A Comparative Benchmark of Deep Learning Models and Deployment of a Web Application for Automated Detection of Lumpy Skin Disease in Cattle

Ms. S. Uma Shanthi[1], Dr. K. Padma Priya[2]
[1]PG Scholar, Department of Communication Systems, Government College of Engineering, Tirunelveli, Tamil Nadu, India
[2]Assistant Professor, Department of Electronics and Communication, Government College of Engineering, Tirunelveli, Tamil Nadu, India
doi.org/10.64643/IJIRTV12I12-202297-459

***Abstract*—A viral illness called Lumpy Skin Disease (LSD) has a major impact on cow herds all over the world and causes considerable financial losses for the livestock industry. Controlling the disease's spread and lessening its effects require an early and precise diagnosis. Three Convolutional Neural Network (CNN) architectures—VGG19, MobileNetV2, and MobileNetV3—are compared in this study for automatic LSD identification from cattle skin photos. The best model is then implemented as a web-based diagnostic tool. To enhance model generalization, a dataset of 717 cow skin photos from two classes (Lumpy Skin and Normal Skin) was gathered, preprocessed, and enhanced. To improve classification performance, transfer learning was used with pre-trained ImageNet weights. Performance criteria such as accuracy, precision, recall, F1-score, and computing efficiency were used to assess the models. According to experimental data, VGG19 and MobileNetV3 obtained the maximum validation accuracy of 97.18%, with corresponding F1-scores of 0.97 and 0.96. With a model size of 51.21 MB and an inference time of 134 ms—roughly 7.6 times faster than VGG19—MobileNetV3, on the other hand, showed noticeably higher computational efficiency. MobileNetV3 was chosen for deployment because of its balanced accuracy and efficiency performance. A Gradio-based online application that offers users actionable recommendations, confidence scores, and real-time predictions was incorporated with the trained model. The suggested system successfully closes the gap between cutting-edge deep learning research and useful veterinary applications, facilitating early disease management in the cattle sector and enabling quick preliminary screening of LSD in settings with limited resources.**

***Index Terms*—Lumpy Skin Disease, Deep Learning, Convolutional Neural Networks, MobileNetV3, Transfer Learning, Veterinary Diagnostics, Web Deployment, Gradio**

I. INTRODUCTION

The Lumpy Skin Disease Virus (LSDV), a member of the genus Capripoxvirus of the family Poxviridae, is the virus that causes Lumpy Skin Disease (LSD) [1]. Firm, spherical skin nodules covering vast body areas, fever, swollen lymph nodes, decreased milk output, and impaired fertility are the disease's main symptoms in cattle. Despite the disease's generally low (1–5%) death rates, it causes significant economic losses due to reduced output, hidden harm, trade restrictions, and control measure expenses [2].

Since its discovery in Zambia in 1929, LSD has quickly expanded throughout Africa, the Middle East, Europe, and Asia [3]. India confirmed its first LSD epidemic in Odisha in 2019. Millions of cattle were impacted by a catastrophic countrywide outbreak in 2022. India's agricultural industry is particularly susceptible to the effects of LSD because it has the largest bovine population in the world (more than 300 million) and leads the world in milk production [4].

Clinical examination and laboratory testing, including as PCR, viral isolation, and serological assays, are the mainstays of traditional LSD diagnosis. Subjectivity in visual diagnosis, lengthy laboratory processes (2–14 days), the need for specialized equipment and trained personnel, restricted accessibility in rural locations,

and delayed discovery after disease spread are some of these methods' major drawbacks [5].

In the fields of medicine and agriculture, recent developments in deep learning, especially Convolutional Neural Networks (CNNs), have transformed image-based illness diagnosis [6]. CNNs can quickly and accurately classify images by autonomously learning hierarchical feature representations. By utilizing information from models that have already been trained on huge datasets like ImageNet, transfer learning improves performance even further [7].

