International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

RabiesScan: A Multimodal Deep Learning Framework for Non-Invasive Rabies Detection in Dogs via CNN-LSTM Visual Analysis and Behavioural Assessment

Krithi Y Shetty, Manya Jain, Praneetha Padmashali, Prathik Prashanth Nullipady
Dept. of Computer Science & Engineering
Mangalore Institute of Technology & Engineering
Moodabidri, India

Mr. Shivaraj B G
Assistant Professor (Guide)
Dept. of Computer Science & Engineering
Mangalore Institute of Technology & Engineering Moodabidri, India

***Abstract*—** Rabies is a deadly zoonotic virus that causes encephalitis and is thought to claim an average of 59,000 lives worldwide annually, the majority of which occur in Asia and Africa due to lack of control over their respective canine reservoirs. Current diagnostic techniques prior to death are not feasible in low-resource contexts, as they are highly lab-dependent and time-consuming. This paper introduces a new multimodal approach to non-invasive rabies detection, named RabiesScan. The proposed methodology combines two complementary sources of diagnostic information: a spatiotemporal pipeline for visual analysis of input video streams and a structured survey of the symptoms of the disease in terms of aggressive behaviour, paralysis, and excessive salivation. Visual analysis utilizes an InceptionV3 CNN as a spatial extractor of 2,048-dimensional features per input frame, after which classification is performed via a masked Long Short-Term Memory network with two layers to distinguish between the states of Normal and Rabies behaviour. The resulting classification decision together with responses to the structured questionnaire is aggregated to produce a HIGH- or LOW-Risk diagnosis alongside veterinary instructions. The whole model is deployed as a mobile application based on React Native framework and runs inference using Flask-based backend service in Python. When tested

***Keywords*—**Rabies detection; deep learning; convolutional neural network; InceptionV3; LSTM; transfer learning; multimodal fusion; mobile health; zoonotic disease; animal disease surveillance; public health AI.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

I. INTRODUCTION

The causative agent of rabies is an RNA virus of the Rhabdoviridae family, which causes acute viral encephalomyelitis in all warm-blooded animals. The WHO states that the annual number of human deaths attributed to this virus exceeds 59,000, wherein an approximate 36% and 43% of the deaths are recorded in the regions of Asia and Africa respectively [9]. Rabies accounts for about 20,000 human deaths each year in India, constituting about one-third of the total deaths worldwide, with domestic and stray dogs acting as the chief source of infection in 96% of the cases. Even though proven post-exposure prophylaxis (PEP) regimes and highly effective vaccines are available, mortality is very high as treatment is often delayed as it is difficult to ascertain within minutes if the biting animal was rabid when the bite occurred.

Clinical staging of rabies in dogs occurs in a prodromal stage of one to three days, which is characterised by slight behavioral alterations, fever, and photophobia, followed by either the furious stage, which is characterized by aggression, hypersalivation, and attacks without provocation, or the dumb stage, which involves ascending progressive flaccid paralysis and dropping jaw. As such, these signs become evident visually only at the terminal stages of either the excitatory or dumb stages, when viral shedding from the saliva is at its peak; hence, visual screening requires professional skills that are lacking in rural and peri-urban settings.

The Direct Fluorescent Antibody(DFA) method used on post-mortem samples remains the standard test for diagnosing rabies in animal subjects. Post-mortem samples are, of course, only accessible when the subject has already succumbed to the infection, and thus alternative diagnostic methods like serology and RT-PCR must be considered for an accurate ante-mortem diagnosis. Serology requires special equipment and trained personnel and can take between 24-48 hours to yield results. This leaves communities with little choice but to rely on the naked eye to observe signs that may indicate rabies, with no objective standard for taking follow-up action once a case is diagnosed.

Deep learning models have developed quickly to such an extent that transfer learning involving CNNs and the application of LSTM models to sequence modeling have made it possible to achieve reliable, non-invasive, and automatic screening for diseases based on videos. The hierarchies of visual information learned by CNNs based on large datasets of real-life images enable them to be applied efficiently to medical and veterinary imaging even without sufficient labels related to their fields [10]. Extensions involving temporal modeling using LSTM models are important for disease identification based on the progression of symptoms, which is especially needed in diseases like rabies.

