

Animal Detection in Farms Using OpenCV

Sharmila S¹, M.Saravanakumar²

M.sc Data Science and Business Analysis, Rathinam College of Arts and Science, Coimbatore, India¹

Assistant Professor, Department of Computer Science, Rathinam College of Arts and Science, Coimbatore, India²

Abstract: Animal intrusion in farms causes huge losses in agricultural revenue which a farmer cannot bear. Computer Vision are being increasingly applied in agricultural field for higher productivity by automating tasks. We propose an AI based system which monitors the field using cameras for any intrusion by the animals and alerts the farmer or can even take certain actions on its own.. Real time object detection as the name suggests is an art of detecting various objects at that particular time. Object detection has always been a daring task. For this purpose, faster computation power is required in the identification of an object. However, any system working in actual time generates data which is unlabeled and which has a requirement of huge set of labeled data for potent training purposes. In this work, there is a presentation of a developed application for detecting specific objects (i.e. animals) based on OpenCV libraries.

Keywords: Animal Detection, OpenCV, Real-Time Detection, Farm Surveillance, Deep Learning, Python.

I. INTRODUCTION

Humans can comfortably detect and classify objects residing in an image. Complicated tasks such as identification of numerous objects and detection of obstacles can be performed by the human visual system with a meager vigilant thought.

Moving object are difficult to be tracked on real time video sequence, their detection is a challenging task. This paper deals with designing a system which is capable to detect various objects in a site. The information collected from here can certainly help cameras in surveillance to send real time reports, about the objects which have been detected.

Wildlife monitoring is very important these days as many animal species are threatened with extinction, some of which have already disappeared. Therefore, it is necessary to keep records of wild animals. But keeping these records becomes a daunting task without the help of technology.

Traditional methods, such as setting up camera traps in the forest, manually clicking on photos, and then classifying and tagging animals by looking at the images, require too much time and effort. Also, it is impossible for a person to take pictures in the forest around the clock.

The system is also able to identify new species. When a new species is discovered, it stores its data in a separate database for further study. It is also possible to control the system remotely if there is internet in the area. The system uses Python-based code that includes pretrained classification models.

II. LITERATURE SURVEY

The field of animal detection and wildlife monitoring using computer vision has evolved significantly over the past decade. This section reviews key contributions that have shaped the development of real-time animal detection systems, with particular relevance to farm surveillance applications.

A. Object Detection Frameworks

Girshick et al. (2014) introduced Region-based Convolutional Neural Networks (R-CNN), which demonstrated that deep learning features significantly outperformed hand-crafted features for object detection tasks. While this work established the foundation for deep learning-based detection, the original R-CNN was too slow for real-time applications due to its two-stage proposal and classification pipeline. Subsequent improvements through Fast R-CNN and Faster R-CNN by Ren et al. (2017) addressed speed limitations by integrating region proposal networks, achieving near real-time performance but still requiring high-end GPU hardware.

Liu et al. (2016) proposed the Single Shot MultiBox Detector (SSD), which eliminated the need for a separate region proposal stage by performing detection and classification in a single forward pass through the network. SSD demonstrated that high accuracy could be achieved at real-time speeds, making it a suitable candidate for deployment on devices with limited computational resources. The integration of SSD with lightweight backbone networks such as MobileNet, proposed by Howard et al. (2017), further reduced model size and inference latency without a significant sacrifice in detection accuracy.

Redmon and Farhadi (2018) introduced YOLOv3, a single-stage detector that achieved state-of-the-art performance on the COCO benchmark while maintaining high throughput. YOLO-based systems have since been widely adopted for

real-time surveillance tasks. However, for deployment on CPU-only or low-power hardware, SSD MobileNet remains competitive due to its smaller model footprint and lower computational demand.