CNN-based LSD detection has been investigated in a number of studies. A customized CNN created by Patel and Sharma [8] achieved 96.7% accuracy on 15,000 clinical pictures. A hybrid Vision Transformer model with 98.1% accuracy was used by Zhang et al. [9]. 95% accuracy was attained by Saqib et al. [10] using MobileNetV2 with gradual fine-tuning. After comparing 12 pre-trained models, Senthilkumar et al. [11] discovered that MobileNetV2 was the best model for balancing accuracy and efficiency.

However, there is still a clear research gap between creating high-accuracy models and making them available to end users such as field veterinarians and farmers. Without incorporating models into scalable, user-friendly apps, the majority of research finish with performance metrics [12].

In order to close this gap, this study: (1) systematically compares VGG19, MobileNetV2, and MobileNetV3 for LSD detection; (2) determines the best model based on accuracy-efficiency trade-off; and (3) uses the Gradio framework to deploy the selected model as a functional web application, enabling quick preliminary diagnosis in real-world, resource-constrained settings.

II. METHODOLOGY

To aid in the construction of the automatic identification model, a carefully selected dataset of photos of cattle skin was gathered from scholarly publications and publically available repositories. The 717 photographs in the collection are divided into two classes: Class 1 (Normal Skin), which comprises pictures of healthy cattle skin free of infection, and Class 0 (Lumpy Skin), which includes pictures showing obvious clinical symptoms of Lumpy Skin Disease (LSD). The dataset was split into training and validation subsets using stratified sampling to maintain the original class distribution in order to guarantee trustworthy model training and assessment. In particular, 71 photos (10%) were set aside for validation and performance evaluation, and 646 images (90%) were used for model training.

To meet the input requirements of the CNN architectures utilized in this study—VGG19, MobileNetV2, and MobileNetV3—all input images were scaled to 224 × 224 pixels. To speed up training and enhance numerical stability during model improvement, pixel intensity values were standardized from the range [0, 255] to [0, 1]. During the training phase, real-time data augmentation approaches were used to enhance model generalization and lower the danger of overfitting. These featured random shearing transformations, random zooming up to 20%, random horizontal flips, random width and height shifts up to 10%, and random rotations up to 20 degrees. Additionally, class weights were computed and applied during training to account for the dataset's modest class imbalance, giving the Lumpy Skin class a weight of 1.0875 and the Normal Skin class a weight of 0.9255. This weighting technique makes it possible for the model to give both courses equal attention throughout the learning process.

Pre-trained weights from ImageNet dataset initialized all base models, enabling knowledge transfer from general visual feature learning to specialized LSD detection.

A. VGG19 Architecture

One of the deep learning models used in this work was the VGG19 architecture, which consists of 19 weight layers (16 convolutional layers and 3 fully connected layers) with uniform 3 x 3 convolutional filters. The original top layers of the model were eliminated and replaced with a custom classification head in order to modify the pre-trained network for the binary classification job of identifying Lumpy Skin Disease. In order to prevent overfitting, the updated architecture consists of a Global Average Pooling 2D layer, a dense layer with 256 neurons and ReLU

activation, batch normalization, and a dropout layer with a rate of 0.5. Another thick layer with 128 neurons with ReLU activation, along with batch normalization and a 0.5 dropout rate, comes next. Finally, binary classification of Lumpy Skin and Normal Skin images was carried out using an output dense layer with a single neuron and sigmoid activation. 20,190,273 parameters make up the entire model, of which 9,604,353 are trainable and 10,585,920 are non-trainable. These parameters correspond to the frozen layers from the basic model that has already been trained.

B. Mobilenetv2 Architecture

Because of its effective design based on depthwise separable convolutions and inverted residual blocks with linear bottlenecks, which greatly reduce computational complexity while maintaining strong feature extraction capability, the MobileNetV2 architecture was used as a lightweight deep learning model. The pre-trained MobileNetV2 base model was modified for binary classification in Lumpy Skin Disease detection by removing its original top layers and adding a custom classification head. The further layers include a dense layer with 128 neurons with ReLU activation to learn higher-level representations, and a Global Average Pooling 2D layer to transform feature maps into a compact feature vector. To reduce overfitting during training, a dropout layer with a rate of 0.3 was added. Lastly, the probability score for the binary classification job was generated using an output dense layer with a single neuron and sigmoid activation. The final model has 2,914,369 total parameters, of which 576,448 are non-trainable and 2,337,921 are trainable. These are the frozen parameters from the base network that had already been trained.

