

Smart Farm Animal Intrusion Detection Using Multi-Modal IoT, Edge AI and Kafka Streaming

Ashish Gawande^{*}

IIOT Cloud Engineer (SLB)

Email: ashish.gawande@gmail.com

Abstract

Wild animals such as deer or rabbits cause significant crop losses worldwide and create major problems for farmers. Traditional methods, such as fences, are often too expensive, difficult to maintain and not consistently effective. This paper provides details about an animal intrusion avoiding system that uses a combination of sensors, including thermal cameras, microphones, and motion detectors. The sensors work in conjunction with artificial intelligence running on edge devices, sending alerts through a Kafka streaming system. Unlike existing systems that rely on only one type of sensor or fixed models, our method combines several signals and can also learn to recognize new animals with only a few examples. In simulation tests, the system achieved about ninety-four percent accuracy, reduced false alarms by more than a third, and responded in less than two hundred milliseconds. When compared with systems that used only motion sensors or only cameras, our approach proved to be more reliable. The work is still limited because it is based on simulations rather than real-world farm testing, but plans include real-world trials, adding long-range communication such as LoRaWAN, and utilizing advanced techniques like federated learning to make the system even stronger.

Keywords: IoT; Kafka; Edge AI; Few-Shot Learning; Smart Agriculture; Wildlife Detection; TinyML.

1. Introduction

Wild animals remain a significant problem for farmers. Deer trample young plants, boars dig up large parts of fields, and rabbits can destroy crops very quickly. These intrusions cause billions of dollars in losses every year. Traditional methods such as fences, alarms, or human patrols are costly and often not flexible enough to deal with changing threats. With the rise of new technologies, such as edge artificial intelligence and IoT sensors, farms now have the opportunity to detect and respond to animal threats in real-time. In this paper, we describe a system that brings together several layers of technology. It utilizes various types of sensors, applies rapid detection on edge devices, sends alerts via Kafka, and relies on cloud-based few-shot learning to recognize new animals without retraining the entire model.

Received: 9/14/2025

Accepted: 11/14/2025

Published: 11/26/2025

^{*} Corresponding author.

In our simulation tests, the system was able to detect animals with more than ninety-four percent accuracy and respond in less than two hundred milliseconds. Wildlife intrusions continue to reduce farm productivity on a global scale. Existing approaches, such as fences, manual patrols, and simple motion sensors, are either too costly or not adaptable enough for the wide variety of animal behaviors. Advances in edge AI, IoT sensor fusion, and event streaming now provide the possibility of real-time adaptive intrusion detection.

Current systems have apparent gaps. Most rely on only one type of sensor, such as a PIR detector or a camera, and do not combine multiple sources of information. They also do not include adaptive learning to handle new animals. Many depend only on cloud processing, which can lead to delays and requires constant connectivity.

The contributions of our work are as follows. We propose a multimodal IoT system that combines thermal, visual, and acoustic data for stronger detection. We can design a low-latency edge AI framework with YOLOv8 lite models that are optimized for NVIDIA Jetson Nano. [2] We integrate Kafka-based streaming to provide scalable and reliable message delivery. We introduce a few-shot learning method to allow the system to recognize new animals with very little data. Finally, we provide a comparison of our approach against systems that use only PIR sensors or only cameras.

2. Related Work

Most animal detection systems on farms today use only cameras or simple motion sensors such as PIR. Some systems apply artificial intelligence, but these usually depend on fixed models that cannot easily adapt to new situations. Very few approaches make use of sound or heat sensors, and even fewer are designed to work at a large scale across real farms.

Some studies have explored few-shot learning to detect new objects, but this method is rarely used in farm environments where models must run directly on small edge devices. TinyML has demonstrated promising results in controlled environments, such as greenhouses, but open rural areas with diverse animal populations remain a challenge.

Another weakness of current systems is that they are hard to update or customize once they are deployed. This makes them less useful when conditions change. In contrast, our system is designed to be adaptable, affordable, and easily expandable to different farms. It operates in real-time and can continue to improve over time without requiring full retraining. By combining Kafka for fast data streaming, multiple sensors for enhanced accuracy, and cloud support for machine learning, we aim to provide a practical and scalable solution for modern agriculture.

A. System Improvements Over Existing Work

Earlier research has shown the limits of camera-based systems, PIR detectors, and single-sensor AI models. These methods struggle when there is poor lighting, animal occlusion, or background noise. Few-shot learning has been tested in computer vision tasks but is not commonly applied to farm-based animal detection. Recent work also shows the promise of TinyML and edge AI. Still, most of this research focuses on greenhouse

monitoring or livestock tracking, rather than keeping wild animals out of open fields.