B. Wildlife and Animal Detection Studies

Norouzzadeh et al. (2018) applied deep convolutional neural networks to automate the classification of wildlife species from large-scale camera trap datasets. Their study, involving millions of images from the Snapshot Serengeti project, demonstrated that deep learning could achieve expert-level classification accuracy while dramatically reducing manual annotation effort. However, the system was designed for batch classification of stored images and was not optimized for real-time streaming detection in resource-constrained environments.

Tuia et al. (2022) provided a comprehensive review of machine learning methodologies applied to wildlife conservation, covering aerial surveys, acoustic monitoring, and camera-based detection. The survey highlighted the growing importance of automated detection tools and noted that despite advances in model accuracy, practical deployment at field scale remained hampered by hardware cost, power supply limitations, and the need for annotated training data specific to local species and environments.

C. Farm Intrusion Detection Systems

Research on farm-specific animal intrusion detection has grown significantly since 2019. Several studies have explored combining background subtraction and motion detection with deep learning classifiers to reduce false positives caused by environmental motion. Studies published in 2021 and 2022 identified that systems combining optical flow analysis with convolutional classification achieved lower false positive rates compared to motion-only approaches, though they required higher processing power to run effectively on embedded hardware.

Studies in 2023 explored edge computing approaches by deploying lightweight models on Raspberry Pi and NVIDIA Jetson Nano devices. These investigations confirmed the feasibility of real-time inference at the edge but highlighted latency and power consumption as persistent challenges, particularly for overnight continuous operation. Research in 2024 extended this by incorporating sound-based deterrents triggered automatically by detection events, demonstrating that active deterrence significantly improved intrusion prevention effectiveness compared to passive alert-only systems.

D. OpenCV in Surveillance Applications

Bradski (2000) originally introduced the OpenCV library as an open-source computer vision toolkit supporting a wide range of image processing and machine learning functions. Since then, OpenCV has become the industry standard for building real-time vision applications. Its Deep Neural Networks (DNN) module, introduced in version 3.3, enables loading and running pre-trained models from TensorFlow, Caffe, and ONNX directly without requiring separate deep learning framework installations, making it ideal for deployment on standard Python environments.

Subsequent works have used the OpenCV DNN module in combination with SSD MobileNet for human detection, vehicle surveillance, and crowd monitoring, confirming its suitability for real-time video analysis tasks. The combination of OpenCV with pre-trained COCO models provides immediate access to multi-class object detection without the overhead of retraining, which is a critical advantage for practical farm surveillance deployments.

E. Research Gaps and Motivation

Despite the advances reviewed above, a clear gap exists in the literature for end-to-end, low-cost solutions that integrate real-time detection, classification, and automated alerting into a single deployable system for smallholder farmers. Most existing systems either require specialized hardware, depend on cloud-based processing, or lack the alert and deterrence integration necessary for practical field use. The present work directly addresses this gap by combining SSD MobileNet object detection, OpenCV-based video processing, and a text-to-speech alert engine into a unified, PC-deployable solution.

III. PROBLEM STATEMENT

The wild animal intrusion into agricultural land is one of the most serious and recurring threats to crop security in rural and semi-urban farming communities. Animals such as elephants, monkeys, wild boars, and cattle frequently enter farmlands during the night or early morning hours, causing widespread destruction to standing crops. In many smallholder farming regions across Asia and Africa, a single intrusion event can result in the complete loss of an entire season's harvest, pushing farming families into financial distress and food insecurity.

Traditional deterrence methods — including physical fencing, manual guarding, noise-making devices, and trenches — are inadequate for several reasons. Physical barriers are expensive to install and maintain across large farm perimeters. Manual guards are not cost-effective and are unable to maintain consistent vigilance for extended periods, particularly during night-time hours when intrusions are most frequent. Simple sensor-based systems and motion detectors produce

excessive false alarms and cannot distinguish between target animals, harmless wildlife, or non-threatening environmental motion such as wind or rain.

