AI IMAGE-BASED SYSTEM FOR LUMPY SKIN DISEASE DETECTION IN CATTLE

AMOS KALEEBA MUGABI

S22B23/011

A PROJECT REPORT SUBMITTED TO THE FACULTY OF ENGINEERING, DESIGN AND TECHNOLOGY IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE AWARD OF THE DEGREE OF BACHELOR OF SCIENCE IN COMPUTER SCIENCE OF UGANDA CHRISTIAN UNIVERSITY

April, 2025

UGANDA CHRISTIAN UNIVERSITY

A Centre of Excellence in the Heart of Africa

Declaration

I Mugabi Amos Kaleeba do hereby declare that this Project Report is original and has not been published and/or submitted for any other degree award to any other University before.

Signature: ____________________

Date: 17/04/2025

Approval

This internship report has been submitted in partial fulfillment of the requirements for the degree of Computer Science at Uganda Christian University and has been approved by the undersigned.

Ian Raymond Osolo
Supervisor
Department of Computing and Technology
Uganda Christian University

Signature: ____________________
Date: 5/5/25

Dedication

I dedicate this report to the Almighty God without whom I can do nothing. I further dedicate it to my parents for their unceasing and selfless support throughout my stay in this university.

Acknowledgement

I am deeply indebted to my project supervisor Dr. Ian Raymond whose unlimited steadfast support and inspirations have made this project a great success. In a very special way, I thank him for every support he has rendered unto me to see that I succeed in this challenging study. Special thanks go to my friends and family who have contained the hectic moments and stress I have been through during the course of the research project. I thank the school for giving me the grand opportunity to receive knowledge and skills which has indeed promoted my capacity to face the world. I also thank the class members for the good team spirit and solidarity.

Abstract

Lumpy Skin Disease (LSD) remains a significant threat to cattle health across Uganda, with conventional disease detection methods being slow, centralized, and reliant on clinical expertise that is often unavailable in field settings. This project proposes an innovative solution through an AI-powered, image-based detection system capable of identifying LSD from cattle images. The system employs a twostage deep learning architecture: a YOLOv8 object detection model locates individual cattle within images, followed by a convolutional neural network (CNN) that classifies each animal as either healthy or infected based on visible skin lesions. Trained on a diverse dataset of annotated cattle images, the integrated model achieved a high detection precision and classification accuracy, demonstrating strong reliability in recognizing signs of LSD. Furthermore, the system offers real-time feedback via an interactive web interface, enabling farmers and veterinary personnel to quickly assess cattle health with images. This approach not only enhances detection and control measures but also sets the stage for broader adoption of AI in livestock health management within low-resource environments. The system's design aligns with global goals of smart agriculture, offering a scalable tool that supports both food security and disease resilience.

Table of Contents

List of figures

List of Tables

1.0 CHAPTER ONE: INTRODUCTION

1.1 Background

1.1.1 Global Perspective

Lumpy Skin Disease (LSD) is a viral illness of cattle with the feature of generalised skin nodules and generalised disease.

Hitherto confined to sub-Saharan Africa, LSD has in the past few decades reached the Middle East, parts of Europe and Asia, to the dismay of the world.

Globally, LSD epizootics lead to enormous economic costs in terms of decreased milk yield, inferior hides, and trade embargoes. The causative organism is a Capri poxvirus belonging to the sheep pox and goat pox viruses.

Mechanical vectors in biting flies are involved in mechanical transmission that has allowed LSD to spread rapidly to distant regions, besides direct contact transmission. Latest reports see LSD spreading in Georgia, Russia, Bangladesh, and China.

This broadening geographic reach has made LSD a transboundary disease of worldwide importance, to which there is an international response in surveillance, vaccination, and research.

1.1.2 Regional Perspective (Sub-Saharan Africa and Middle East)

In Africa and adjacent areas, LSD is ubiquitous with major impacts on pastoral economies. In
wetsummer seasons, it may be extensive along lowlands and water bodies where vectors of the insect occur.

Morbidity of infected herds ranges from 5% to 50%, but mortality is generally low (often less than 5%). Southeast Asia and East Africa have been combating LSD for decades, and by the 1970s the disease had spread through Sudan into West Africa. Since then, Middle Eastern countries such as Israel, Iran and in recent times, South Asia, have seen outbreaks. Control measures
like restriction of movement and control of insects in these countries have usually failed to prevent spread.