Our work addresses these gaps by combining three kinds of sensing: thermal, acoustic, and motion. This multi-modal setup improves reliability. We also use sensor fusion scoring to cut down on false alarms. Finally, we integrate few-shot learning so that the system can quickly adapt to new animal species while still running efficiently on farm-ready edge devices.

B. Dataset

For our simulations, we can use thermal images of deer, boars, and rabbits collected from public repositories, along with sound recordings of wildlife. The data was divided into multiple sets like training, validation and test sets. To improve the robustness of the model, preprocessing steps included normalization and data augmentation techniques such as adding noise, rotating images, and adjusting contrast.

C. Model Training

We used a lightweight version of YOLOv8 as the detection model. [2] The model was trained on the combined thermal and acoustic datasets and optimized to run on the NVIDIA Jetson Nano using TensorRT. Few-shot learning was applied with as few as three to five labeled examples, which allowed the system to recognize new animals without needing a large amount of additional data.

D. Evaluation Metrics

To measure system performance, we tracked four main metrics: detection accuracy, processing latency in milliseconds, false positive rate, and throughput measured as the number of messages processed per second. These results provided a clear view of how the system performed under simulated farm conditions.

3. System Architecture

Our system is built using four interconnected layers that work together to detect animals in real-time, classify them accurately, and update the model as new species emerge. This setup is designed to work across the farm.

A. IoT Sensor Layer- The first layer is the IoT Sensor Layer, which includes thermal cameras, motion sensors (PIR), and directional microphones. These sensors help detect heat from animals, movements in the field, and sounds such as footsteps or growls. Using more than one type of sensor makes the system more reliable, especially in poor lighting or noisy conditions.

B. Edge AI Layer- Next is the Edge AI Layer, where small devices, such as the Jetson Nano, run lightweight AI models, including YOLOv8-lite. [2] These models quickly analyze data from the sensors and identify the type of animal present. The system combines data from all sensors to make more informed decisions, rather than relying on just one.

C. Kafka Streaming Layer- The Kafka Streaming Layer handles the transmission of alerts and data. Kafka brokers receive detection events and forward them to other parts of the system, including cloud servers

and local deterrent tools. Messages are organized based on animal type, location, and time, which helps in taking quick and accurate action.

D. Cloud Learning Layer-Finally, the Cloud Learning Layer is responsible for keeping the system up to date. If an unfamiliar animal is detected, the cloud collects a few examples, retrain the detection model, and sends the improved version back to edge devices using over-the-air updates. This allows the system to learn and adapt without needing a complete reset.




Figure 1

4. Edge implementation

The edge part of our system is built using the NVIDIA Jetson Nano, which offers a good mix of computing power and energy efficiency. This makes it an excellent choice for farms where power use and cost are essential. We run a lighter version of the YOLOv8 model on the device, which has been optimized for speed and efficiency, utilizing minimal memory and power. [2]

At the edge, the AI processes data in real time using input from thermal cameras, regular cameras, and microphones. It checks for heat signatures, movement, and animal sounds simultaneously. To reduce false alarms, the system employs a technique called temporal smoothing, which helps confirm that the animal is truly present and not merely passing by briefly or making noise.

When a threat is confirmed, the device sends an alert almost instantly using Kafka, with the total delay from detection to response being under 200 milliseconds. This fast reaction time is essential for triggering deterrents before damage occurs.

The system also uses a scoring method to combine results from different sensors. For example, if thermal data

indicates heat, motion is detected, and a possible animal sound is present, the system will treat it as a high-confidence alert. If the score is too low or uncertain, the event is logged for later review, but no action is taken. This helps strike a balance between speed and accuracy, avoiding the waste of resources on false positives.




Figure 2

5. Cloud and Few-Shot Training

The cloud system is designed to support few-shot learning using a method called Prototypical Networks. When edge devices detect an animal that the model does not recognize, they mark it as "unknown" and send a few labeled images to the cloud. These images are stored in a temporary dataset for review and evaluation. If a new type of animal is confirmed, such as a porcupine, the system initiates a brief learning process.

First, it creates image embeddings using a pre-trained ResNet model, which turns the images into numerical patterns the system can understand. Then it uses only three to five images to teach the model about this new animal. After learning, the model is updated to include the new animal type. This improved model is then sent to all edge devices via over-the-air updates.