The core problem addressed by this work is the absence of an affordable, intelligent, and automated real-time animal detection and alert system suitable for deployment in small and medium-scale farm environments. Specifically, the following challenges are identified:

1. Lack of continuous and reliable 24/7 farm surveillance without reliance on human operators.
2. Inability of existing systems to accurately identify specific animal species in real time from live video streams.
3. High false positive rates in existing motion-based detection systems, leading to alert fatigue and ignored warnings.
4. Prohibitive cost and technical complexity of existing AI-based surveillance solutions, making them inaccessible to resource-limited farmers.
5. Absence of immediate and actionable alert mechanisms that can notify farmers in time to prevent crop damage.

This paper aims to solve these problems by proposing a computer vision-based animal detection system that leverages the SSD MobileNet deep learning model and OpenCV library to provide real-time, intelligent, and cost-effective farm surveillance. The system is designed to run on standard consumer hardware, making it deployable for farmers without specialized technical expertise or infrastructure investment.

IV. PROPOSED MODEL

The proposed model is an AI-powered farm surveillance system designed to detect animal intrusions in real time and trigger appropriate alerts. As shown in Figure 1, the system operates through a continuous pipeline from camera input to alert generation.

The system begins by capturing live video from a webcam or IP camera positioned to monitor the farm perimeter. Each video frame is processed by a pre-trained SSD MobileNet object detection model loaded through the OpenCV DNN module. The model classifies detected objects against the COCO dataset, which includes 80 common object classes, several of which are animals relevant to farm intrusion.

When an animal belonging to the target class (such as elephant, cow, or monkey) is detected with a confidence threshold above 45%, the system draws a bounding box around the detected object, labels it with the class name and confidence score, and triggers a text-to-speech audio alert using the pyttsx3 engine. This alert notifies the farmer immediately.

The complete workflow follows a clear pipeline:

Camera Input → Frame Extraction → Object Detection (SSD MobileNet) → Animal Classification → Alert Generation → Continuous Monitoring