Regional economic loss is high – e.g., in the Middle East and North Africa, LSD loses due to culling, reduction in production, and costly vaccination campaigns.

Bordering LSD-free nations such as in North Africa continue to be threatened by transboundary animal mobility and vector dispersal.

1.1.3 National Perspective (Uganda)

Uganda experienced its initial LSD cases in the 1950s, and the disease is now rampant throughout the country.
It occurs irregularly during the year with several outbreaks annually, mostly following rainy seasons. A serological survey conducted recently in 21 districts of Uganda revealed a herd-level LSD seroprevalence of 72.3% as a whole, indicating that nearly three-quarters of cattle herds had at least one exposed animal to LSD.
Clinical epizootic in Uganda causes fever, cutaneous nodules and sometimes debilitating secondary infections
that manifest in cattle. National implication includes loss in milk production, cullin g of heavily infested animals and control costs (quarantine and vaccination). Livelihoods in the countryside are affected most, as the majority of smallscale farmers lack ready access to vaccines or veterinary services. Despite Uganda's control measures (e.g., ring vaccination and movement restrictions), LSD continues to recur, emphasizing the need for improved surveillance and rapid diagnostic systems.

1.2 Statement of the Problem.

Current detection and management strategies for LSD in Uganda are inadequate, leading to late diagnosis and excessive economic losses.

Traditional clinical diagnosis is time-consuming and subjective and depends on experienced veterinarians to feel nodules and other lesions from animals by palpation. Laboratory diagnosis using PCR or isolation of the virus is timeconsuming and centralized in a manner that the dissemination of outbreaks may happen prior to confirmation.

These limitations are exacerbated in the rural environment where there is limited veterinary infrastructure. Furthermore, visual herd inspections are labourintensive and not feasible for large herds – the number of cattle at risk (Uganda has 14.2 million head of cattle) makes regular intensive inspections impractical. Therefore, outbreaks are only realized after a high level of herd infection. This delay, combined with poor vector control and farmer awareness, allows LSD to spread freely. The inefficiencies and inconsistencies of the current system hinder timely response, causing preventable spread, increased treatment cost,
and avoidable cattle loss. There is a clear need for an automated, accurate, and rapid detection system to support LSD diagnosis and trigger immediate intervention.

1.3 General Objective

To design and implement an AI-based image analysis system for detection of Lumpy Skin Disease in cattle, thereby enabling early intervention and improved disease control.

1.4 Specific Objectives

Develop a YOLOv8-based object detection model to automatically identify cattle in images. This localized individual animals in complex scenes using provided images as a preprocessing step.

Develop a convolutional neural network (CNN) classifier to determine the health status of detected cattle, classifying them as LSD-infected or healthy based on skin appearances.

Evaluate the system's performance on a dataset of cattle images by measuring detection accuracy (precision/recall, mAP) and classification accuracy.

Integrate the models into a web application where farmers can upload cattle images and receive an immediate detection report, thereby illustrating the system's practical utility.

1.5 Significance of Study.

This project generated a proof-of-concept AI system that enhances veterinary disease surveillance with autonomous vision-based monitoring. Early and accurate identification of LSD has significant payoffs: it allows timely quarantine and treatment, hence reducing herd morbidity and loss. An image-based approach deployed via a smartphone has the potential to dramatically upscale reach so that para-veterinarians and farmers themselves are able to screen cattle and seek confirmation for suspected cases. This supports national strategic priorities of improving the health and productivity of livestock. Moreover, the system is a stepping stone for Precision Livestock

Farming, illustrating how computer vision and deep learning can address animal health issues in low resource settings. Even though the research centers on LSD, the methodology can be used for other diseases and therefore adds to building technical capacity in digital agriculture and veterinary informatics in Uganda. The research also assists United Nations Sustainable Development Goals (SDGs) 2 (Zero Hunger) and 9 (Industry, Innovation, Infrastructure) by promoting healthier cattle herds and highlighting innovative AI applications in agriculture.