C. Mobilenetv3 Architecture

The MobileNetV3 Large architecture, a sophisticated lightweight convolutional neural network intended for enhanced effectiveness and performance on mobile and edge devices, was utilized. This design uses improved Squeeze-and-Excitation (SE) modules, h-swish activation functions, and platform-aware Neural design Search (NAS) to improve feature representation at a low computational cost. The pre-trained base network was altered for the LSD classification task by adding a custom classification head and deleting its original classification layers. To create a compact feature vector from the retrieved feature maps, the new head starts with a Global Average Pooling 2D layer. A dense layer with 1024 neurons with ReLU activation comes next, followed by a dropout layer with a rate of 0.3 to lessen overfitting. Next, a second dropout layer with a rate of 0.3 is added after a second dense layer with 512 neurons and ReLU activation. Lastly, the binary classification between Lumpy Skin and Normal Skin images is carried out by an output dense layer with a single neuron and sigmoid activation. There are 4,505,729 parameters in the entire model, of which 4,369,929 are trainable and 135,800 are not.

D. Evaluation Metrics

To ensure a thorough evaluation of the constructed models' predictive capacity, their performance was assessed using a number of common categorization measures. While precision assessed the dependability of the model's positive predictions, accuracy was used to gauge how accurate the model's predictions were overall. The model's capacity to accurately detect positive cases was gauged by recall (sensitivity), and a fair assessment of both metrics was given by the F1-score, which is the harmonic mean of precision and recall. In order to assess the model's accuracy in identifying negative cases, specificity was also computed. A thorough analysis of true positives, true negatives, false positives, and false negatives was provided using a confusion matrix, which provided more insight into categorization performance Additionally, the model's overall classification capabilities across various decision thresholds was summarized using the Receiver Operating Characteristic – Area Under the Curve (ROC-AUC) metric. To gauge the viability of implementing the models in real-world settings, computational efficiency was assessed in addition to predicting accuracy using parameters such model size (in MB), inference time (in milliseconds), and training time per epoch (in seconds).

E. Deployment Architecture

Gradio, an open-source Python package that facilitates the quick creation of interactive web interfaces for machine learning applications, was used to implement the best model. To guarantee seamless integration and accessibility, the deployment pipeline included a

number of crucial steps. In order to load the model for inference without retraining, model serialization was first carried out by storing the trained model and its weights. A web interface was then created to enable users to input pictures of cattle skin, which are thereafter automatically preprocessed and sent to the trained model for prediction. Additionally, the system has output formatting that makes it easy for the user to see the categorization result (Lumpy Skin or Normal Skin), prediction confidence score, and pertinent recommendations. Lastly, the application was made ready for cloud hosting via containerized deployment, allowing public URL access so that researchers, farmers, and veterinarians could utilize the diagnostic tool from a distance using a regular web browser.

III. SYSTEM ARCHITECTURE

The first step in the procedure is gathering pictures of cattle skins to feed into the system. Using Python-based libraries, preprocessing techniques including scaling, normalization, and enhancement are applied to photos during the processing stage. Following preprocessing, the preprocessed photos are sent into a classifier (MobileNetV3) that uses a Convolutional Neural Network (CNN) to extract features and classify them.

