PAWSCAN – AI-POWERED EARLY DETECTION OF SKIN DISEASES IN STRAY DOGS

Project Reference No.: *48S_BE_0626*

College : Dayananda Sagar University, Ramanagara
Branch : Computer Science and Engineering
Guide : Prof. Sasikala N
Student(s) : Mr. Venkatesh P
Mr. Nandeesh P Math
Mr. Samarth S S
Ms. Sree Vibha G

Keywords:

Skin Disease Detection, Deep Learning, EfficientNet, MobileNetV2, Stray Dogs.

Introduction:

India has a large population of stray dogs that often suffer from untreated skin diseases due to lack of early diagnosis and accessible care. Diseases like Mange, Hotspots, Mast Cell Tumors (MCT), and Ringworm are prevalent, causing severe discomfort and health deterioration in animals. Many animal welfare Non-Government Organizations (NGOs) face challenges in identifying and managing such cases at scale. In this project, we address this issue using a deep learning-based system called PawScan. The solution aims to detect and classify four major types of skin conditions in stray dogs through image analysis. The system is designed to aid NGOs, veterinarians, and rescue volunteers in identifying infections at an early stage using a mobile or web-based interface. The project integrates image segmentation and classification with geolocation-based alert mechanisms, creating a comprehensive and scalable support tool for the stray dog ecosystem in urban India

Objectives:

1. To develop an AI-based system capable of detecting four common skin diseases in stray dogs: mange, MCT, hotspot, and ringworm.

2. To create a robust binary segmentation model that isolates infected skin regions.
3. To classify segmented regions accurately using lightweight CNNs suitable for mobile/web deployment.
4. To build a geolocation-enabled platform that alerts nearby NGOs/hospitals upon detection.
5. To provide a cost-effective and scalable tool for real-time animal health monitoring.

Methodology:

Images were collected from NGOs and rescue volunteers across South Bangalore, focusing on four specific skin diseases in stray dogs. During preprocessing, the images were resized to 256x256 pixels, normalized, and augmented using techniques such as flipping, zooming, rotation, and brightness adjustments. For segmentation, an EfficientNet-based model was implemented for binary segmentation of the affected skin areas, with EfficientNet as the encoder and transposed convolutional layers as the decoder. The segmented patches were then classified using a MobileNetV2-based model to identify the specific disease among the four targeted conditions. In the training setup, a combination of binary cross-entropy, focal loss, and dice loss (in respective ratios) was used for segmentation, while categorical cross-entropy was applied for classification. The models were trained using the Adam optimizer over multiple epochs. Evaluation of model performance was carried out using metrics such as accuracy, precision, recall, F1-score, dice coefficient, AUC, mean IoU, and confusion matrices. For deployment, a mobile application and server were developed using Flutter and Flask. The app enables users to upload images and receive disease predictions. Additionally, Flutter Maps integration with geolocation allows for locating nearby veterinary hospitals and clinics, while also sending geotagged alerts to nearby NGOs for immediate action.