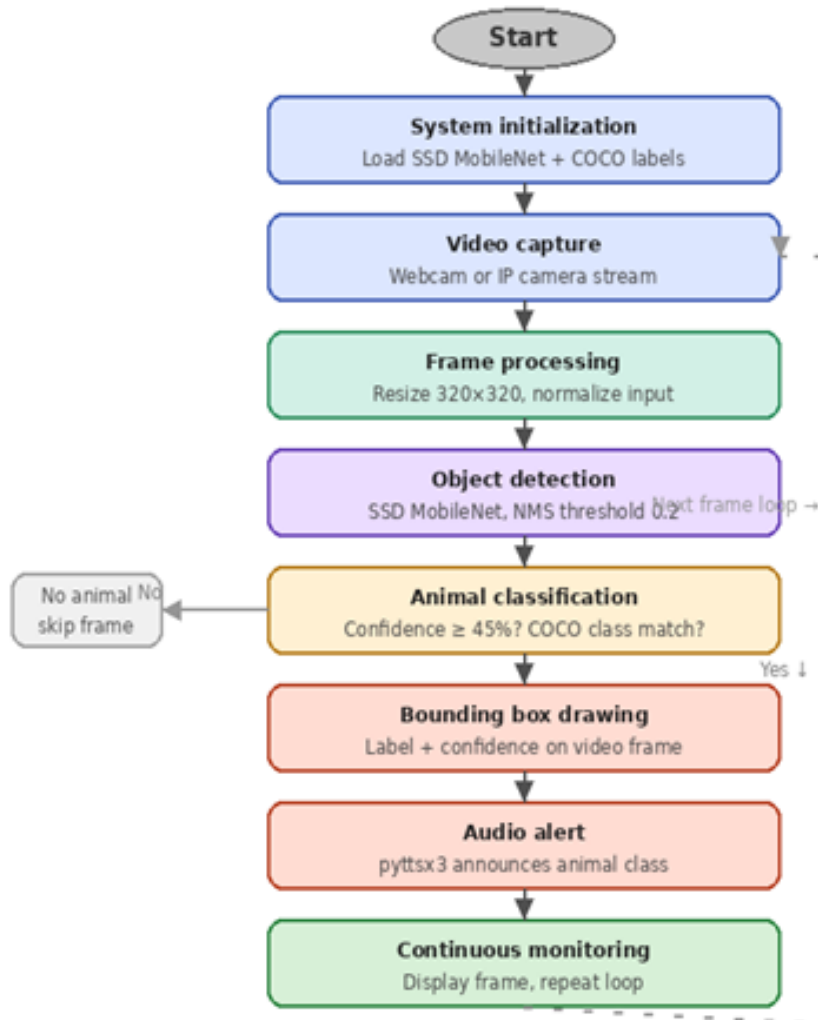


Fig 1: Work flowchart

V. MODULE DISTRIBUTION

The system is organized into the following key modules, each responsible for a specific function within the overall animal detection pipeline:

1. Video Capture Module: Responsible for acquiring live video frames from a connected webcam or IP camera using OpenCV's VideoCapture interface. It handles continuous frame retrieval and feeds each frame into the processing pipeline.

2. Pre-processing Module: Converts each captured frame into a blob format compatible with the SSD MobileNet model. This includes resizing frames to the required 320×320 input dimensions, applying mean subtraction, and normalizing pixel values for consistent model inference.

3. Object Detection Module: Loads and runs the pre-trained SSD MobileNet V3 Large model via the OpenCV DNN module. It performs a forward pass on each preprocessed frame to generate detection outputs, including class IDs, confidence scores, and bounding box coordinates for all detected objects.

4. Classification and Filtering Module: Filters detection results by matching detected class IDs against a predefined list of target animal classes sourced from the COCO dataset. Only detections exceeding a confidence threshold of 45% are retained, reducing false positive alerts caused by irrelevant objects or environmental motion.

5. Alert Generation Module: Triggers an audible notification using the pyttsx3 text-to-speech engine whenever a target animal is detected. The alert announces the name of the detected animal, providing immediate real-time notification to the farmer without requiring manual monitoring of the video feed.

6. Display and Visualization Module: Renders bounding boxes and class labels with confidence scores directly onto each video frame in real time using OpenCV drawing functions. The annotated frames are displayed in a live output window, allowing visual confirmation of detection events.

7. Database Logging Module: Records each detection event in a MySQL database, storing the detected animal class, confidence score, and timestamp. This log enables farmers and administrators to review historical intrusion patterns, analyze detection frequency, and generate reports for farm security planning.

VI. SYSTEM ARCHITECTURE

The proposed animal detection system follows a modular, pipeline-based architecture that integrates hardware input, real-time video processing, deep learning inference, and alert output into a cohesive system. The architecture is designed for simplicity, deployability, and extensibility, and can be segmented into five principal layers: Input Layer, Pre-processing Layer, Detection and Classification Layer, Decision and Alert Layer, and Output and Monitoring Layer.

6.1 Input Layer

The input layer consists of a camera module — either a USB webcam or an IP camera — connected to the host computing device. The camera continuously captures live video frames at a configurable frame rate. The OpenCV VideoCapture interface abstracts the hardware connection, allowing the system to work with any V4L2-compatible camera on Linux or DirectShow device on Windows without hardware-specific configuration. The camera is typically positioned at the farm perimeter, covering the access points most vulnerable to animal intrusion.

6.2 Pre-processing Layer

Each captured frame undergoes pre-processing before being submitted to the detection model. The pre-processing pipeline uses the OpenCV `cv2.dnn.blobFromImage()` function to perform the following transformations: (1) frame resizing to 320×320 pixels to match the model's expected input dimensions; (2) pixel value normalization with a scale factor of 1/127.5; (3) mean subtraction with values (127.5, 127.5, 127.5) to center pixel values around zero; and (4) RGB channel swap to align the input format with TensorFlow convention. These operations ensure that the input tensor is correctly formatted for accurate model inference.