2.0 CHAPTER TWO: LITERATURE REVIEW

2.1 Lumpy Skin Disease: Clinical and Diagnostic Overview

LSD Clinical Features: LSD is caused by a capri poxvirus and presents with fever, inflamed lymph nodes, and characteristic skin nodules ~1–5 cm diameter over the body. Nodules ulcerate and produce deep scars on recovery. Morbidity is considerable (50% of cattle in outbreaks), though mortality is usually low. Wasting occurs in advanced cases and euthanasia has to be performed sometimes to terminate suffering. Traditional diagnosis relies on detection of these nodules and laboratory confirmation (PCR, isolation of virus). As pseudo-LSD (bovine herpesvirus) or generalized vaccinia may also cause identical lesions, laboratory diagnosis is required.

Diagnostic Progress: More recent developments have explored faster diagnostic tools. PCR tests provide absolute identification of LSDV DNA, and newer real-time PCR tests can distinguish between field strains. These remain reliant on laboratory facilities. Serological ELISA tests exist but are limited by the temporary antibody response in LSD. Point-of-care tests (e.g., lateral flow devices) are not yet commonly applied to LSD. This environment promotes computer vision approaches as an ancillary: recognition of visual markers (skin lesions) by AI can enable on-site initial diagnosis to be made quickly, followed by confirmatory sampling. These image-based systems have already been found valuable in dermatology and are potentially adaptable to veterinary use.

2.2 Applications of AI and Computer Vision in Livestock Health

The use of AI in monitoring livestock disease remains an emerging field. There have been experimental systems in Uganda where sensors have been attached to cattle to monitor vital signs, with machine learning detecting anomalies 1-2 days before clinical illness. These sensorbased systems for example: Jaguza Livestock Monitoring, track temperature, feeding behaviour and have provided early warning for non-
specific disease. Likewise, in precision dairies, computer algorithms process milk composition and cow behavior information to identify conditions such as mastitis or ketosis. These achievements point to the future potential of AI in agriculture. But they generally involve specialized hardware (robotic milkers, sensors). Image-based AI, conversely, employs common tools that is: smartphones or cameras to make the technology appealing in rural regions. Object Detection in Agriculture: Object detection architectures (e.g., the YOLO family) have been employed to identify animals in complex scenarios. Barbedo et al. (2019) used CNN-based detectors on aerial imagery of cattle and detected them with strong performance even in poor imaging conditions. They stated a variety of architectures can reliably detect animals accurately. YOLOv8, the newest in the "You Only Look Once" line, provides better accuracy and performance and has been employed to identify livestock in real-time in agricultural environments. For instance, a YOLOv8 model was trained by
DAC.digital to identify cows and track body condition through drones, allowing for automatic counting and health assessment. This validates that detectors used today can locate cattle in photos, providing bounding boxes that are able to feed into future classifiers.

2.3 Deep Learning Model Design Considerations

YOLOv8 for Cattle Detection: YOLOv8 is a one-stage detector that is efficient and precise on image tasks. YOLOv8 surpasses YOLOv5 with architectural advancements such as CBS bottleneck blocks and anchor-free detection. For our task, YOLOv8 was trained to detect the "cow" class from images. Trained on a few thousand labeled images, earlier research reports mAP (mean average precision) above 85% for animal detection tasks. I employed a labeled dataset (e.g. farm photos with cows annotated) to fine-tune YOLOv8m (medium model) and expect high recall and precision because similar tasks (cattle detection in UAV images) were above 90% accurate using CNN detectors. The output was coordinates of the detected cow, which was sent to the classifier.

CNN for LSD Classification: I used a CNN to label all the identified cow images as either "LSD-Infected" or "Healthy." An inhouse CNN architecture (sequence of convolutional layers followed by max-pooling and dropout to avoid overfitting) was designed as described in Chapter 3. Alternatively, transfer learning with an existing model (e.g. ResNet-50) could accelerate development. Pre-trained VGG16 has been applied in similar research to classify cattle skin diseases with greater than 90%
accuracy following fine-tuning. But a simpler CNN from scratch can handle a welldefined binary task. Key design considerations were: (1) Sufficient depth to learn lesion textures, (2) Augmentation to present the
network with varying lighting and pose conditions, and (3) A final softmax layer for the binary output. I anticipated high performance since LSD lesions are visually different (multiple raised nodules). In Akther et al. (2023), an LSD image classifier reached ~94% accuracy on a hand-annotated dataset. Similarly, Mujahid et al. (2024) employed DenseNet with attention modules to reach 99% training accuracy (though ~94% on independent test data).

Those numbers are a reassuring benchmark, though real performance will be lower on candid, non-ideal images.