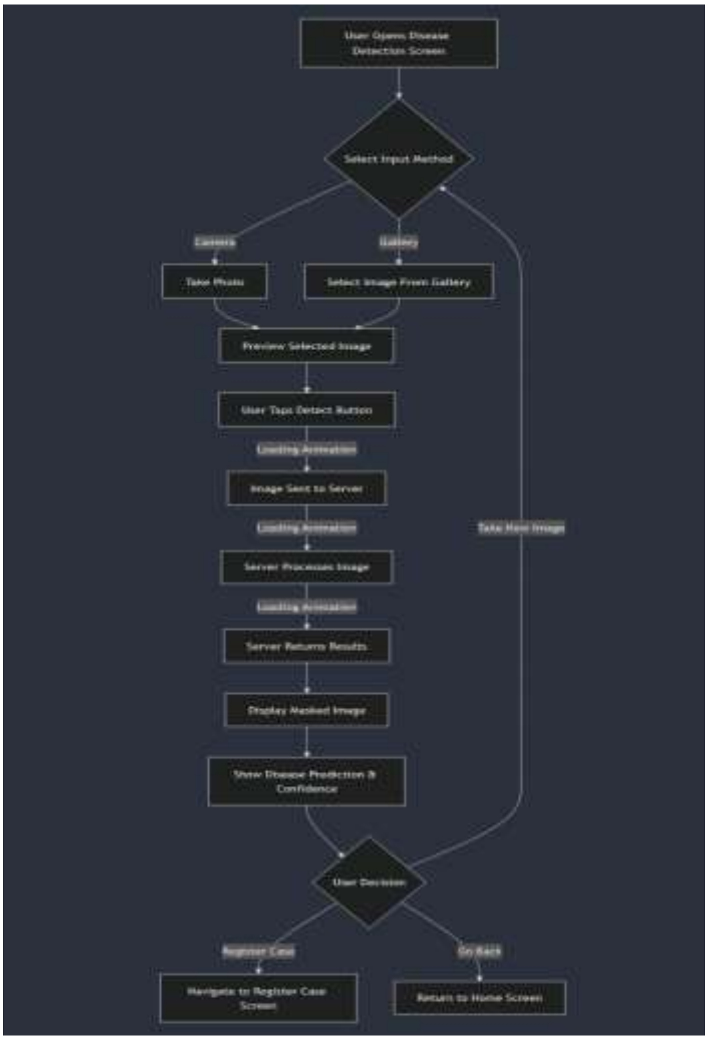

Figure 1: Disease Detection Flow in Flutter Application

Result and Conclusion:

In conclusion, the hybrid EfficientNetB0 model demonstrated strong performance, achieving a Dice coefficient of 0.66 and a mean IoU score of 0.42. The MobileNetV2 classifier attained a test accuracy of 60.61%, with a precision score of 0.59 and recall scores of 0.60 across all four disease classes. The system proved to be robust when tested on noisy, real-world images, reinforcing its practical applicability in field conditions. Visual validation further confirmed the reliability of both the segmented masks and classification outputs. Real-time predictions through the web and mobile interfaces were successfully tested using images from actual rescue scenarios. The integration of segmentation, classification, and geotagging within the system has shown significant utility in supporting and streamlining rescue workflows.

