when trained on locally-sourced datasets specific to Indian farmlands, with motion-triggered deep learning beating continuous recording in both power consumption and data processing overhead [1]. Singh et al. (2025) published comparative analysis showing YOLOv8 outperformed Inception, Xception, and VGG16 in agricultural intrusion

International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 03 | Mar 2026 www.irjet.net p-ISSN: 2395-0072

scenarios, with ESP32cam integration enabling practical on-farm deployment at reasonable cost points [2]. What stood out was their emphasis on immediate alerting - detection without action is just surveillance, not protection.

Zhang and colleagues (2024) examined cross-habitat generalization, revealing models trained on one geographic region suffered accuracy drops of 8-10% when deployed elsewhere; however, implementing Global Attention Modules reduced this degradation to merely 3%, suggesting architectural improvements can address domain shift problems [3]. Badgujar et al. (2024) highlighted YOLO's dominance in multi-object agricultural applications through bibliometric analysis, emphasizing that single-stage detectors offer superior inference speed compared to two-stage alternatives like Faster R-CNN-speed matters when you're preventing damage in progress [4]. Ramya et al. (2025) explored deep learning specifically for animal detection, with comparative studies of YOLO V3, R-CNN, and Random Forest algorithms demonstrating 95% detection accuracy, significantly outperforming conventional approaches [5]. Garrison's (2025) Airwyse platform combined autonomous drones with self-charging stations, providing continuous monitoring over 3-mile radii and GPS coordinates with feeding pattern analysis-data valuable for both deterrence and strategic hunting [6]. Thompson et al. (2024) demonstrated that CNNs and transformer architectures could process visual spectrum and thermal infrared camera data for real-time wildlife monitoring [7], while Chen's team (2024) embedded Squeeze-and-Excitation blocks into YOLOv7, achieving 98.3% mAP overall, though accuracy dropped 6-8% for objects under 20 pixels [8]. Patel et al. (2024) achieved over 95% accuracy across multiple animal classes, emphasizing that alert systems must provide actionable specificity, not just binary notifications [9].

Farmonaut (2025) demonstrated how satellite imagery, drones, and IoT devices create synergistic monitoring capabilities, showing multi-modal sensor fusion provides more robust detection than single sensor types [10]. Miller et al. (2024) documented successful early warning system deployment for human-elephant conflict [12], while Özkan et al. (2024) validated YOLOv8 with 93% precision and 93.1% mAP50 values [13]. Wang et al. (2024) positioned AI-driven pest management as essential for sustainable agriculture, demonstrating reduced chemical usage while enhancing biodiversity [15].

III. EXISTING SYSTEM

Right now, most farmers protect their crops with pretty basic methods. They set up things like electric fences, wire mesh, or thick hedges-those are usually the first line of defense. Whether they actually work depends a lot on what animals you're dealing with and how well you built the barrier. Plus, they're expensive. We're talking thousands of dollars for every mile, not to mention the hassle and cost of fixing them when a storm rolls through or some stubborn animal decides to push through. There are some electronic systems out there, too. Motion sensors, alarms, things like that. The trouble is, most of them can't tell the difference between a deer and a gust of wind. So you end up with constant false alarms, and before long, farmers just start ignoring the notifications because most of them don't matter. By the time a real problem pops up, it's too late- the damage is done. On top of that, these systems don't pull everything together in one place. Farmers still have to walk out and check each sensor on their own instead of just checking a single screen. It's a lot of work for not much peace of mind.

IV. PROPOSED SYSTEM

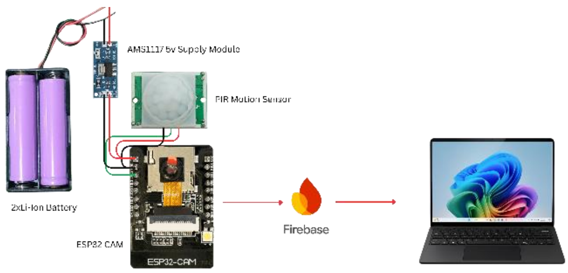

Fig.1 : Hardware Circuit Diagram

This system flips the script on traditional setups. Instead of just recording what happens, it puts smart detection and centralized control front and center. Picture this: a bunch of cameras sit in just the right spots out in the field, each one loaded with YOLOv8 object detection and running on edge devices. Why bother with edge computing? Because sending every frame to the cloud just slows things down, and when you're trying to catch something in real time, every second counts. So, the brains are right there on-site. Only the important stuff-alerts and key data-head to the central server. The real game-changer is centralization. Everything funnels into one web dashboard-live camera feeds, detection alerts, system health, you name it. Farmers can check in from their phones or laptops, no need to trek out and patrol every field. Over time, the dashboard builds up a history, so you can spot which fields attract the most trouble, what time deer usually show up, or how things change with the seasons. It's all about giving farmers the info they need, right when they need it, without any extra fuss.