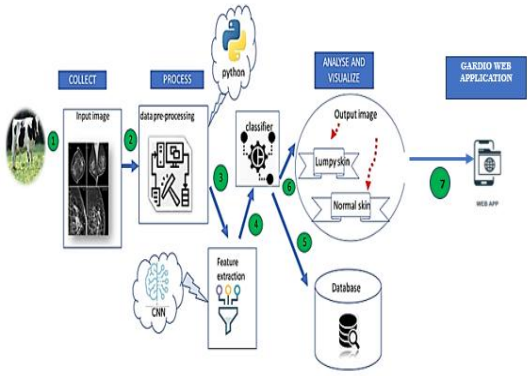

Fig. 1. PROPOSED SYSTEM ARCHITECTURE

The model classifies the image as either Lumpy Skin or Normal Skin based on an analysis of lesion characteristics. In order to maintain records and conduct additional analysis, the prediction results are kept in a database. The output label and confidence data are produced throughout the analysis and visualization phase. Ultimately, a Gradio web application incorporates the trained model, allowing users to upload photos and obtain real-time diagnostic predictions via an easy-to-use web interface. For field-level veterinary applications, this end-to-end platform guarantees automatic, effective, and easily accessible LSD detection.

IV. RESULTS

A. VGG19 Performance Analysis

The VGG19 model showed good learning characteristics throughout training. Although the model achieved its optimal validation accuracy of 97.18% at epoch 12, training concluded at epoch 18 due to early ending conditions. The learning rate scheduling performed effectively, enabling fine-grained optimization in the final training stages with drops started at epochs 6, 10, and 15.

Figure 2 shows the VGG19 training progression, which shows a smooth convergence with minimal overfitting. The validation loss gradually decreased to a minimum of 0.0841, and the training and validation accuracy curves stayed close to one another throughout the training phase.

The VGG19 model demonstrated exceptional classification performance on the validation set, with 100% recall for the "lumpy skin" class. This is particularly crucial for applications that detect diseases, as false negatives could have detrimental effects.

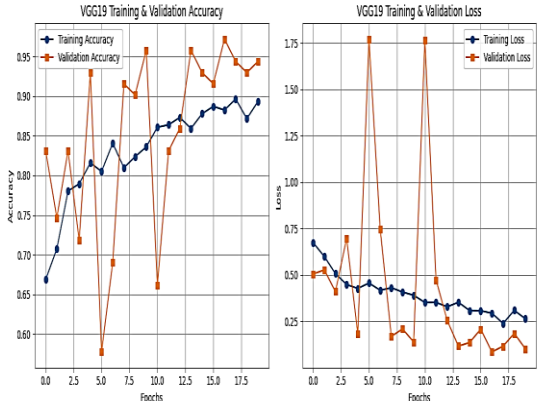

Fig. 2. TRAINING AND VALIDATION ACCURACY/LOSS CURVES(VGG19)

Fig.3 displays the VGG19 model's confusion matrix. Only two misclassifications—one false positive, in which lumpy skin was mistakenly forecasted as normal skin, and one false negative, in which normal skin was mistakenly classified as lumpy skin—occurred out of 71 validation samples.

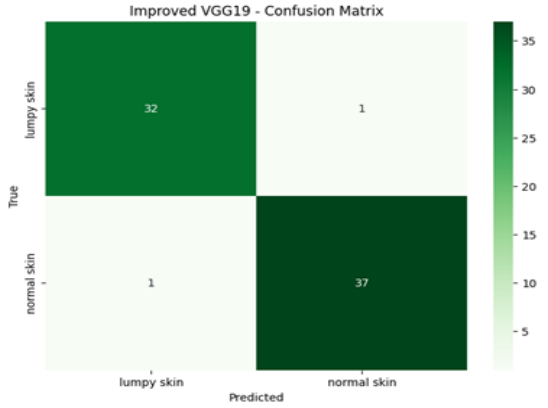

Fig. 3. Confusion Matrix of VGG19

B. Mobilenetv2 Performance Analysis

MobileNetV2 showed remarkable training efficiency, around 6.6 times quicker than VGG19. The model's best validation accuracy of 93.06% was reached at epoch 10, but training concluded at epoch 17 due to early halting. The model demonstrated consistent growth with effective learning rate adjustments during training.