This loop runs continuously, helping the system stay up to date without requiring full retraining from scratch. The retraining jobs on the cloud are scheduled to run automatically every night using Kubernetes. This way, the system continues to improve over time, even with limited data, and farmers don't have to update anything manually.




Figure 3

6. Kafka Streaming and Deterrent Systems

Kafka serves as the central communication channel between edge devices, deterrent systems, and the cloud-based learning engine. Every time an edge device detects an animal, an event is created in JSON format and sent to a specific Kafka topic. These topics are organized by animal type and location, such as alerts for deer in zone 1, alerts for unknown animals in zone 3, or audio events during night shift. Each message contains helpful details, including the time of detection, GPS location, animal type (if identified), and the level of confidence in the detection.

To support older sensors that still use MQTT, Kafka Connect is used to bring their data into the same system. Kafka Streams helps process data in real-time, for example, by checking if the same animal keeps returning to the same place and deciding whether to escalate the alert.

Several necessary settings are also in place to ensure the system's reliability. Events are stored with a replication factor of three, ensuring that the data is not lost in the event of an error. Logs are saved for seven days, allowing the system to replay old data for retraining or debugging purposes. Access to Kafka is managed through permissions, ensuring that only trusted devices and services can send or receive data.

Once a detection event is published, the deterrent layer listens to the alert and takes immediate action. Depending on the type of animal, it may turn on flashing lights, play loud sounds, activate ultrasonic emitters, or spray water to scare the animal away. These actions are chosen based on what works best for each animal and the time of day. This fast reaction helps protect the crops before any damage is done.




Figure 4

7. Results and Evaluation

The performance of the proposed system was evaluated through simulations using publicly available thermal image datasets and animal sound recordings for species such as deer, boars, and rabbits. These simulations were designed to estimate how the system might perform under typical farm conditions, but no real-time or controlled testing has been conducted yet.

In simulation, the edge AI model running on the Jetson Nano was able to process each image in approximately 45 milliseconds. The estimated time from detection to deterrent response, including Kafka message delivery, was under 200 milliseconds. With three rounds of simulated retraining using few-shot learning, the model reached a projected detection accuracy of 94.2 percent. Sensor fusion helped reduce false positives by around 35 percent.

Simulation Results (with base assumptions)

Detection Accuracy (94.2%): Projected accuracy based on simulation experiments using publicly available thermal image datasets and wildlife audio recordings. The system was not validated on a real-world test set. The figure comes from model training on mixed modalities (thermal + acoustic) under simulated farm conditions.

Latency (187 ms end-to-end): Estimated using published Jetson Nano inference benchmarks (~45 ms per frame), Kafka broker delivery (~80 ms), and assumed deterrent response times (~60 ms). No physical deployment was conducted.

False Positive Reduction (35%): Expected improvement derived by comparing simulated single-modality detection logs (PIR-only, vision-only) against multi-modal fusion runs. No real-world false positive testing has been performed.

Kafka Throughput (15,000 messages/sec): Based on Apache Kafka performance benchmarks and simulation of 1,000 sensors generating events, not measured in a farm deployment.

Limitation –

These results are based entirely on simulated inputs and estimated performance benchmarks. While the findings are promising, future work will involve controlled experiments and real-world testing to validate the system in actual farm environments.

Security And Cost Analysis

The system includes several built-in security measures to keep data and devices safe. All Kafka messages are encrypted using TLS, which protects information as it moves between devices. Over-the-air updates sent from the cloud to the edge devices are digitally signed and checked before installation to make sure they have not been changed. Sensors and edge devices use mutual authentication, meaning both sides verify each other before sharing data. Kafka topics and system components are protected using access control lists, which prevent unauthorized users from sending or receiving messages. These security features help ensure that only trusted devices and services can access the system.

Cost:

For a small 10-acre test farm, the total hardware cost is estimated at around \$1,230. This includes one Jetson Nano for processing, five PIR sensors for motion detection, two thermal cameras, two microphones, a Kafka broker device, and basic networking equipment.

Table 1

Item	Quantity	Estimated Cost
Jetson Nano	1	\$120
PIR Sensors	5	\$100
Thermal Cameras	2	\$500
Microphones	2	\$60
Kafka Broker Device	1	\$300
Networking Equipment	-	\$150
Total		\$1,230

This cost is much lower than building a full fence around the same area. Unlike fencing, which cannot adapt to new threats, this smart system can grow and improve over time. In many cases, the return on investment can be achieved within a single growing season due to reduced crop damage and lower labor costs.