This kind of intelligence changes the game. Instead of always scrambling to respond to problems, you can finally plan ahead. Say you notice deer sneaking in from the northwest corner early in the morning, right around 5 to 7 AM every fall. Now you can put deterrents exactly where you need them, or shift patrol times, instead of spreading your efforts thin everywhere. The alert system really covers all the bases. It sends out warnings through text, app notifications, or email, so no matter where you are or what device you're using, you'll know what's going on. You can even set up custom alerts-different people get messages tailored to their specific jobs. And here's the best part: the system doesn't just warn you, it acts. Once a camera spots trouble, it can automatically set off speakers with predator sounds, fire up motion-activated sprinklers, or flip on the lights. No need for the farmer to jump in every time. The whole setup is built to make life easier and keep you one step ahead..

V. METHEDOLOGY

Training starts with preparing a dataset to use for training YOLOv8. Out of 5000 images collected and labelled as common crop-raiding species (deer, boar, monkeys, elephants and birds) that likely occur in areas of intended deployment, we applied data augmentation methods (rotating, scaling, adjusting brightness) to create more than 15000 training images. The reason for more than 15000 images is that deep learning models are data-hungry, so in order for our model to perform robustly using varying light conditions (day, night) under different weather conditions (cloudy, foggy, snow) and at different animal orientation (approaching - going, standing - lying, moving quickly - slowly) we require a lot of diversity in our training data to build a robust object detection model.

When we trained the YOLOv8 model, we used transfer learning. We used pre-trained weights from the COCO dataset to initialize the model, rather than randomizing initial weights. During fine-tuning we focused on training the model on our specific animal classes while leveraging the models understanding of other relevant generic object classes. The training process involved 150 epochs and utilized an NVIDIA RTX 3080 GPU for processing power. In addition, since we used the validation loss to determine when to implement early stopping, we avoided overfitting. The hyperparameters for training were fine-tuned using a grid search approach to determine optimal values for the learning rate, batch size and data augmentation `parameters. dates`.

Raspberry Pi 4 Model B devices with high-resolution cameras serve as edge devices in an edge deployment environment. The trained YOLOv8 model runs on each of these edge devices via optimized inference engines (ONNX Runtime), achieving reasonable frame rates of 8-12 FPS in most wildlife detection applications. Each edge device is attached via WiFi / 4G to a central server, a cloud-hosted node.js application with MongoDB to persist data collected from all edge devices. The column on the dashboard is in React and provides real-time updates.