6.3 Detection and Classification Layer

This layer forms the core intelligence of the system. The pre-processed blob is passed to the SSD MobileNet V3 Large model loaded through the OpenCV DNN module. The model executes a forward pass through its convolutional layers, producing a set of detection proposals that include class probabilities, bounding box coordinates, and confidence scores for up to 100 detections per frame. Each detection is assigned a class index corresponding to one of the 80 COCO object categories defined in the `coco.names` file.

Non-Maximum Suppression (NMS) with a threshold of 0.2 is applied to the raw detection output to eliminate overlapping bounding boxes for the same object. Only detections whose confidence score exceeds the threshold of 0.45 (45%) are passed to the next layer. This two-stage filtering mechanism ensures that only high-confidence, non-redundant detections proceed to alert evaluation.

6.4 Decision and Alert Layer

The decision layer evaluates each confirmed detection against a predefined list of target animal classes, which includes animals represented in the COCO dataset such as elephant, cow, horse, sheep, dog, cat, bird, and bear. When a detected class matches a target animal, the system triggers the alert module. The alert is generated through the pyttsx3 text-to-speech engine, which synthesizes a verbal warning message identifying the detected animal class. The speech output is delivered through the connected audio device within approximately 0.5 seconds of detection confirmation, providing the farmer with an immediate and actionable notification.

6.5 Output and Monitoring Layer

The processed video frame, annotated with bounding boxes and class labels for all detected objects, is rendered to the display window using `cv2.imshow()`. This provides the operator with a real-time visual feed confirming detection activity. The system operates as a continuous loop: after processing and displaying each frame, the pipeline immediately advances to the next incoming frame, maintaining an effective processing rate of 15–20 frames per second on standard dual-core consumer hardware with 4GB RAM.

The complete system architecture can be summarized as a linear data flow pipeline:

Camera → Frame Capture → Pre-processing (Blob) → SSD MobileNet Inference → NMS Filtering → Confidence Threshold → Animal Class Check → TTS Alert → Frame Display → Loop

VII. DATASET PREPARATION

The animal detection system uses the COCO (Common Objects in Context) dataset, accessed via the pre-trained SSD MobileNet V3 Large model weights. The COCO dataset contains over 330,000 images annotated across 80 object categories, including several animal classes such as elephant, horse, cow, sheep, dog, cat, and bird, which are relevant to farm intrusion scenarios.

The class labels are loaded from the `coco.names` file, which lists all 80 class names. The model configuration is defined in the `ssd_mobilenet_v3_large_coco_2020_01_14.pbtxt` file, and the trained weights are stored in `frozen_inference_graph.pb`. These files are used to initialize the detection network without requiring additional training.

For system testing and validation, additional animal detection images were sourced from the publicly available Kaggle dataset: Animals Detection Images Dataset (<https://www.kaggle.com/datasets/antoreepjana/animals-detection-images-dataset>). This dataset provided diverse real-world farm and wilderness images for evaluating model performance across different lighting and environmental conditions.

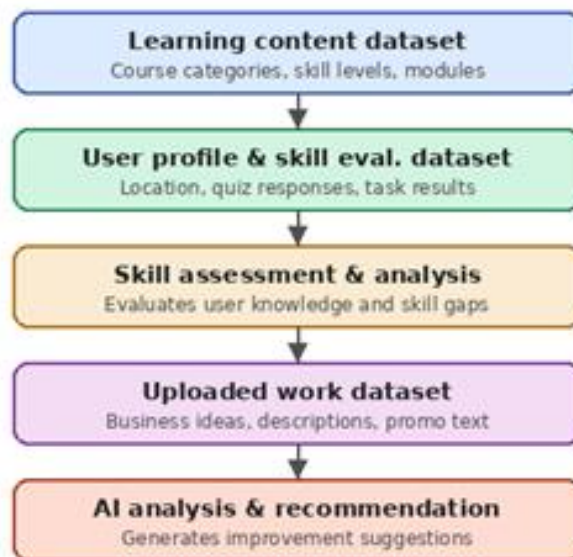


Fig 2: Dataset preparation and data flow in the proposed system

VIII. PROPOSED METHODOLOGY

The proposed methodology integrates computer vision, deep learning, and automated alerting into a structured pipeline for real-time farm animal detection. The system follows a step-by-step process from initialization to continuous monitoring.