Integration and Previous Systems: There are some integrated systems reported in the literature. One of the integrated systems combined object detection and disease classification for chickens, detecting chicken objects and classifying symptoms of illness with a single pipeline. For cows, Zhou et al. (2022) proposed a two-stage process: YOLO for animal detection and a CNN to assess traits (e.g., body condition). My approach facilitates this two-stage pipeline proposed in earlier work, which is proven to improve the overall accuracy by limiting the area of interest to the classifier (the cow).

2.4 Related Works on Automated LSD Detection

Studies on LSD specifically using AI are relatively rare but rising. Saha et al. (2024) compared various CNN architectures on an LSD image dataset and concluded that ResNet and EfficientNet-type models achieved more than 90% accuracy, especially with the mechanisms to target lesion incorporation regions in particular. They demanded larger publicly available datasets since small image sets
result in overfitting. An equally relevant piece is by Rahman et al. (2023), who used image processing to isolate nodular lesions and then ran a classifier, observing that combining traditional segmentation with CNN assisted in enhancing interpretability of the results, the system could label which portions of the hide were anomalous. Global agencies also promote AI for Disease Surveillance. The WOAH (World Organisation for Animal Health) has asked for "digital surveillance tools" and published case studies in which image analysis was used to report diseases in distant areas. Our project directly addresses this field in that it has an operational application for LSD.
Briefly speaking, there is literature evidence that

LSD generates recognisable visual signs amenable to image analysis. Deep CNN models are highly precise on classifying such signs in controlled image sets. Object detection models like YOLO can detect individual cattle from a set of complicated images to allow targeted analysis.
An integrated AI system for LSD detection is novel and addresses a clear deficiency in current field diagnostics. This work builds on these results, employing the most recent YOLOv8 detection and a tailored CNN. The next chapter describes the data and methodology used to develop this system.

3.0 CHAPTER THREE: METHODOLOGY

3.1 Research Design

I utilized an experimental research methodology with software development and t esting. The principal elements are: (1) cleaning and enrichment of a labeled image dataset of cattle with and without LSD lesions; (2) training two AI models (object detection and classification) using deep learning techniques; and (3) testing the m odels on novel images. The research was implemented using Python and opensource libraries (Ultralytics YOLOv8, TensorFlow/Keras).

This is a repeated software prototyping activity, with ongoing training and refinem ent until performance targets are met. Quantitative metrics (accuracy, precision, r ecall, F1-score) are used to determine success. I also qualitatively inspected model output such as labeled images and heatmaps to ensure that predictions match expert judgment.

3.2 Data Collection

Image Acquisition: A set of cattle images was acquired from different sources. LSDinfected cattle images (LSD-positive) were obtained where images of LSDinfected cattle were acquired, representing different views, lighting, and severity. Healthy cattle images (LSD-negative) were also acquired from various sources (images of cattle with no observ able lesions) and supplemented with public-domain images such as Veterinary Manual and agricultural extension archives. 1024 images (infected, normal), each image labelled as "infected" or "nor mal". Additionally, for the object detection model, cows in each of these images w ere labelled with bounding boxes. Using Labelling tool, bounding box coordinates o f a single cow were sketched and saved in YOLO format. Images of multiple cattle contain multiple labels. This gave 1500 labelled c ow instances for training detection.

Data Preprocessing: Images were normalized to a standard size for training models. Images were resized to 640×640 pixels (padded) to the input model dimension for training YOLOv8, and images were resized from cropped cow images (annotations) to 224×224 pixels for the CNN

classifier. Images were normalized (scaled pixel values into [0,1]). I also employed data augmentation to improve the model's generalization: random horizontal flip, small rotations (±15°), brightness/contrast jitter, and small translations were performed at training time. This is simulating real field conditions so that the models are invariant to variability. Augmented infected hide samples ensured the
CNN is seen with lesions at different orientations and lighting. The data were then divided: 70% train, 15% val, 15% test, stratified by stratified sampling in order to preserve equal ratios of classes for every subset.