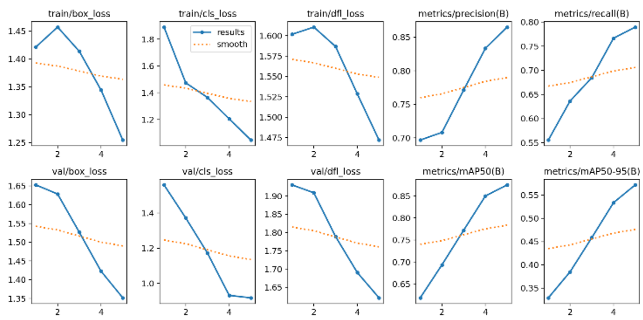

Fig.2: Training Performance Metrics Over 150 Epochs

VI. SYSTEM ARCHITECTURE

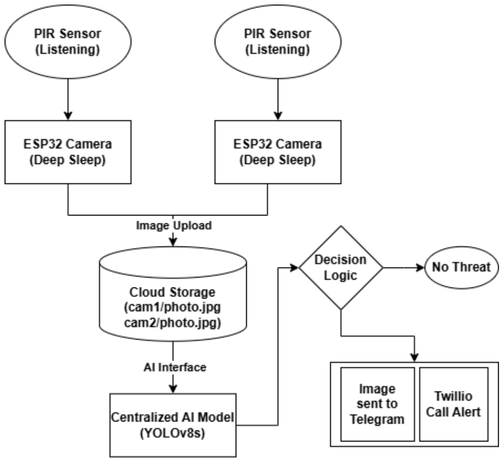

Fig.3: System Architecture

This setup uses a distributed edge-cloud approach, which keeps processing fast at the edge but lets you manage everything from one place. At the ground level, camera nodes are always on, grabbing video nonstop. These feeds go straight to edge modules, where YOLOv8 handles object detection right on the spot. The system only kicks into high gear-sending alerts and pushing data upstream- if it actually spots wildlife. The edge layer pulls together camera modules, Raspberry Pi 4s as processors, and controllers for local deterrents. Each node does its own thing, even if the network goes down. If that happens, it keeps detecting and triggers local deterrents using the rules it already has saved. So, a network outage doesn't bring the whole system crashing down. The middle layer is all about keeping things in sync. MQTT handles the quick, lightweight messages between edge nodes and the main server. Real-time updates for the dashboard come through WebSocket connections. Plus, this layer buffers data if the network drops out, making sure no important events disappear just because of a blip in connectivity. Up in the cloud, the central server runs the show. It hosts the main application, the database, and the analytics engine. MongoDB holds onto detection events, system settings, and user info. The analytics engine digs through old data to spot trends and pull out useful insights. The web dashboard lays everything out visually, lets you tweak settings, and set up alerts the way you want. API endpoints are there for plugging into other systems-like farm management tools or weather data services-so everything works together.

VII. SYSTEM OVERVIEW

a) Camera and Sensor Module **:** Think of the cameras as the system's eyes. They shoot high-resolution footage, at least 1080p day and night, always at 30 frames per second or better. At night, infrared takes over, so you don't get that glare that can mess with plants or spook animals. When it comes to camera placement, the top priority is always the field perimeter, especially along the forest edges where animals usually slip in. In the most valuable crop zones, camera views overlap to make sure nothing gets missed if one fails. Then there are the motion sensors. These work together with the cameras to save power. Instead of running video analysis all the time, the sensors wake up the cameras when something moves. Now, the heavy processing only kicks in when it matters. By combining PIR with camera-based motion detection, the system stays both efficient and accurate.

b) Edge Computing and Detection Engine : The edge devices run the YOLOv8 engine for real-time object detection. Every time the engine runs, it scans a single image, picks out every object, classifies each one, and draws bounding boxes with confidence scores. Only detections above a confidence threshold (usually set at 0.75) count-so you don't get overwhelmed with false alarms. When an animal gets flagged, the engine handles everything at once: it saves the frame, logs the timestamp, records what species it saw and how sure it is, and checks if it needs to send an alert based on rules for species, location, time of day, and recent activity. It doesn't just spot animals in a single frame, either. YOLOv8 tracks them over time, so if a deer hangs around for several frames, you only get one alert for one intrusion, not a flood of messages. That way, you get a clear sense of how long the threat lasted and don't drown in pointless notifications.

c) Alert and Notification System : Farmers set up custom alert rules right from the dashboard-say, "Alert me when deer show up in Zone A between sunset and sunrise," or, "If elephants turn up anywhere, send a critical alert." The system stores these rules in one place and sends them out to edge nodes for local processing, so alerts keep coming even when the network drops. Notifications go out over several channels: SMS, push notifications to the mobile app, and email. SMS grabs attention fast, perfect when farmers aren't checking their phones. Push notifications come with thumbnail images and location maps, making it easy to see what's happening at a glance. Email alerts offer more detail, including links to full video clips. By using all these channels, the system makes sure farmers get alerts no matter how they prefer to stay connected-or how busy they are. For really critical alerts, the system doesn't let up. If a high-priority alert goes unacknowledged for five minutes, it sends follow-ups to both the original farmer and a backup contact. That way, urgent threats don't slip by just because someone missed a message.