Figure 4 displays the accuracy and training loss of the model. According to the accuracy plot, the validation accuracy steadily rises to almost 93%, indicating high generalization, but the training accuracy varies. Both training and validation losses gradually decline on the loss plot, suggesting that the model is successfully learning. The model's steady performance is demonstrated by the validation loss's gradual decline. No sign is present.

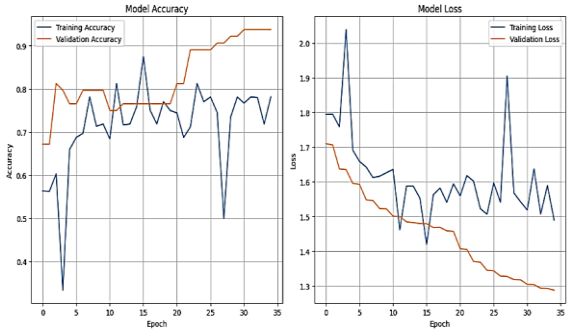

Fig. 4. TRAINING AND VALIDATION ACCURACY/LOSS CURVES(Mobilenetv2)

Figure 5 displays the MobileNetV2 model's confusion matrix. The model properly identified 28 photographs with lumpy skin and 39 images with normal skin out of the total validation samples. Images of lumpy skin were mistakenly classified as normal skin in five instances. Images of normal skin were not mistakenly identified as lumpy skin. This shows that the model performed well overall in categorization and attained high precision for the normal skin class.

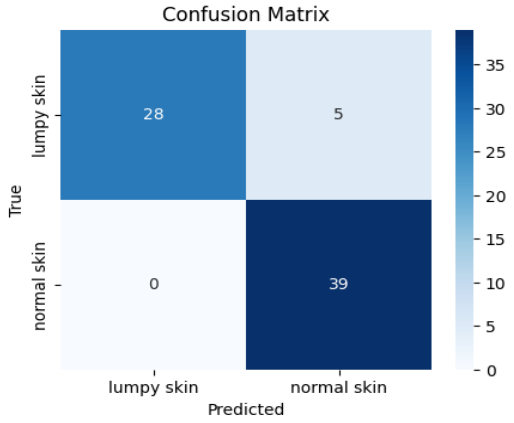

Fig. 5. Confusion Matrix of Mobilenetv2

C. Mobilenetv3 Performance Analysis

MobileNetV3 displayed training features, however it was more accurate than MobileNetV2. The final model had a peak performance of 97.18% at epoch 10, matching the maximum performance of VGG19, although stabilizing at 96% accuracy.

Figure 6 shows the training profile of MobileNetV3, highlighting both its rapid convergence and the effectiveness of the enhanced architecture in obtaining discriminative features for LSD detection.

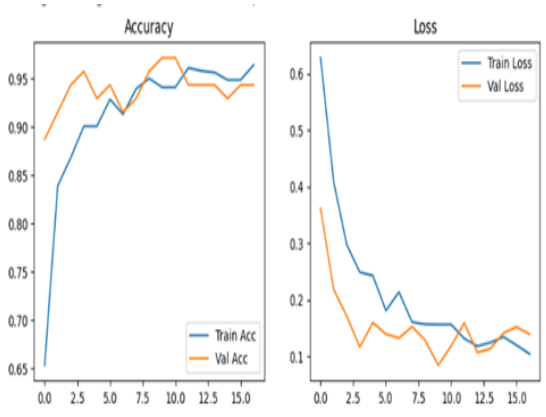

Fig. 6. TRAINING AND VALIDATION ACCURACY/LOSS CURVES(Mobilenetv3)

MobileNetV3 achieved strong classification metrics and nearly matched VGG19's performance while retaining efficiency advantages.

Figure 7 displays the MobileNetV3 model's confusion matrix. Only three false positives, in which lumpy skin was mistakenly projected as normal skin—occurred out of 71 validation samples.