In our research, we propose RabiesScan – a multi-modal deep learning framework that combines (i) a spatio-temporal InceptionV3-LSTM visual pipeline using normal smartphone videos, and (ii) a three-question structured behavioral assessment survey, in order to generate an accurate binary risk categorization as either HIGH RISK or LOW RISK, with additional symptom-based explanation of clinical implications and actions. We deploy the framework as an Android app with a lightweight Flask back-end, which works offline without any specialized hardware or network connection.

A. Principal Contributions

The primary contributions of this work are as follows:

1. Development of a specialized curated set of videos related to symptoms of canine rabies and normal dog behavior with symptom-stratified classification and frame-based processing algorithms for transfer learning purposes.
2. Introduction of a masked InceptionV3-LSTM model for binary classification task along with loss function adjusted according to data imbalance issues and implementation of sequence masking due to variable video length.
3. Adequately designing the multimodal decision fusion framework that will integrate the output probabilities of the image-based classifier with structured behavioral symptom data obtained from a validated three-question questionnaire.
4. Design and development of the full-fledged mobile application and performance evaluation with 92% accuracy, 0.99 precision in case of Rabies class and 0.93 AUC on test samples.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

II. RELATED WORK

A. CNN-Based Animal Health and Behaviour Analysis

According to Hermawan et al. [2], the VGG-19 Transfer Learning algorithm was used to categorize dog facial expressions into four emotional states (happiness, sadness, anger, and neutrality) with 95% and 90% accuracy scores during training and validation, respectively. Although the model proves successful in transferring knowledge learned from a large image dataset for recognizing small animals, the study notes that static facial expressions alone cannot be relied upon for pathological behavior detection, advising the use of posture and sound features alongside the model. The noise removal, normalisation, and augmentation processes that made up their preprocessing technique improved classification robustness, although poor generalisation resulted due to the relatively homogeneous evaluation set.

In Chen et al.'s work [5], a more advanced pipeline was constructed for surveillance videos of dogs, consisting of YOLOv3 for detection, Mask R-CNN for instance segmentation, DeepSORT for tracking, and LSTM for emotion recognition in a sequence of 16 frames. The detection achieved an accuracy of 98%, while emotion recognition achieved an accuracy of 76% to 82%, thereby confirming the efficacy of the multi-stream architecture when it comes to making inferences on canine behaviour.

B. Temporal Sequence Modelling for Behaviour Recognition

In their study, Tong [7] proposed a model called multi-stream CNN-LSTM for fine classification of canine behavior from 916 labeled videos involving four different behavior types. Among various two-stream models that combine InceptionV3 based on spatial feature extraction and optical flow-based temporal feature extraction, the one with optimal performance delivered test accuracy of 45.5%. Thus, this experiment sets up an empirical benchmark and stresses the importance of detecting short-term animal behaviors from videos. The poor performance calls for larger datasets with stratification for symptoms and improved temporal modeling. These challenges will be addressed in the proposed project. Chae et al. [8] proved that LSTM models outperformed the ARIMA and OLS approaches when using data about search engine usage, social media, and weather conditions for predicting infections' prevalence. This evidence is relevant to the current research since it justifies using LSTM for predicting sequential evolution of symptoms.

C. Rabies-Specific Computational Approaches

The agent-based model developed by Brookes et al. [6] based on GPS telemetry data from community dog populations found that rabies-caused behavioral modifications, such as a quadrupling in biting activity during the furious phase, increased contact range, and changed movement patterns allowed rabies transmission despite 70% vaccination rates. These results have important applications for designing detection systems, since vaccines will not prevent outbreaks, and therefore, quick identification of behavioral abnormalities is required to break transmission cycles. Joshy et al. [4] designed a web-based multiple module system combining CNN-based canine identification with the VIDHOP protein signature detection model for viruses such as rabies and distemper. Nonetheless, the latter requires biological samples in a lab environment and cannot be applied at the community level.

D. Identified Research Gaps

From an analysis of existing research, there emerge four consistent challenges: (1) Lack of ante-mortem non-invasive rabies diagnosis techniques capable of implementation using inexpensive computing technology; (2) Absence of integration between spatio-temporal visualization analysis and symptomatology questionnaires into one output for diagnostics; (3) Poor practicality because of testing on small and restricted data samples; and (4) Inadequate translation of AI-based predictions into comprehensible risk advisories for non-specialist users. RabiesScan solves all these problems.