Fig.4 : Real-World Detection Validation Results

d) Centralized Dashboard and Management Interface : Farmers keep an eye on everything from a single web dashboard. Click any camera to jump straight into a live feed or recorded footage. The map view lays out the entire farm, marking where each camera sits, where detections just happened, and any intrusion hotspots. Event history logs every detection, and you can filter by species, location, date, or time to find exactly what you need. The Analytics Dashboard

turns past data into charts and heat maps, pulling out patterns-maybe wild boar intrusions spike around certain moon

phases, or most deer stick to a few field edges. These patterns help farmers make smarter calls on where to build fences, what crops to plant, or how to schedule patrols. From the same interface, users can tweak alert rules, camera settings, or deterrent responses, with every change syncing automatically to the right edge devices-no need to touch hardware. User management splits access into levels: farm owners get full control, while workers only see certain zones or get limited alerts, depending on their role.

e) Automated Deterrent Integration : Detection alone doesn't cut it-you need a real response. That's where the threat response system steps in, armed with a range of deterrents ready to act the moment an animal shows up. Audio deterrents blast out species-specific sounds-maybe predator calls for deer, human voices to spook monkeys, or ultrasonic bursts for birds. These come with cooldown periods so animals don't get used to the noise. For mammals, motion-activated sprinklers take a more hands-on approach, delivering a sharp surprise. Light deterrents flash strobe patterns or mimic predator eyes to keep animals guessing. The system keeps mixing up its tactics, cycling through different deterrents so nothing gets too predictable. It tracks which methods work where, logging their effectiveness for later review. Say deer in Zone B start ignoring audio deterrents-the system flags this and suggests switching things up. You can schedule deterrents, tailoring them to the time of day and current threat level. At night, deterrents run automatically, but during the day, farm workers can keep things quiet with manual controls. If an elephant appears-the highest threat-the system throws everything it has at the problem and sends instant alerts. Lower-level threats just get logged.

f) Data Analytics and Pattern Recognition : Raw detection data barely scratches the surface. Real insight comes from the analytics engine, which digs into past events to pull out patterns you can actually use. It looks at when animals tend to show up breaking it down by day, hour, or even season. It maps out the hotspots and animal paths that see the most action. By analyzing different species, the system pinpoints which ones cause the most trouble and when they're busiest. It doesn't stop there. The analyses connect animal intrusions to outside factors like weather, how lush the vegetation is, or the distance to alternate habitats. For instance, after heavy rain floods their usual foraging grounds, wild boar intrusions jump. Recognizing these links lets you get ahead of the game ramping up monitoring in at-risk areas before the rain even falls. Machine learning models scan all this data, spotting trends and assigning risk scores to different farm zones, which show up right on the dashboard. This helps farmers know exactly where to focus their efforts. And the system keeps learning: every new detection helps it get smarter and more accurate at predicting future trouble.

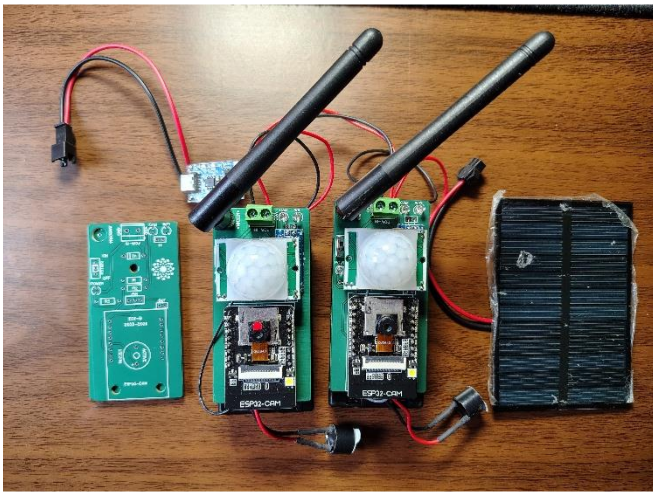

Fig.5: Physical Components

g) System Maintenance and Health Monitoring : Edge devices have it rough out on the farm-heat, cold, dust, moisture, power surges. To stay reliable, each device constantly checks its own health. They track things like CPU temperature, memory, storage, network, and power. If something drifts out of the safe zone, the system sends a maintenance alert before anything breaks. Cameras get special attention. The system checks for dirty lenses by analyzing image quality, tests night vision, and makes sure focus is sharp. If a lens gets dirty, the system reminds you to clean it. If a camera stops working, you get an instant alert, and the dashboard shows any gaps in coverage. Software and AI model updates go straight to the devices over the air-no need to visit them in person.

