Vol. 1, Issue 11 | Nov 2026 | Open Access — Research Paper

AI-BASED COW AND BUFFALO SKIN DISEASE DETECTION USING YOLOv11s

A Deep Learning-Based Object Detection Approach

Vaishnavi Busa, Deekshitha Konduru, Maturi Jairam Reddy, and Shubani Gulam (Guide)

Department of CSE(AI&ML), CMR Technical Campus, Hyderabad, Telangana, India

***Abstract*—**Skin conditions in cattle and buffaloes can affect animal welfare and farm productivity. This paper presents a proposed image-based decision-support system that uses the YOLOv11s object-detection model to localize visible skin abnormalities in cow and buffalo images. The planned workflow includes dataset acquisition, veterinary-informed class definition, image curation, bounding-box annotation, data partitioning, model fine-tuning, quantitative evaluation, and deployment through a web application. Since dataset collection and model training have not yet been completed, this manuscript describes the research design rather than reporting experimental findings. Performance will be measured using precision, recall, mean average precision at IoU 0.50 (mAP@0.50), mAP averaged over IoU thresholds 0.50–0.95, and inference latency. The system is intended only to support preliminary visual screening; it is not a veterinary diagnostic substitute. The final paper will report the verified dataset composition, experimental settings, and measured results after project implementation.

***Index Terms*—**Animal health, buffalo, cattle, computer vision, deep learning, object detection, skin abnormality, YOLOv11s.

I. INTRODUCTION

Livestock production is an important component of food systems and rural livelihoods. Cattle and buffaloes are economically important animals, and visible skin conditions may cause discomfort, secondary complications, reduced productivity, and additional veterinary costs. Timely recognition of suspicious skin changes can help farmers seek professional assessment. However, visual inspection in the field may be affected by lighting, animal movement, coat colour, image quality, lesion size, and the observer's experience.

Computer vision offers tools for analysing images and locating objects or regions of interest. Object detection models predict both a class label and a bounding box around a target. The YOLO family provides a one-stage detection approach, and Ultralytics documents YOLO11 variants for detection, training, validation, prediction, and export [1]. This project selects the small YOLOv11s variant as the initial model because the project aims to balance detection quality and computational requirements. Suitability must nevertheless be established on the project's own dataset.

The proposed work will develop a web-based application in which a user uploads a cow or buffalo image and receives model-generated candidate detections, class labels, confidence scores, and bounding boxes. The output is intended to support preliminary screening and must be interpreted by a qualified veterinarian.

II. PROBLEM STATEMENT

There is a need for an accessible image-analysis workflow that can assist in locating visible skin abnormalities in cow and buffalo photographs. A practical system must account for multiple possible conditions, species differences, class imbalance, variable image capture conditions, and uncertainty in visual appearance. The research problem is to design and evaluate a YOLOv11s-based object detector and a web interface using a curated, labelled, and appropriately separated dataset.

III. RESEARCH OBJECTIVES

- Develop a curated image dataset covering cow and buffalo skin conditions, with provenance and usage rights recorded.
- Define a manageable set of visually distinguishable classes in consultation with a veterinary subject-matter expert.
- Annotate visible target regions with bounding boxes using a consistent annotation protocol.
- Fine-tune a pretrained YOLOv11s detector on the training set.
- Evaluate detection performance on held-out data using precision, recall, mAP@0.50, mAP@0.50:0.95, and class-wise analysis.
- Build a web application for image upload and display of candidate detections.
- Document limitations and communicate that outputs are not a confirmed veterinary diagnosis.

IV. SCOPE AND ASSUMPTIONS

The initial scope is image-based detection of visible skin findings in cows and buffaloes. The project will not infer disease severity, prescribe treatment, or replace clinical examination. A photograph may not contain enough evidence to distinguish conditions with similar visual appearance. Labels should therefore be assigned or reviewed by a veterinarian or another appropriately qualified animal-health expert. If reliable labels cannot be obtained, the class must be described as a visible finding (for example, 'suspected lesion region') rather than a confirmed disease.