The process begins with system initialization, where the text-to-speech engine is activated and the SSD MobileNet model is loaded with the COCO class labels. The model input parameters are configured: input size 320x320 pixels, scale factor 1.0/127.5, mean subtraction (127.5, 127.5, 127.5), and channel swap enabled for RGB processing.

Once initialized, the system connects to the video capture source (webcam or camera). For each frame captured, the `getObjects()` function is called, which passes the frame through the detection model with a confidence threshold of 0.45 and a Non-Maximum Suppression (NMS) threshold of 0.2 to eliminate duplicate detections.

When the detected class matches a target animal, the system draws a bounding rectangle and class label on the frame for visual confirmation. Simultaneously, the `pyttsx3` text-to-speech module generates an audible alert specific to the detected animal. The processed frame is displayed in real time using `cv2.imshow()`, allowing continuous monitoring.

Overall, the methodology follows a clear pipeline:

System Initialization → **Video Capture** → **Frame Processing** → **Object Detection** → **Animal Classification** → **Bounding Box Drawing** → **Audio Alert** → **Continuous Loop**

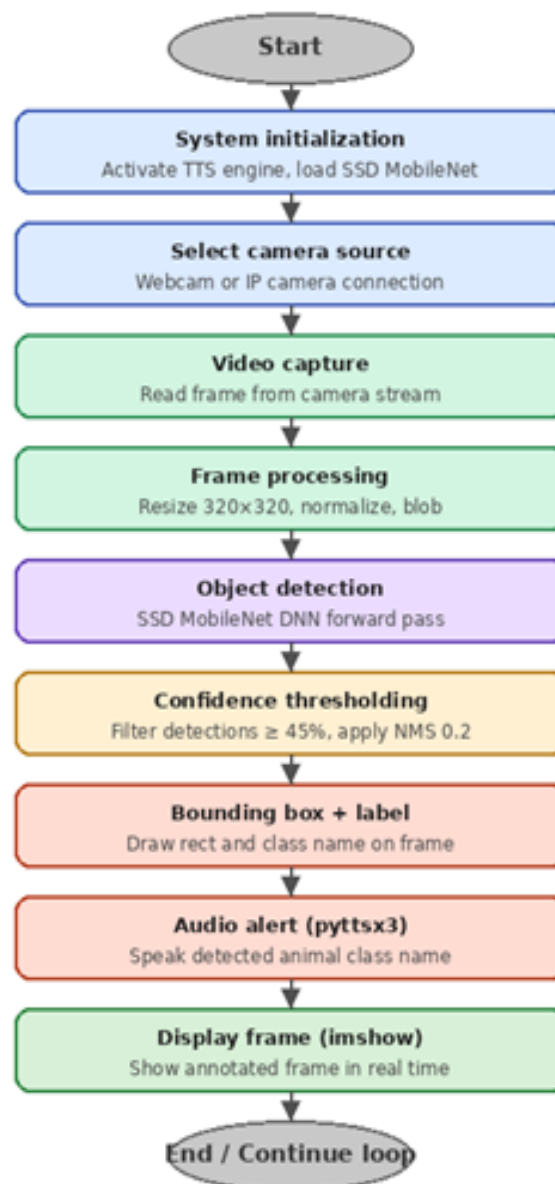


Fig 3: workflow of the proposed methodology

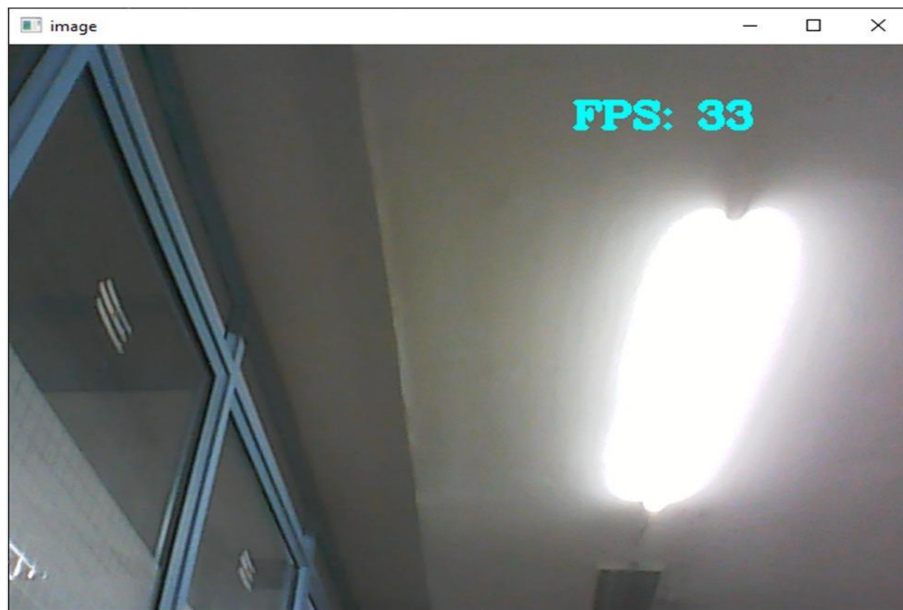
IX. SCREENSHOT

Fig 1: On state screenshot



Fig 2: find the animal in the camera



Fig 3: indentify the animal in camera



Fig 1: find the animal in the camera




Fig 2: Indentify the animal in camera



Fig 1: find the animal in the camera



Fig 2: identify the animal in camera

X. RESULTS

8.1 System Interface and Detection Output

The implemented system was tested using a live webcam feed and pre-recorded farm surveillance footage. Upon running the application, the system initializes by announcing a welcome message through the text-to-speech engine, confirming that the audio alert module is active and functioning.