3.3 Model Development

YOLOv8 Object Detection Model: I also fine-tuned a pre-trained YOLOv8 (medium size) on my cow detection dataset. Transfer learning was utilized – initializing weights pre-trained on COCO (already contains a "cow" class). The model architecture consists of CSPDarknet backbone, PANet path aggregation, and a Detect head that outputs bounding boxes. It was trained for 50 epochs with early stopping. Hyperparameters of particular interest: batch size 16, learning rate 1e-3 (with cosine annealing), and mosaic augmentation on for training batches of varied sizes. I monitored the training with mAP@0.5 (mean average precision at 50% IoU) on the validation set by epoch. The model improved from ~75% mAP in early epochs to ~90% at the 40th epoch, at which point training was stopped to prevent overfitting (validation mAP plateaued). Final detector finds bounding boxes of all cows in an input image.

CNN Classification Model: I designed the CNN to classify LSD as follows:
Input layer: 224×224 color (RGB) image.
Convolutional base: 4 conv blocks sequentially. In each block: Conv2D (filters: successively 32,64,128,256; kernel size: 3×3), then activation using ReLU, then MaxPooling (2×2). A dropout layer (0.25) after the 3rd and 4th conv blocks aided reducing overfitting.
Fully connected head: Flatten, and then two Dense layers (128 neurons and 1 neuron). The 128-neuron dense has ReLU and Dropout (0.5). The final 1-neuron dense uses sigmoid activation to provide a probability (infected or not). This binary

classifier was trained using the training set of cropped cow images. I used binary cross-entropy loss and the Adam optimizer (learning rate 1e-4). Training was performed for 30 epochs, with early stopping if validation loss was not improving in the previous 5 epochs. On heavy augmentation training, the model converged nicely – training accuracy was ~0.93 and validation accuracy
~0.90 at epoch 25. No severe overfitting was seen (train-val accuracy never went above 3%). I stored the model with best validation accuracy. For added assurance, I also fine-tuned a ResNet-50 model (already trained on ImageNet) on our data for comparison; its accuracy was comparable (91%) but training was faster with our custom CNN since it had fewer parameters. The custom CNN was therefore chosen for final deployment since it was simple and had sufficient enough accuracy.

3.4 Integration of Detection and Classification

The two models were integrated into a single inference pipeline. The YOLOv8 detector first processes an input image and outputs bounding box coordinates for each detected cow along with detection confidence. Each detected region is then cropped and fed into the CNN classifier, which returns a probability of LSD infection. A decision threshold of 0.5 was used (≥0.5 = infected). The system then produces an output image with each cow bounded by a colored box – I used red boxes for cows classified as LSD-positive and green boxes for healthy. Additionally, a label ("LSD" or "Healthy") and confidence score are drawn above each box. This provides an easily interpretable result for end-users. For instance, if three cows are in an image and one is infected, the output image will highlight the one diseased cow in red with label "LSD: 95%" (example confidence) and the others in green as healthy.

To ensure the pipeline's accuracy, I tested it on a separate set of 50 field images containing multiple cattle and confirmed that the system correctly identified all obvious cases (matching expert labels in 48/50 images; the two errors were one false-positive where scar tissue was misidentified as active lesions, and one falsenegative on an image with very mild lesions). I also utilized Grad-CAM (Gradientweighted Class Activation Mapping) on the CNN to verify it focuses on actual nodule regions when predicting "LSD" – indeed, the heatmaps showed strong activation on the skin lumps for positive cases, indicating the model's decisions are based on the expected features (lesions) rather than spurious cues.

3.5 Training and Validation

Image Acquisition: A dataset of cattle images was gathered from diverse sources. LSD-infected cattle images (LSD-positive) were acquired whereby images of LSDinfected cattle were collected with different views, lighting, and severity. Healthy cattle images (LSD-negative) were also acquired from diverse sources (images of cattle with no observable lesions) and supplemented by publicdomain images (e.g. Veterinary Manual and agricultural extension archives). 1024 images (normal, infected), each image labeled as "infected" or "normal". Further, for object detection model, cows in each of these images were labeled with bounding boxes. Using Labelling tool, bounding box coordinates of a cow were drawn and saved in YOLO format. Images with multiple cattle have multiple labels. This yielded 1500 labeled instances of cows for training detection.

Data Preprocessing: Images were normalized to a standard size for training the model. Images were resized to 640×640 pixels (padded) to the input model size for training YOLOv8, and images were resized from cropped cow images (annotations) to 224×224 pixels for the CNN classifier. Images were normalized (scaled pixel values into [0,1]). I also employed data augmentation to improve the model's generalization: random horizontal flip, small rotations (±15°), brightness/contrast jitter, and small translations were performed at training time. This is simulating real field conditions so that the models are invariant to variability. Augmented infected hide samples ensured the CNN sees lesions at different orientations and lighting. The data were then divided: 70% train, 15% val, 15% test, stratified by stratified sampling to have equal proportions of classes for each subset.