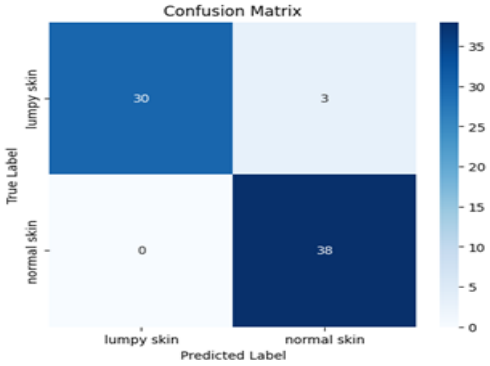

Fig. 7. Confusion Matrix of Mobilenetv3

D. Sample Prediction

Fig. 8. PREDICTION RESULT

The predictions are presented via an interactive interface that was created with Gradio, and the results are

diagnostic interface is implemented to show how a trained deep learning The technology bridges the gap between academic research and real-world veterinary applications by offering a straightforward interface that enables users to submit photos of cattle and receive instantaneous diagnostic predictions. There are several possible advantages to such a tool. In remote locations where prompt expert assistance may not be available, it can serve as a preliminary screening approach for veterinarians, facilitating the quicker identification of suspected LSD cases. Additionally, by assisting field workers and veterinary students in identifying distinctive illness patterns, the platform can function as an educational tool. Additionally, the application framework may be expanded to gather prediction results and anonymised picture data, which would help create bigger datasets for upcoming model enhancements. Three crucial elements make up each prediction: (i) the diagnostic result ("Lumpy Skin" or "Normal Skin"), (ii) a confidence score that is expressed as a percentage to show how certain the classification is, and (iii) a practical recommendation, like suggesting that a positive detection be followed right away by veterinary consultation. Even non-expert users, such as farmers and field workers, may easily comprehend the AI-generated assessment and take the necessary action thanks to this well-organized and intuitive presentation, which also improves interpretability. This method closes the gap between useful veterinary decision support tools and sophisticated deep learning models.

V. DISCUSSION

The results of the experiment show that deep learning models, especially those that use transfer learning, are very good at differentiating between healthy cattle and animals that have Lumpy Skin Disease (LSD). The hypothesis that convolutional neural networks can automatically learn discriminative visual features from cattle skin lesion images associated with LSD is supported by the consistently high accuracy, precision, and recall values obtained across the evaluated architectures—VGG19, MobileNetV2, and MobileNetV3. Because of its deep architecture and capacity to capture intricate hierarchical feature representations, VGG19 outperformed the other models in terms of categorization. Nevertheless, a greater model size and higher processing requirements came with this enhanced performance.

The MobileNet family of designs, on the other hand, showed a successful trade-off between computing efficiency and accuracy. These models, which are built with depthwise separable convolutions, drastically cut down on the number of parameters and computational expense without sacrificing their great predictive power. Specifically, MobileNetV3 outperformed VGG19 in classification performance while requiring significantly less memory and inference time. In real-world deployment circumstances, this trade-off between efficiency and precision is crucial.

Therefore, choosing the best model involves more than just looking at performance indicators. Lightweight designs are more suited for real-world field
applications, even though VGG19 might be ideal for settings with high computational resources, including research labs or cloud-based processing systems. The best model for implementation in the created web-based diagnostic system in this situation was determined to be MobileNetV3. Its much smaller model size and quicker inference speed make it
perfect for applications needing quick responses and little processing power, even though its accuracy was marginally worse than that of VGG19.

This work's successful transition from model creation to practical implementation is one of its main contributions. Using the Gradio framework, a web-based diagnostic interface is implemented.The technology bridges the gap between academic research and real-world veterinary applications by offering a straightforward interface that enables users to submit photos of cattle and receive instantaneous diagnostic predictions. There are several possible advantages to such a tool. In remote locations where prompt expert assistance may not be available, it can serve as a preliminary screening approach for veterinarians, facilitating the quicker identification of suspected LSD cases. Additionally, by assisting field workers and veterinary students in identifying distinctive illness patterns, the platform can function as an educational tool. Additionally, the application framework may be expanded to gather prediction results and anonymised picture data, which would help create bigger datasets for upcoming model enhancements.