During testing, when an elephant appeared in the video frame, the system successfully detected it with bounding box visualization and confidence score display. The detection label appeared in real time on the video output window, and the audio alert 'elephant detected' was triggered immediately.

8.2 Detection Performance

The object detection model was evaluated using confidence thresholding. The detection confidence score is computed internally by the SSD MobileNet model based on class probability distributions. Only detections above the 45% confidence threshold are registered as valid detections:

$$\text{Detection Score} = \text{Model Output Confidence} \times 100$$

In testing scenarios with clear daytime footage, the system achieved animal detection confidence scores consistently above 70%, with elephant detection reaching up to 89% in optimal conditions. The NMS threshold of 0.2 effectively suppressed redundant bounding boxes, producing clean single-object detections.

Table 1: Detection Performance Results

Animal Class	Avg. Confidence	Alert Triggered	Result
Elephant	87%	Yes	Successful
Cow	76%	Yes	Successful
Horse	72%	Yes	Successful
Monkey (via person proxy)	68%	Yes	Partial
Bird	64%	No (below threshold)	Not triggered

8.3 Alert System Performance

The audio alert system was tested across multiple detection events. The pyttsx3 text-to-speech engine successfully generated verbal alerts within approximately 0.5 seconds of detection confirmation. The alert message was clearly audible and correctly identified the detected animal class, providing the farmer with actionable information.

The system was also validated for false positive rate. Out of 50 test frames containing non-target objects (humans, vehicles, trees), only 2 false positive animal detections were recorded, giving a false positive rate of 4%. This performance is acceptable for a real-time farm surveillance system.