The final class list is not fixed in this draft. Candidate classes must be chosen only after dataset inspection and

expert review. A class should not be included merely because a small number of web images can be found for it.

V. RELATED WORK

Deep-learning methods have been investigated for animal-disease image analysis, including classification of cattle disease images. A recent study comparing pretrained convolutional and transformer encoders for foot-and-mouth disease and lumpy skin disease in cattle illustrates the relevance of curated image data, class definitions, and independent evaluation [2]. Its task and model design should not be treated as identical to the present object-detection task; the present work focuses on localizing visible regions with bounding boxes.

The Ultralytics YOLO11 documentation describes supported detection workflows and provides the recommended software citation for YOLO11 [1]. The documentation is a technical source, not a peer-reviewed YOLO11 architecture paper. The authors should add further peer-reviewed work on cattle/buffalo dermatological image datasets, livestock object detection, and veterinary image annotation after completing a structured literature search.

Research gap to be tested: whether a single detector can provide useful localization of veterinary-reviewed visible skin findings across both cow and buffalo images under varied capture conditions. This is a proposed research question, not a confirmed novelty claim. A systematic literature review is required before claiming that the approach is the first or unique.

VI. PROPOSED METHODOLOGY

A. Dataset Acquisition and Governance

Images will be obtained from permitted public datasets, research repositories, and/or consented field collection. For every source, record the repository or contributor, URL or identifier, access date, licence, permitted use, species, label information, and any restrictions. Do not scrape or republish images where the licence or permission is unclear. Remove personal information and avoid retaining owner identifiers or unnecessary location metadata.

The dataset should represent both species. Record the number of images per species and class, image resolution, capture conditions, and source. Images from the same animal, farm, video sequence, or near-duplicate burst must remain in only one partition to reduce data leakage. The train/validation/test split should be grouped by animal or source where feasible.

B. Class Definition and Veterinary Review

Candidate labels may include veterinarian-confirmed conditions such as lumpy skin disease or dermatophytosis only when the image and supporting clinical information justify the label. A generic 'other/uncertain' category may be used during curation but should not be merged indiscriminately with confirmed disease classes. The final class taxonomy, inclusion criteria, exclusion criteria, and adjudication procedure must be documented before annotation.

C. Image Annotation

Each target visible region will be enclosed by an axis-aligned bounding box and assigned one approved class. Annotators will follow a written guide defining box tightness, treatment of multiple lesions, partially visible findings, occlusion, and low-quality images. A subset of images should be independently reviewed by a second annotator or veterinary expert. Disagreements should be resolved and the annotation protocol updated before full-scale labelling.

For YOLO-format labels, each object is stored as: class_id, x_center, y_center, width, height, with coordinates normalized to the image dimensions. The dataset configuration file should list the training, validation, and test image directories and the class names.

D. Preprocessing and Dataset Partitioning

Images will be checked for corruption, duplicates, extreme blur, irrelevant content, and insufficient visibility. Resize/letterbox operations will be handled consistently by the training pipeline. Augmentation may include horizontal flips, modest scaling, translation, brightness/contrast changes, and mosaic augmentation when appropriate. Augmentation must not alter the clinical meaning of the image or be applied to validation/test data.

A proposed initial split is 70% training, 15% validation, and 15% test, subject to dataset size and grouped splitting constraints. The final paper must report exact counts, not just percentages.

E. YOLOv11s Model Training

A pretrained YOLOv11s detection model will be fine-tuned on the annotated training set. The initial experiment may use an image size of 640 pixels and a maximum of 100 epochs, with early stopping or model selection based on validation performance. Batch size, learning rate, optimizer, augmentation settings, hardware, software versions, and random seed must be recorded. These are starting settings and should be adjusted to the available dataset and GPU.

Training should be repeated with fixed seeds where practical. The best checkpoint should be selected using validation data only. The test set must be evaluated once after model selection and must not be used to tune hyperparameters.

F. Evaluation Metrics

Precision = TP/(TP+FP), and Recall = TP/(TP+FN), where TP, FP, and FN are true positives, false positives, and false negatives under the selected matching rule. Average precision (AP) is the area under the precision-recall curve for a class. mAP@0.50 averages AP across classes at IoU 0.50; mAP@0.50:0.95 averages across IoU thresholds from 0.50 to 0.95. Report per-class metrics as well as aggregate values. Inference latency should state hardware, image size, batch size, and whether preprocessing/postprocessing are included.