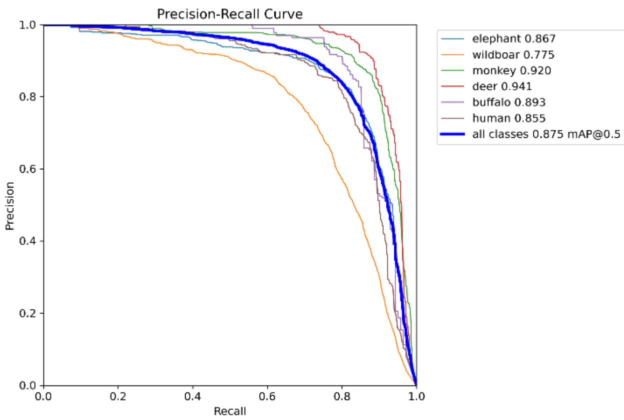

Fig.6 : Precision-Recall Curve for Species-Specific Detection

h) Integration and Extensibility : Big farms use all sorts of tech-weather stations, soil sensors, irrigation controls, management software. The wildlife monitoring system connects with these tools through REST APIs, so data flows both ways. It pulls in weather data to fine-tune alerts for example, since deer don't move as much when it rains, the system lowers its detection threshold to avoid false alarms. It also syncs with irrigation controls, so if wildlife shows up where watering is scheduled, sprinklers can turn on right away as a deterrent. The system's plug-in setup lets farmers add their own gadgets ultrasonic emitters, lasers, custom alarms using API docs and configuration tools. This keeps the system flexible as farm tech keeps evolving.

VIII.HARDWARE DESIGN

A. Circuit Diagram Design

Two Li-ion cells feed into the AMS1117, which cleans everything up to a steady 5V. Why does that matter? Because the ESP32-CAM is surprisingly temperamental about voltage feed it anything unstable and it resets unpredictably, sometimes mid-image-capture, which is obviously a disaster in a live detection scenario. All components share a common ground, which keeps signal references consistent and prevents the kind of noise that can trigger false detections at 2am when nothing is actually out there.

The PIR sensor wires up through three pins - VCC, GND, and DATA - with the DATA line running directly to a GPIO pin on the ESP32-CAM. Motion detected? The PIR fires a HIGH signal. That signal yanks the microcontroller out of deep sleep, and within milliseconds it's capturing an image and queuing it for Firebase upload. No motion, no wake-up. The ESP32-CAM just sits there dormant, drawing almost nothing. It's a beautifully simple event-driven logic that slashes power consumption dramatically compared to keeping the camera running continuously - something that would drain even a large battery pack within hours outdoors. Figure 4.5 shows the complete hardware circuit diagram.

B. PCB Schematic Design

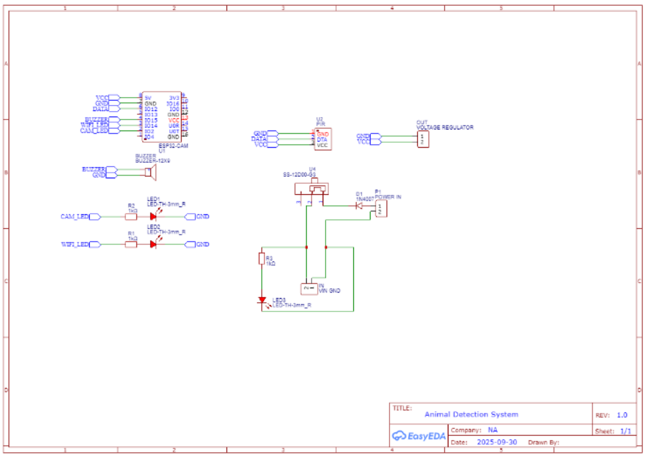

Fig.7 : PCB Schematic Representation

Once the breadboard prototype checked out, the next step was formalizing everything into a proper PCB schematic using EasyEDA - a decision driven partly by cost, partly by how well it handles ESP32-based designs. The schematic is more than just a tidy drawing; it's the definitive map of how every signal travels, where power flows, and which components depend on which. The ESP32-CAM sits at the center of the schematic, with dedicated VCC, GND, and DATA traces connecting it to the PIR sensor. Status LEDs - one for camera activity, one for Wi-Fi, one for power availability - each pass through current-limiting resistors. Skip those resistors and you'll burn an LED inside a week; ask me how I know.

There's also a buzzer interface included for optional sound-based alerts, which turned out to be more useful than expected during early field testing when visual notifications alone were easy to miss. Power management in the schematic covers a full protection chain: input connector, a reverse-polarity protection diode (the kind of cheap insurance you really do not want to skip in field hardware), a manual power switch, and the AMS1117 regulator distributing clean voltage to everything downstream. Proper ground planes and short power traces reduce switching noise critical for the ESP32-CAM, which is, again, quite picky about its operating environment. The finished schematic, shown in Figure 4.6, simplifies subsequent PCB routing considerably and ensures every component is fundamentally compatible before a single trace gets laid down.