TABLE I

Comparative Summary of Related Work

Reference	Modality	Architecture	Target Task	Limitation vs. RabiesScan
Balaji et al. [1]	Image	CNN + ResNet-50	Wound severity	No temporal modelling; not disease detection
Hermawan et al. [2]	Image	VGG-19 (TL)	Dog emotion	Static cues only; no clinical application
Saoud et al. [3]	Video/ Review	Multiple DL	Behaviour analysis	Review only; no deployable system
Joshy et al. [4]	Image + Bio	CNN + VIDHOP	Rabies/distemper	Requires biological sampling; web-only
Chen et al. [5]	Video	YOLO+RCNN+LSTM	Dog emotion in CCTV	No rabies focus; no mobile deployment
Tong [7]	Video	CNN-LSTM + OF	Canine behaviour	45.5% accuracy; no disease context
RabiesScan (Ours)	Video + Quiz	InceptionV3-LSTM	Rabies detection	— (proposed system)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

III. PROPOSED METHODOLOGY

A. System Architecture Overview

RabiesScan is an architecture composed of three layers: (1) Frontend Layer implemented using React Native that deals with user interaction and media recording; (2) Backend Layer based on Python Flask that provides RESTful API endpoints and executes inference using the model, and (3) Deep Learning Model Layer, which is comprised of the pre-trained InceptionV3 feature extractor and the trained LSTM sequence classifier. The two major paths for diagnosis—analysis of visual media and completion of the questionnaire—are performed in parallel, and the results are then unified at the ReportScreen component.

Note:

Fig. 1. End-to-end system architecture: user uploads media → Flask server extracts frames → InceptionV3 generates embeddings → LSTM classifies sequence → quiz fusion → risk report.

B. Dataset Construction and Preprocessing

A dataset specific to the medical field was compiled from freely accessible videos from veterinary databases, animal welfare sources, and laboratory conditions. The dataset is divided into two main directories, named Rabies and Normal, further segmented into symptom-based subdirectories. Sub-directories related to the Rabies directory include clinically observed symptoms, such as furious phase aggression, hypersalivation, dropped jaw paralysis, hind limb incoordination, and photophobia-induced avoidance behaviour. Healthy dogs carrying out activities in their surroundings and outdoors make up the sub-directory in the Normal directory. There are two benefits of organising the dataset into symptom-based sub-directories: ensuring diversity in the Rabies category, and analysis of symptom-specific model accuracy in the future.

Each video was preprocessed using an algorithm for deterministic extraction of frames using OpenCV. Up to $N_{max} = 20$ frames per video were extracted randomly. Each image was preprocessed using three consecutive operations: (i) center square cropping — extracting a square that fits into the geometric center of the frame to remove uninformative areas; (ii) resampling to the resolution of 224 x 224 px as the InceptionV3 input size requirements; and (iii) normalization using InceptionV3 specific function $\varphi(x) = (x / 127.5) - 1$. Data augmentation included only flipping horizontally with the probability $p = 0.5$ to mimic the bilateral symmetry of dogs while avoiding any unrealistic features. The data was divided into 80:20 training-validation split using stratified randomization; class imbalance was addressed using inverse-frequency weighting of classes.

C. Feature Extraction: InceptionV3

In order to extract the spatial features from the images, a pre-trained InceptionV3 network [11], trained using the ImageNet corpus with one million images, is used. The convolutional layers used in InceptionV3 have a factored architecture; this involves replacing the large $n \times n$ convolution operations with smaller $1 \times n$ and $n \times 1$ convolutions. These networks reduce parameters while keeping the same field of view, thus making them computationally effective for mobile devices. Instead of the classification layer, a Global Average Pooling (GAP) layer is used which generates a fixed 2,048 dimensional feature vector $\psi_t \in \mathbb{R}^{2048}$ from each image t. All InceptionV3 model parameters are frozen during training.

The extraction process of the features for a given video V of n images results in a feature matrix $\Psi \in \mathbb{R}^{(n \times 2048)}$. If a video consists of fewer than N_{max} number of frames, a feature matrix is created in $\mathbb{R}^{(N_{max} \times 2048)}$ along with a binary mask $m \in \{0,1\}^{N_{max}}$ to indicate the actual frame indices.

D. Sequence Classification: Masked LSTM

The Sequence Classifier model takes in the masked feature matrix (Ψ, m) and outputs a posterior $P(y \mid V)$. architecture is:$LSTM_1$: 128 units, return_sequences=True, masked input.