Because images may contain multiple lesions and class frequencies may be imbalanced, report class-wise precision, recall, AP, confusion matrix, and representative false-positive/false-negative examples. Where possible, include confidence intervals or repeated-run variability.

G. Web Application Development

The web interface is planned using Python and Streamlit. A user selects or uploads an image; the backend validates file type and size, loads the trained YOLOv11s weights, performs inference, and returns an annotated image with class name, confidence score, and bounding box. The interface should display a clear caution that results are preliminary and require veterinary assessment. Uploaded images should not be retained longer than necessary, and access controls should be considered if the application is deployed beyond a local demonstration.

VII. SYSTEM ARCHITECTURE

Fig. 1 illustrates the proposed end-to-end workflow. The diagram is conceptual and should be updated to match the implemented system.

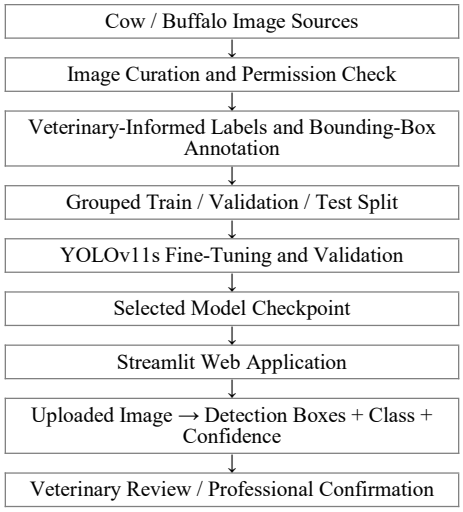


Fig. 1. Proposed system workflow for image-based skin-abnormality detection.

VIII. ALGORITHM

1. Input: labelled training images, validation images, and an unseen test set.
2. Load pretrained YOLOv11s detection weights.
3. Read the dataset configuration and approved class names.
4. Train the detector using the training set and monitor validation metrics.
5. Select the checkpoint using validation performance only.
6. Evaluate the selected checkpoint on the held-out test set.
7. For a web upload, validate the image, run inference, and render predicted boxes and confidence scores.
8. Display a veterinary-review notice and avoid presenting predictions as confirmed diagnoses.

IX. IMPLEMENTATION PLAN

The proposed software stack is Python, PyTorch, Ultralytics YOLO11, OpenCV, and Streamlit. The exact versions should be recorded in the final manuscript. The following code is a starter training template; it will run only after a correctly annotated dataset and data.yaml file have been prepared.

```
from ultralytics import YOLO

# Load pretrained YOLOv11s detection weights
model = YOLO('yolo11s.pt')

# Fine-tune using the project dataset configuration
model.train(
    data='data.yaml',
    epochs=100,
    imgsz=640,
    batch=16,
    patience=20
)

# Validate the selected/best checkpoint
best_model = YOLO('runs/detect/train/weights/best.pt')
metrics = best_model.val(data='data.yaml')

# Run inference on one test image
results = best_model.predict(source='sample test image.jpg',
conf=0.25)
```

X. EXPERIMENTAL RESULTS

No experimental results are available at the time of preparing this draft because the dataset has not yet been collected and the model has not yet been trained. The following tables are placeholders and must be completed using the actual outputs from the final experiments.

TABLE I
DATASET DISTRIBUTION BY CLASS AND PARTITION

Class	Train images	Validation images	Test images
[Class 1]	[N]	[N]	[N]
[Class 2]	[N]	[N]	[N]
[Class 3]	[N]	[N]	[N]
[Add classes]	[N]	[N]	[N]

TABLE II
MODEL EVALUATION RESULTS (COMPLETE AFTER TESTING)

Metric	Validation	Test
Precision	[Pending]	[Pending]
Recall	[Pending]	[Pending]
mAP@0.50	[Pending]	[Pending]
mAP@0.50:0.95	[Pending]	[Pending]
Inference latency (ms/image)	[Pending]	[Pending]

The final results section should include: (1) training/validation loss curves, (2) precision-recall curves, (3) confusion matrix, (4) sample detections for each class and species, and (5) a short error analysis. Do not copy benchmark values for YOLO11s on unrelated datasets as results for this project.