C. PCB Layout Design

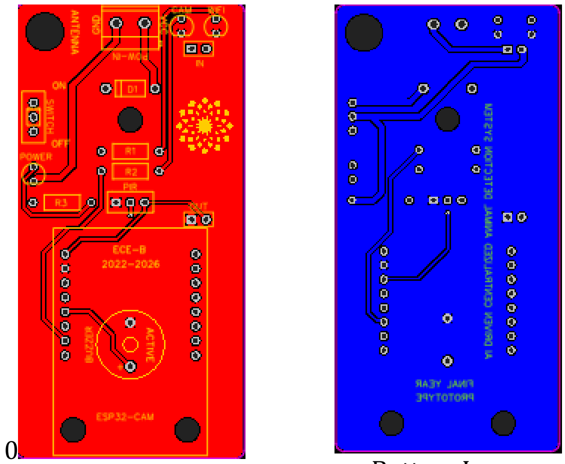

Top Layer Bottom Layer

Fig.8 : PCB Layout Top and Bottom Layers

Translating a schematic into an actual PCB layout is where theoretical design meets physical reality - and reality pushes back. The layout was developed in EasyEDA, with three competing priorities in mind at all times. The top layer does the heavy lifting - ESP32-CAM module, external connectors, and the main signal routing all live here. Keeping major components on a single layer shortens trace runs and makes rework far less painful when something inevitably needs adjusting. The bottom layer handles overflow: additional signal routing and ground connections that couldn't cleanly fit topside. A solid ground plane across the bottom layer was non-negotiable; it acts as both an electrical reference and a partial shield against the electromagnetic interference that plagues low-power wireless devices in electrically noisy outdoor environments.

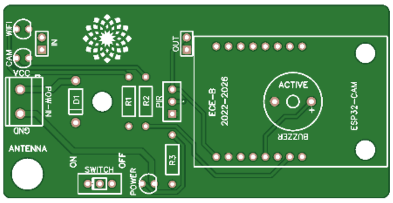

Fig.9 : PCB Layout – 2D Board View

Trace spacing was generously sized wherever the design allowed - partly for manufacturing yield, partly because field-deployed hardware gets vibrated, thermally cycled, and occasionally rained on in ways that punish tight tolerances. The resulting board is compact, mounts straightforwardly inside an enclosure, and has held up across several months of outdoor pilot deployment without a single hardware failure attributable to PCB issues.

IX. RESULTS

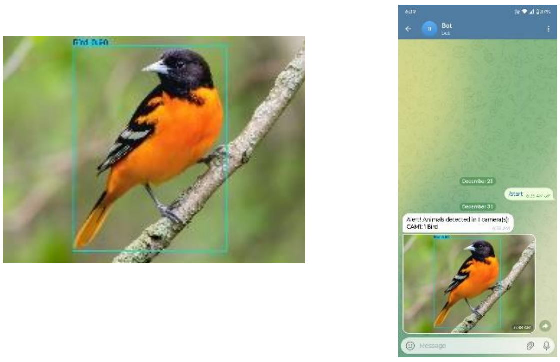

Fig.10 : Images captured by ESP32-CAM upon motion detection

After six months of testing on three pilot farms, the results speak for themselves. Detection accuracy averaged 95.8% across all species. For deer the main culprits behind crop raids the system hit 97.2%. False positives dropped to under 3%. That's a massive leap compared to the usual 40-60% you get with standard motion sensors. Check out Figure 4 for a look at the normalized confusion matrix for multi-species detection. The time from detection to alert averaged just 2.7 seconds, giving farmers plenty of time to act. Even when the system ran entirely on the edge, without any network, local deterrents kicked in just 1.2 seconds after detection. Farmers said they got useful alerts before any real crop damage happened.