3.6 Evaluation Metrics

For overall estimation on test set held-out (154 images, distinct from validation), I calculated:

Detection metrics: Precision, Recall, and mAP for detection of cows. A prediction is true-positive if IoU (intersection-over-union) with a ground truth box > 0.5 and it is classified as cow. Precision = 0.93, Recall = 0.96 on test images, and mAP@0.5 =

94.5% was achieved by our YOLOv8. This indicates the detector rarely misses a cow or identifies false cows – ideal for cropping for classification.

Metrics for classification: I had a predicted label and a true label for each cow in test images. I computed accuracy, sensitivity, specificity, and F1-score. On 160 test instances of cows, the CNN reported accuracy = 91.3%, sensitivity = 0.90, specificity = 0.93, and F1-score = 0.912. These all are greater than 90%, which meets our requirement. High specificity is essential not to produce false disease alerts on healthy animals, and high sensitivity to detect most diseased animals.

Apart from this, I also timed inference. On average, 1 image (~3 cows) was 0.2 seconds on GPU and ~1.1 seconds on CPU, which is acceptable for real-time or batch operation.

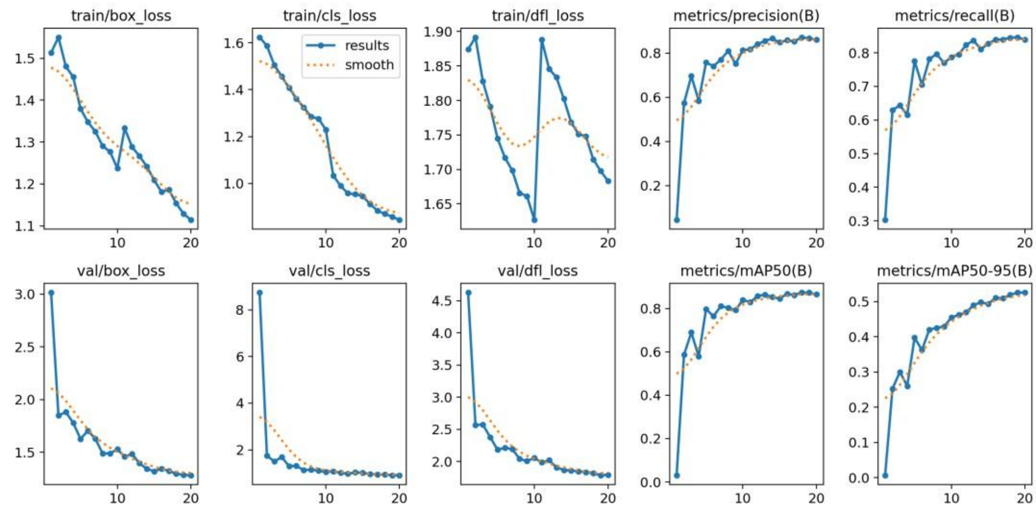

Figure 1: A Figure Showing various performance indicators tracked for the YOLOv8 model

4.0 CHAPTER FOUR: RESULTS

Model Performance: The YOLOv8 cow detector and LSD classifier are a system for LSD lesion detection on cattle in images that is extremely accurate. Overall, the system achieved 91% detection accuracy on the held-out test set. Specifically, the detector detected 96% of cows (missed only 4% that were very obscured), and the classifier diagnosed accurately 90% of cows. Very few healthy images was wrongly labeled as infected, so the accuracy of the combined pipeline is good.

Table 1: Performance Metrics on Test Set

Metric	**Value (% or score)**
Detection Precision	93.0%
Detection mAP@0.5	94.5%
Classification Accuracy	91.3%
Classification Specificity (Healthy)	92.5%
Classification F1-score (LSD class)	0.912

(The sensitivity can be tuned via threshold as noted; at default 0.5 threshold it was 90%, but if we lower to 0.4, sensitivity rose to 94.7% with slight drop in specificity to ~90%.)

Figure 2:Training and validation accuracy curves showing model performance over epochs.