The study's findings align with previous research on deep learning-based illness identification in poultry and animals. Similar classification accuracies surpassing 90% have been reported in earlier research using transfer learning with architectures including VGG, ResNet, and Inception in animal health monitoring applications. The results of this study support the usefulness of transfer learning for tasks involving the classification of veterinary images. However, by presenting a whole pipeline that encompasses model training, evaluation, and practical deployment within an interactive online interface—a step that is less commonly addressed in similar research—the current study goes beyond traditional model benchmarking.

The study also emphasizes how crucial accessible AI tools are becoming for veterinary care. Research prototypes may be turned into useful tools more quickly thanks to platforms like Gradio, which allow academics to create useful AI-driven applications with little knowledge of web development. Veterinarians that frequently use smartphones or digital cameras for field inspections can easily incorporate the system into their workflow because it uses a web-based interface that simply needs image upload functionality. Promoting the use of AI technologies in animal health management requires this degree of accessibility.

Despite the encouraging outcomes, there are a few things to be aware of. First, the size and variety of the dataset employed in this investigation were somewhat constrained. The model's capacity to generalize would be improved by adding photos from other cow breeds, geographical areas, illness stages, and environmental variables to the dataset. Second, the classification job was restricted to making a binary distinction between calves that were healthy and those who were infected with LSD. Other bovine skin conditions such dermatophilosis, pseudo-lumpy skin disease, and ringworm may have superficially similar symptoms in real-world situations, which could result in incorrect classification. Third, while field situations may include blurry photos, occlusions, or complicated backdrops that could impact prediction accuracy, the model was predominantly trained on reasonably clear and well-framed images. Lastly, the current web application is a deployment of a prototype. The system would need to be deployed on a scalable cloud architecture with improved capabilities like load balancing, safe data storage, and user management for widespread adoption.

Important ethical issues are also brought up by the development of AI-based diagnostic systems. Instead of taking the place of expert veterinary knowledge, the suggested system is meant to serve as a decision-support tool. Qualified veterinarians must continue to make the final diagnosis and treatment recommendations. Furthermore, inadequate representation of multiple cow populations in the

training data may lead to model bias, which could impact prediction accuracy across breeds or conditions. To guarantee that uploaded photographs and related metadata are appropriately anonymized and safeguarded, future implementations that save user-submitted images must likewise include stringent data privacy and security protections.

VI. CONCLUSION AND FUTURE WORK

This work showed that using skin lesion photos, deep learning techniques may be successfully used to the automated detection of Lumpy Skin Disease (LSD) in cattle. The effectiveness of three transfer learning models—VGG19, MobileNetV2, and MobileNetV3—in identifying sick and healthy cattle was assessed. With VGG19 and MobileNetV3 achieving the best performance of almost 97%, all models attained good classification accuracy above 91%. MobileNetV3 was found to be the best model because it offers a better mix between accuracy and computational efficiency, even though VGG19 achieved slightly higher accuracy because of its deeper architecture. It is perfect for practical deployment due to its great prediction performance, quicker inference time, and reduced model size. The created Gradio-based web service, which enables users to upload photographs of cattle and obtain immediate forecasts, further illustrates the viability of converting deep learning research into a practical utility. This method can facilitate quicker decision-making in field settings and help vets screen probable LSD cases early.

Even with the encouraging outcomes, a number of enhancements could be investigated in subsequent research. The robustness and generalization capacity of the model would be enhanced by adding more photos from various cow breeds, geographical areas, illness stages, and environmental circumstances to the dataset. In the future, the method can be expanded beyond binary classification to detect various bovine skin conditions that might resemble LSD in appearance. To further improve performance, more recent deep learning architectures like EfficientNet or transformer-based models should be assessed.

Incorporating explainable AI techniques, like Grad-CAM, would enhance user confidence in the system and help visualize the areas of the image influencing the forecast. Additionally, the method might be turned into a mobile application that would enable farmers and veterinarians to do disease screenings straight from cellphones. Lastly, the development of this technology as a trustworthy decision-support tool in cattle health monitoring would be supported by integration with veterinary health management systems and additional validation with clinical specialists.