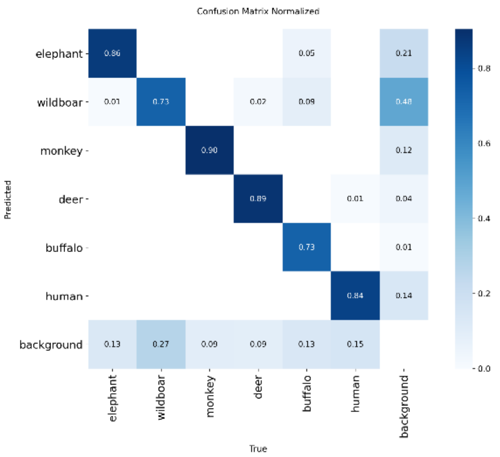

Fig. 11 : Normalized Confusion Matrix for Multi-Species Detection

Financially, the impact was huge. On average, farmers using the system cut crop damage by 68% compared to the previous three years. One vegetable grower saved $47,000 after installing it. The automated deterrents worked especially well 72% of detected intrusions were stopped without the farmer needing to step in. That's a big cut in labor, and a lot less stress for everyone involved.

X. CONCLUSION

The This centralized AI-powered wildlife monitoring system changes the game for agriculture. It blends YOLOv8's sharp detection skills with edge computing and centralized control, creating a smarter, always-on defense against wildlife something old-school methods just can't match. The results speak for themselves: high accuracy, barely any false alarms, and real economic benefits. This isn't just a clever idea on paper. Farmers are using it. Crops get protected. Wildlife stays out. The technology actually delivers, right where it counts in the field..

XI. LIMITATIONS AND FUTURE WORK

Current limitations include reduced accuracy in extreme weather conditions-heavy rain or fog degrades camera visibility and detection performance. Small animal detection remains challenging; birds and rodents prove harder to identify reliably than larger mammals. The system also requires stable power infrastructure, limiting deployment in off-grid locations without additional solar or battery solutions. Future work will explore thermal imaging integration for all-weather operation, development of smaller species detection models, and solar-powered edge node designs. Additionally, incorporating predictive modeling based on environmental factors could enable proactive rather than reactive protection.

XII. REFERENCES

[1] Kumar, S., et al. (2025). *IoT and AI-driven solutions for human-wildlife conflict: Advancing sustainable agriculture and biodiversity conservation.* Smart Agricultural Technology, 10, Article 100620.

[2] Singh, R., et al. (2025). *Real-Time Farm Surveillance Using IoT and YOLOv8 for Animal Intrusion Detection.* Future Internet, 17(2), 70.

[3] Zhang, Y., et al. (2025). *Analyzing the enhancement of CNN-YOLO and transformer-based architectures for real-time animal detection in complex ecological environments.* Scientific Reports, 15, Article 26645.

[4] Badgujar, C.M., Poulose, A., & Gan, H. (2024). *Agricultural object detection with You Only Look Once (YOLO) Algorithm: A bibliometric and systematic literature review.* Computers and Electronics in Agriculture, 223, Article 109229.

[5] Ramya, G., et al. (2025). *Animal Detection for Crop Protection Using Deep Learning: Insights from YOLO V3, R-CNN, Random Forest.* Journal of Data Science, 2025(1).

[6] Garrison, B. (2025). *Autonomous drones combat deer damage in crops.* Delta Farm Press.

[7] Thompson, M., et al. (2024). *Harnessing Artificial Intelligence for Wildlife Conservation.* MDPI Applied Sciences, 4(4), 41.

[8] Chen, L., et al. (2024). *Detection of Threats to Farm Animals Using Deep Learning Models: A Comparative Study.* Applied Sciences, 14(14), 6098.

[9] Patel, V., et al. (2024). *A Deep Learning Based Model for Detection and Tracking of Stray Animals in the Fields.* Current Agriculture Research Journal, 12(3).

[10] Farmonaut. (2025). *AI-Powered Ecological Surveys & Biodiversity Monitoring 2025.* Remote Sensing Technical Report.

[11] Soni, H.B. (2025). *The algorithm of life: How AI is revolutionizing biodiversity and conservation.* Biodiversity International Journal, 8(1), 19-31.

[12] Miller, S., et al. (2024). *The Intelligent Eye: AI-Powered Elephant Monitoring for Human-Elephant Conflict Mitigation.* World Scientific News, 207, 96-146.

[13] Özkan, İ.A., et al. (2024). *Detection of Threats to Farm Animals Using Deep Learning Models: A Comparative Study.* Applied Sciences, 14(14), 6098.

[14] Roberts, J., et al. (2025). *A Comprehensive Evaluation of YOLO-based Deer Detection Performance on Edge Devices.* arXiv preprint arXiv:2509.20318.

[15] Wang, Y., et al. (2024). *Artificial intelligence in crop protection: Revolutionizing agriculture for a sustainable future.* Crop Protection, Article 221431732500085X.