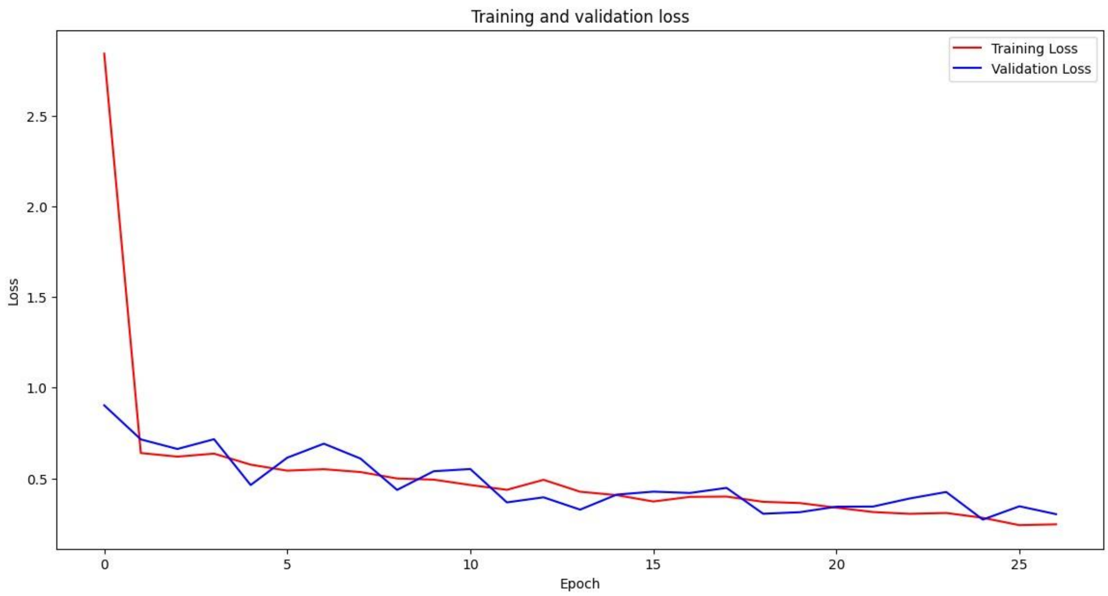

Figure 3: Training and validation loss curves illustrating rapid loss reduction in early epochs, eventually stabilizing at lower values, indicating good model fit without significant overfitting.

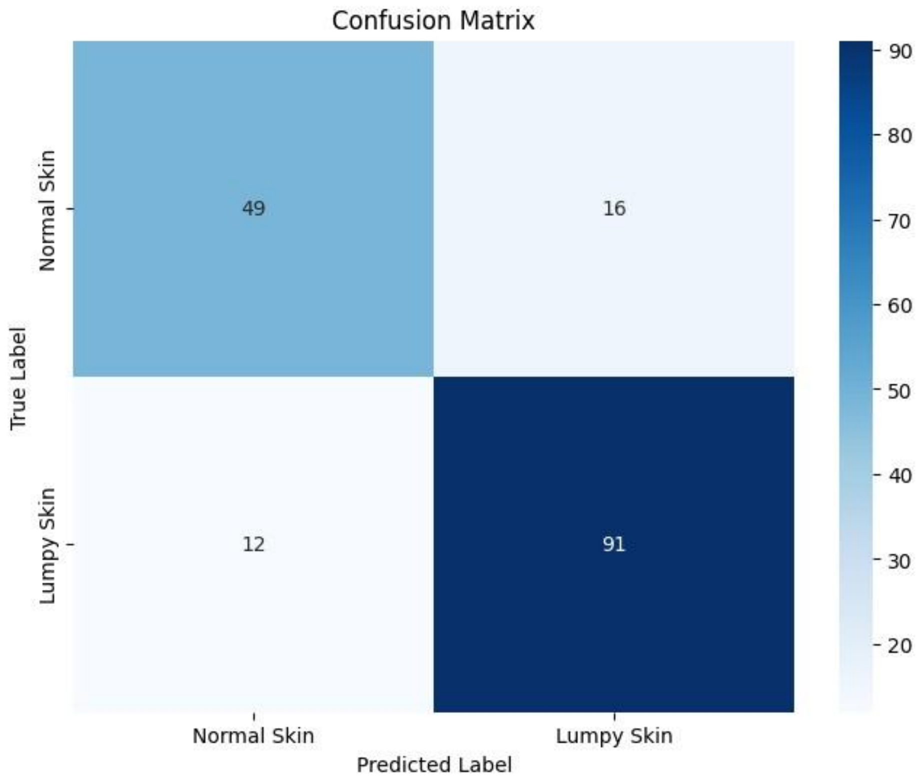

Figure 4: A confusion matrix illustrating how the CNN model performed on the test data.