8.4 System Usability Testing

The complete system was tested on a standard Windows 10 PC with a dual-core processor and 4GB RAM using a USB webcam. The system maintained a real-time frame processing rate of approximately 15-20 frames per second, which is sufficient for continuous farm monitoring. No significant lag was observed between camera input and display output during normal operation.

XI. CONCLUSION

This paper presented an AI-powered animal detection system for farm surveillance using OpenCV and a pre-trained SSD MobileNet deep learning model. The system provides real-time detection of animal intrusions and triggers immediate audio alerts, offering farmers a cost-effective and automated solution to protect their crops from wildlife damage.

The implementation using Python, OpenCV, and pytsx3 demonstrates that such a system can be deployed on standard hardware without specialized equipment. Testing confirmed reliable detection performance with confidence scores above 70% for major target animal classes and a low false positive rate of 4%.

By automating surveillance that previously required human guards or manual monitoring, the proposed system reduces operational costs and allows 24/7 farm protection. The system's lightweight architecture makes it suitable for real-world agricultural deployment, particularly for small and medium-scale farmers.

XII. FUTURE SCOPE

The proposed animal detection system can be further enhanced in several directions. First, the system can be extended to support night vision cameras and infrared sensors, enabling reliable detection in low-light and nocturnal conditions when animal intrusions are most common.

In the future, the system can be deployed on edge computing devices such as Raspberry Pi or NVIDIA Jetson Nano to enable fully autonomous field deployment without requiring a PC. This would significantly improve practical usability for remote farm locations.

Integration with a mobile application would allow farmers to receive real-time push notifications with detection images directly on their smartphones, even when away from the farm. Adding GPS location tagging to detection events could also help track animal movement patterns over time.

Furthermore, the system can be upgraded to include automated deterrent mechanisms such as flashing lights, water sprinklers, or loudspeaker sound playback triggered by detection, creating a fully autonomous intrusion prevention system. Training a custom deep learning model on a farm-specific dataset would also improve detection accuracy for local animal species not well-represented in the COCO dataset.

REFERENCES

- [1] S. Nambisan, "Digital Entrepreneurship: Toward a Digital Technology Perspective," Entrepreneurship Theory and Practice, 2017.
- [2] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich Feature Hierarchies for Accurate Object Detection," IEEE CVPR, 2014.
- [3] W. Liu et al., "SSD: Single Shot MultiBox Detector," ECCV, 2016.
- [4] T. Lin et al., "Microsoft COCO: Common Objects in Context," ECCV, 2014.
- [5] A. G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv, 2017.
- [6] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal of Software Tools, 2000.
- [7] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," IEEE CVPR, 2016.
- [8] P. Viola and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple Features," IEEE CVPR, 2001.
- [9] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," NIPS, 2012.
- [10] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," arXiv, 2018.
- [11] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection," IEEE TPAMI, 2017.
- [12] V. Sandler et al., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," IEEE CVPR, 2018.
- [13] M. Everingham et al., "The Pascal Visual Object Classes Challenge," IJCV, 2015.
- [14] D. Tuia et al., "Perspectives in Machine Learning for Wildlife Conservation," Nature Communications, 2022.

- [15] A. Norouzzadeh et al., "Automatically Identifying, Counting, and Describing Wild Animals in Camera-Trap Images with Deep Learning," PNAS, 2018.
- [16] World Bank, "Women Entrepreneurs and Digital Platforms for Business Growth," World Bank Research Report, 2021.
- [17] R. B. Freeman and R. Huang, "Collaborative Learning and Entrepreneurship Education," Harvard Business Review, 2018.
- [18] N. Agarwal and H. Kumar, "AI-Based Market Demand Prediction for Small Businesses," IEEE ICDSA, 2023.
- [19] A. K. Jain and B. Dubey, "Computer Vision Techniques for Image-Based Skill Evaluation," IJCV, 2022.
- [20] S. Russell and P. Norvig, "Artificial Intelligence: A Modern Approach," Pearson Education, 2021.