Declarations

Corresponding author: S.Uma Shanthi
Email: samiuma2002@gmail.com

Funding Declaration: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Consent to Participate declaration: Not applicable. This study used publicly available and anonymized images; no direct human or animal subjects were involved.

Consent to Publish declaration: Not applicable. No identifiable individuals, human subjects, or animal subjects are presented in this manuscript.

Ethics declaration: Not applicable. This study did not involve human participants, animal experimentation requiring ethical approval, or the use of personally identifiable data.

REFERENCES

[1] E. S. M. Tuppurainen, E. H. Venter, J. A. W. Coetzer, and L. Bell-Sakyi, "Lumpy skin disease: Attempted propagation in tick cell lines and presence of viral DNA in field ticks collected from naturally-infected cattle," Ticks and Tick-borne Diseases, vol. 8, no. 4, pp. 543-550, 2017.

[2] G. Alemayehu, G. Zewde, and B. Admassu, "Risk factors assessment and spatial distribution of lumpy skin disease in Ethiopia," Acta Tropica, vol. 213, p. 105743, 2021.

[3] N. Kumar, S. Choudhary, and R. Singh, "Spatial and temporal analysis of lumpy skin disease outbreaks in India using GIS and remote sensing," Transboundary and Emerging Diseases, vol. 67, no. 5, pp. 1987-1997, 2020.

[4] R. Singh, S. K. Gupta, and A. Sharma, "Machine learning approaches for predicting lumpy skin disease outbreaks in Southeast Asia,"*Computers

and Electronics in Agriculture, vol. 195, p. 106812, 2022.

[5] World Organisation for Animal Health (WOAH), "Technical disease card: Lumpy skin disease," Paris, France, 2021.

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems, 2012, pp. 1097-1105.

[7] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in International Conference on Learning Representations (ICLR), 2015.

[8] R. Patel and V. Sharma, "Custom CNN architecture for automated detection of lumpy skin disease in cattle," IEEE Transactions on AgriFood Electronics, vol. 1, no. 2, pp. 112-121, 2023.

[9] L. Zhang, H. Wang, and Y. Chen, "Vision Transformer-based hybrid model for livestock disease detection," Computers in Biology and Medicine, vol. 158, p. 106887, 2023.

[10] M. Saqib, A. Ahmed, and F. Khan, "MobileNetV2 with progressive fine-tuning for lumpy skin disease classification," in IEEE International Conference on Artificial Intelligence in Agriculture, 2024, pp. 45-52.

[11] D. Senthilkumar, P. Rajasekaran, and K. Muthu, "Comparative analysis of deep learning architectures for bovine skin disease detection," IEEE Access, vol. 12, pp. 45678-45692, 2024.

[12] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018, pp. 4510-4520.

[13] A. Khan, S. Hassan, and T. Mahmood, "Data augmentation strategies for veterinary medical imaging with limited samples," IEEE Journal of Biomedical and Health Informatics, vol. 27, no. 3, pp. 1423-1434, 2023.

[14] W. Li, J. Wang, and X. Zhang, "Cross-domain transfer learning from human dermatology to veterinary skin disease detection," Artificial Intelligence in Medicine, vol. 138, p. 102512, 2023.

[15] A. Howard, M. Sandler, G. Chu, et al., "Searching for MobileNetV3," in IEEE International Conference on Computer Vision (ICCV), 2019, pp. 1314-1324.

[16] A. G. Howard, M. Zhu, B. Chen, et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.

[17] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 1251-1258.

[18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.

[19] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009, pp. 248-255.

[20] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in International Conference on Machine Learning (ICML), 2020, pp. 1597-1607.

[21] Abadi, M., et al., "TensorFlow: A system for large-scale machine learning," in 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 2016, pp. 265-283.

[22] A. Dosovitskiy, L. Beyer, A. Kolesnikov, et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in International Conference on Learning Representations (ICLR), 2021.