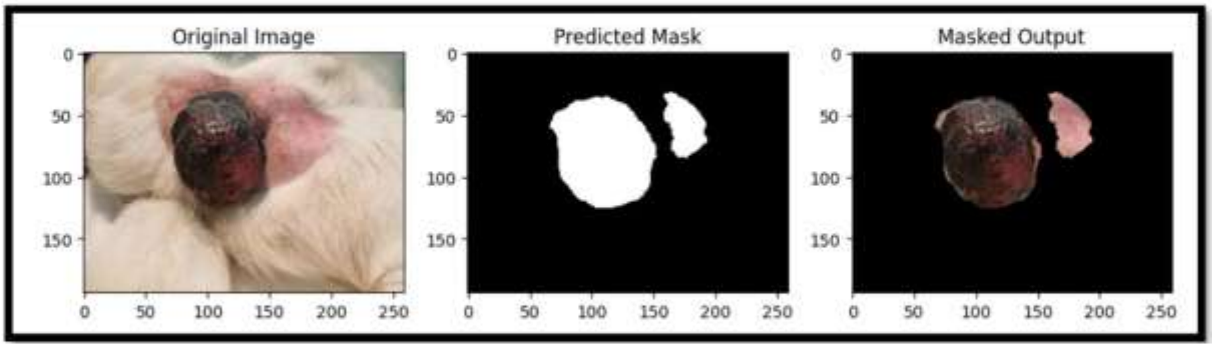

Figure 2. Example Of U-Net Segmentation Model Working For Detection

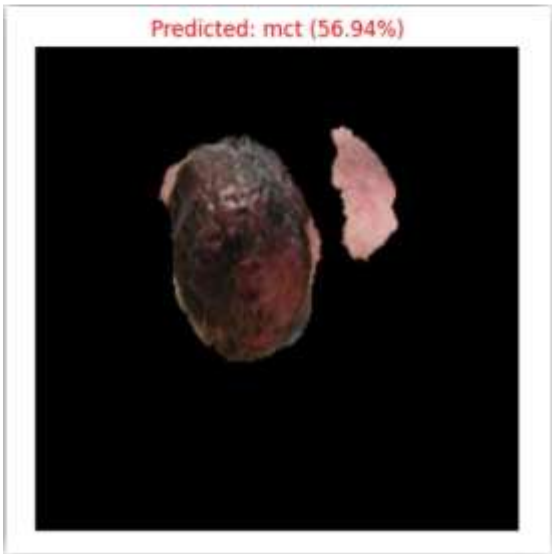

Figure 3. Example Of Classification Using MobileNetV2 Model

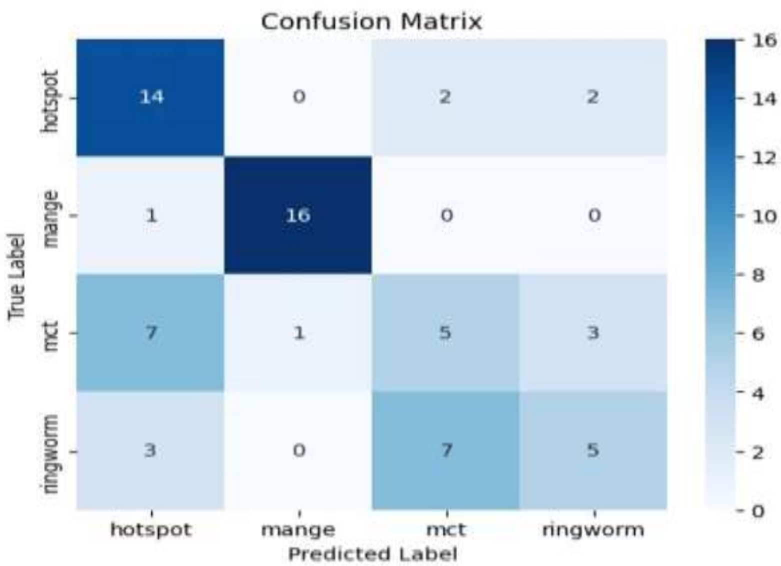

Figure 4. Confusion Matrix Of Classification Model

Project Outcome & Industry Relevance:

This project offers significant practical implications by delivering a deployable AI-based solution tailored for the welfare of stray animals, particularly in urban environments. It contributes to the field of veterinary technology and animal health by integrating lightweight deep learning models—such as EfficientNetB0 and MobileNetV2—that are optimized for mobile and web platforms, making the system accessible and efficient in real-world scenarios. The incorporation of geolocation-based alerts enables real-time rescue coordination, which is especially valuable for veterinary NGOs, animal rescue teams, and urban municipalities managing stray populations. By automating early disease detection and streamlining the reporting process, the project enhances operational efficiency and response times in field rescues. Furthermore, the modular architecture allows for easy scalability—new diseases can be added, and the model can be adapted for other animals, making it highly versatile across regions and use cases. This real-world applicability and potential for expansion make the project a meaningful contribution to both technological research and societal impact in the domain of animal welfare.

Working Model vs. Simulation/Study:

This project resulted in a working software model, including both a deployed mobile application and a trained AI pipeline for disease detection.

Project Outcomes and Learnings:

1. Built practical expertise in using hybrid EfficientNetB0 and MobileNetV2 for medical image tasks.
2. Learned how to balance model complexity with deployment constraints.
3. Gained experience integrating ML models with frontend/backend frameworks.
4. Understood real-world challenges in field data collection and standardization.
5. Developed a scalable system with both social and technological relevance.

Future Scope:

The future scope of this project includes:

1. Implement severity detection to prioritize critical cases in the rescue workflow.
2. Enable real-time notifications and emergency alerts to NGOs or medical responders.
3. Integrate user-submitted images for model retraining and improvement over time.
4. Partner with large-scale organizations like PETA, CUPA, or BBMP for wider adoption.
5. Expand model capabilities to detect additional skin or health conditions.
6. Create an analytics dashboard for centralized case tracking and disease mapping.
7. Incorporate community-based feedback for further model refinement and usability.