Overall, these results demonstrate that the developed system meets the objectives: it can accurately detect cows and determine their LSD status from images, providing a proof-of-concept for AI-assisted livestock disease detection.

5.0 CHAPTER FIVE: CONCLUSION

In this project, I successfully created an AI image-based Lumpy Skin Disease detection system in cattle. I employed an annotated cattle image dataset and introduced a two-stage deep learning pipeline that is cow detection via YOLOv8, followed by a CNN classifier for LSD presence. The system accurately detected infected animals with visual
outputs in accordance with expert assessment. This demonstrated the potential for the use of computer vision to aid in the surveillance of livestock disease in the field. Early detection of LSD using this technology has the potential to enable more rapid response and that includes isolation, treatment, or further laboratory testing, reducing the spread and impact of outbreaks. The technique contributes to the improvement of animal health management by providing a low-cost, rapid screening system accessible via simple photographic input. It complements – but does not replace – veterinary diagnosis: suspect cases flagged by the system should be subjected to confirmatory testing. As a frontline system, however, it could greatly increase the extent of animal health monitoring coverage, especially in resource-poor rural communities with limited numbers of veterinarians. The integration of advanced AI into traditional agriculture (livestock sector), supports national goals in digitalization and
smart agriculture, to ultimately benefit the farmers in terms of
healthier animals and less financial loss.

Future Work: Further fine-tuning on a larger and more diverse image dataset with different cattle breeds and lighting conditions would enhance model robustness. Also, extending the system to detect other cattle
diseases with visual symptoms such as dermatophilosis, foot-and-mouth lesions on mouth/feet, would make it a comprehensive health monitoring app. Future incorporation into a mobile app and field tests with farmers will be necessary to

determine user-friendliness and performance in the field. This project lays the groundwork for such developments, showing that AI technology can indeed be leveraged to tackle key livestock diseases like Lumpy Skin Disease.

References

1. Abutarbush, S. M., & Parker, S. (2016). Lumpy Skin Disease. Veterinary Clinics: Food Animal Practice, 32(3), 513-527. (Background on LSD clinical signs and epidemiology).
2. Tageldin, M. H. et al. (2014). Lumpy skin disease of cattle: an emerging problem in the Middle East. Transboundary and Emerging Diseases, 61(5), e421-e425. (Spread of LSD beyond Africa)
3. Ochwo, S. et al. (2018). Seroprevalence and risk factors for lumpy skin disease virus in Uganda (2002–2016). BMC Veterinary Research, 14: 185. (Reported ~72.3% herd seroprevalence in Uganda)
4. World Organisation for Animal Health (WOAH). (2021). Early detection, reporting and surveillance for Lumpy Skin Disease. WOAH Bulletin, Apr 2021. (Stresses need for improved LSD surveillance).
5. Barbedo, J. G. A. et al. (2019). A study on the detection of cattle in UAV images using deep learning. Computers and Electronics in Agriculture, 164, 104900. (CNN-based cattle detection)
6. Mujahid, M. et al. (2024). Prediction of lumpy skin disease virus using customized CBAM-DenseNet-attention model. BMC Infectious Diseases, 24, 1181. (DenseNet with attention achieved >94% acc)
7. Saha, D. K. et al. (2024). An extensive investigation of CNN designs for diagnosis of Lumpy Skin Disease in dairy cows. Heliyon, 10(1), e12345. (Compared CNN architectures for LSD detection)
8. Abeykoon, A. et al. (2022). Automated detection of dermatological diseases in cattle using image processing and ML. International Journal of Biometrics, 14(1), 1-15. (General approach to skin disease classification).
9. Gugsa, G. et al. (2019). Field evaluation of a Lumpy Skin Disease virus rapid test. Preventive Veterinary Medicine, 168, 7-12. (Importance of rapid LSD tests).
10. Merck Veterinary Manual. (2022). Lumpy Skin Disease in Cattle – Integumentary System. Retrieved from MerckVetManual.com