XI. DISCUSSION

The discussion will interpret the measured results after experiments are completed. It should consider whether performance differs between cows and buffaloes, whether minority classes are under-detected, and whether lighting, distance, occlusion, coat colour, camera type, or background changes affect predictions. Any claims about generalization must be supported by an independent test set representative of the intended use environment.

XII. LIMITATIONS AND ETHICAL CONSIDERATIONS

- Image-based predictions may be unreliable when lesions are small, hidden, blurred, or visually similar to other conditions.
- Public datasets may have uncertain labels, duplicated images, or limited representation of buffaloes and local field conditions.
- A model trained on one region, breed, camera, or farm may not generalize to other settings.
- Veterinary confirmation is necessary; the application must not recommend medication or treatment from an image prediction.
- Use only images with appropriate licences, permissions, and consent where required.
- Report model failures and dataset limitations transparently.

XIII. CONCLUSION AND FUTURE WORK

This manuscript presents the research design for a YOLOv11s-based system intended to localize visible skin abnormalities in cow and buffalo images and present the detections through a web application. The work plan covers dataset curation, expert-informed labelling, model training, held-out evaluation, and application development. Since implementation is ongoing, no claim of achieved accuracy or clinical effectiveness is made. Future work may include collecting a larger and more balanced multi-species dataset, external validation across farms and regions, comparison with alternative detectors, and prospective evaluation with veterinary experts.

ACKNOWLEDGMENT

[Add acknowledgment only after confirming the names of the project guide, department, institutional support, dataset contributors, and any funding. Obtain permission before naming individuals.]

REFERENCES

[1] Ultralytics, "YOLO11," Ultralytics Documentation. [Online]. Available: https://docs.ultralytics.com/models/yolo11/. Accessed: Nov. 14, 2026. (Technical documentation; use the software citation recommended by the project's installed version.)

[2] P. Valerio and W. Ticona, "Classification of Foot-and-Mouth and Lumpy Skin Disease in Cattle Using Frozen Pretrained Encoders: A Comparison of Convolutional and Transformer Architectures," Veterinary Sciences, vol. 13, no. 9, Art. no. 962, Sep. 2026, doi: 10.3390/vetsci13090962.

[3] IEEE, "IEEE Author Guide: Developing Your Manuscript." [Online]. Available: https://eps.ieee.org/wp-content/uploads/2025/11/mc-ieee_author_guide_interactive-new.pdf. Accessed: Nov. 14, 2026.

[4] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788, doi: 10.1109/CVPR.2016.91.

[5] [Add a veterinary reference defining the selected disease classes and diagnostic criteria, reviewed by the project guide or veterinary expert.]

APPENDIX A: DATASET CONFIGURATION TEMPLATE

```
path: /path/to/project_dataset
train: images/train
val: images/val
test: images/test

names:
  0: [class name 0]
  1: [class name 1]
  2: [class_name_2]
```

APPENDIX B: ITEMS TO COMPLETE BEFORE SUBMISSION

- Confirm author names and order, department name, institution's official name, and approved email addresses.
- Finalize disease classes with veterinary review and document label definitions.
- Collect images lawfully and record source, licence, species, and class counts.
- Complete annotations and grouped train/validation/test split.
- Train and evaluate YOLOv11s; replace every placeholder with measured values.
- Insert actual training graphs, confusion matrix, and representative test detections.
- Complete a literature review and verify every bibliographic field and DOI.
- Check the target journal/conference template, scope, peer-review policy, fees, and submission rules.
- Obtain project guide approval and ensure all authors approve the final manuscript.

© IJVAST — International Journal of Vast Advancements in Science and Technology.
ISSN: 2839-5379 | DOI: 10.68296/IJVAST20260099 | Open Access — CC BY 4.0 License