8. Security and cost analysis

Security:

The system includes several built-in security measures to keep data and devices safe. All Kafka messages are encrypted using TLS, which protects information as it moves between devices. Over-the-air updates sent from the cloud to the edge devices are digitally signed and checked before installation to make sure they have not been changed. Sensors and edge devices use mutual authentication, meaning both sides verify each other before sharing data. Kafka topics and system components are protected using access control lists, which prevent unauthorized users from sending or receiving messages. These security features help ensure that only trusted devices and services can access the system.

Table 2

Method	Accuracy	False Positives	Latency
PIR-only	76%	High	<100 ms
Camera-only (YOLOv4-tiny)	88%	Medium	250 ms
Proposed Multi-modal	94.2%	Low	187 ms

Cost:

For a small 10-acre test farm, the total hardware cost is estimated at around \$1,230. This includes one Jetson Nano for processing, five PIR sensors for motion detection, two thermal cameras, two microphones, a Kafka broker device, and basic networking equipment.

Table 3

Item	Quantity	Estimated Cost
Jetson Nano	1	\$120
PIR Sensors	5	\$100
Thermal Cameras	2	\$500
Microphones	2	\$60
Kafka Broker Device	1	\$300
Networking Equipment	-	\$150
Total		\$1,230

This cost is much lower than building a full fence around the same area. Unlike fencing, which cannot adapt to new threats, this smart system can grow and improve over time. In many cases, the return on investment can be achieved within a single growing season due to reduced crop damage and lower labor costs.

9. Conclusion

This paper presents a smart system for detecting and deterring wild animals in real time using IoT sensors, edge AI, Kafka-based data streaming, and cloud-supported few-shot learning. By combining different types of sensors such as thermal cameras, motion detectors, and microphones, the system is able to detect animals more accurately and respond faster. The edge AI model processes information quickly, while the cloud helps the system learn about new animals using only a few examples. Together, this setup achieves quick response times under 200 milliseconds and significantly lowers the number of false alarms.

The system is designed to be secure, flexible, and easy to scale. It can be deployed on different types of farms without needing heavy infrastructure. Compared to traditional methods like fencing or manual monitoring, it offers a cost-effective and more intelligent way to protect crops and livestock. The approach can help farmers reduce losses and save time while adapting to changing threats in the environment.

In the future, we plan to explore more advanced features such as using a group of drones to guide animals away from farmland. We also aim to support wider coverage using LoRaWAN for low-power remote communication and to test the use of satellite images to monitor large farms from above. These improvements will help create a fully connected and intelligent farm monitoring system that can work in real-world conditions on a larger scale.

Data Availability: The thermal image dataset used in this research is publicly available through the FLIR

ADAS thermal dataset (<https://www.flir.com/oem/adas/adas-dataset-form/>). The acoustic recordings are from publicly available wildlife sound libraries such as the Macaulay Library (<https://www.macaulaylibrary.org/>) and Xeno-canto (<https://www.xeno-canto.org/>).

References

[1] Edge AI for Agriculture:

<https://www.mdpi.com/2077-0472/12/6/892>

[2] YOLOv8-lite on Jetson Nano:

<https://github.com/ultralytics/ultralytics>

[3] Kafka Event Streaming:

<https://kafka.apache.org/documentation/>

[4] Few-Shot Learning in Vision:

<https://arxiv.org/abs/1904.04232>

[5] MQTT vs Kafka:

<https://www.hivemq.com/blog/mqtt-vs-kafka-real-time-bidirectional-data-processing/>

[6] Secure OTA Updates for IoT:

<https://source.android.com/docs/core/ota>

[7] TinyML for Smart Farms:

<https://www.sciencedirect.com/science/article/pii/S2772375524000959>

[8] Raspberry Pi & Jetson Benchmarks:

<https://developer.nvidia.com/embedded/jetson-nano-developer-kit>

[9] Prototypical Networks (Few-Shot Learning)

<https://arxiv.org/abs/1703.05175>

[10] Lightweight Edge AI with YOLOv4-Tiny

<https://arxiv.org/abs/2004.10934>

[11]Kakfa

<https://docs.confluent.io/kafka/introduction.html>

[12] TinyML for Environmental Sensing

<https://github.com/tinyMLx/courseware>

[13] Animal Detection in Agriculture

<https://ijarcce.com/wp-content/uploads/2022/05/IJARCCE.2022.114159.pdf>

[14] Google cloud

<https://cloud.google.com/architecture/iot-core-reference-architecture>