- $LSTM_2$: 64 units, returns final hidden state $h_T \in \mathbb{R}^{64}$.
- Dropout: rate $\delta = 0.5$, applied to h_T for regularisation.
- $Dense_1$: 64 units, ReLU activation — $f(z) = \max(0, z)$.
- $Dense_2$ (output): 2 units, softmax activation — $\sigma(z_k) = e^{z_k} / \Sigma_j e^{z_j}$, yielding class probabilities [P(Normal|V), P(Rabies|V)].

The network is trained with the sparse categorical cross-entropy loss: $\mathcal{L} = -\Sigma_i y_i \cdot \log(\hat{y}_i)$, where y_i is the ground-truth one-hot label and $\hat{y}_i$ is the predicted probability. Optimisation uses Adam [12] with an initial learning rate $\eta_0 = 1 \times 10^{-4}$. Three callbacks regulate training: (i) EarlyStopping on validation accuracy with patience $P_{es} = 8$ and weight restoration to the best epoch; (ii) ModelCheckpoint persisting the highest-validation-accuracy weights; and (iii) ReduceLROnPlateau with decay factor $\gamma = 0.5$, patience $P_{lr} = 5$, and minimum learning rate $\eta_{min} = 10^{-6}$. The maximum training budget is set to $E = 50$ epochs with a mini-batch size of $B = 32$.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

E. Behavioural Questionnaire and Multimodal Fusion

The structured questionnaire presents three clinically motivated binary items, each answered as Yes, No, or Not Sure:

5. Did the animal exhibit sudden or unprovoked aggression? (Indicator of furious rabies.)
6. Were signs of paralysis, ataxia, or incoordination observed? (Indicator of paralytic rabies.)
7. Was excessive drooling, foaming, or jaw dropping noted? (Indicator of throat-muscle paralysis.)

Each test item has been mapped ahead of time to its associated clinical interpretation in the form of a clinical interpretation string appearing in the report. The fusion algorithm is deterministic; if the output from the LSTM $P(Rabies|V) \geq 0.5$, the final risk level is HIGH RISK regardless of the quiz answers. However, if the output from the LSTM $P(Rabies|V) < 0.5$, then the risk level would be classified as LOW RISK, but in this case, the quiz answers would provide additional context. This approach gives preference to the highly sensitive visual test but uses the questionnaire for added insight.

F. Mobile Application and API Design

The frontend is developed using React Native and provides access to six screens: Login, Signup, Home (media capturing, and entering information about the animal), Quiz, Report, and Past Reports. The backend stores the user accounts securely using bcrypt hash values for password and email stored in an SQLite database with the help of Flask-SQLAlchemy. The media files are uploaded as multipart/form-data POST request payload to the /predict endpoint and the response contains a JSON object consisting of the predicted label along with its softmax value in under 10 to 15 seconds. The data that needs to be stored includes the risk category, prediction, observations made through the quiz, dog information (name, breed, custody), and time stamp.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. Quantitative Classification Performance

The trained RabiesScan model was tested against the stratified held-out validation set through four conventional classification performance metrics: accuracy, precision, recall (sensitivity), and F1-score, calculated per class based on the confusion matrix counts (TP, TN, FP, FN).

The confusion matrix analysis results were: TP_p = 406 (Rabies correctly classified), TN_n = 92 (Normal correctly classified), FP = 6 (Normal incorrectly classified as Rabies), and FN = 12 (Rabies incorrectly classified as Normal). The overall classification accuracy was 92.2%, calculated as (TP + TN) / (TP + TN + FP + FN) = 498 / 516.

Table II outlines the per-class performance metrics. The Rabies class shows almost perfect precision (0.99), which implies that less than 1% of positive samples are false positives. Recall of 0.86 implies that 86% of the actual rabies cases have been classified correctly, while the other 14% have been missed due to being early stage or unusual presentation without clear visual symptoms. The Normal class exhibits a very high recall (0.94), which validates that healthy animals are not misclassified. The macro-averaged F1-score for both classes is 0.875.

Metric	Normal	Rabies	Macro Avg.
Precision	0.89	0.99	0.94
Recall (Sensitivity)	0.94	0.86	0.90
F1-Score	0.91	0.84	0.875
Support (samples)	98	418	516

TABLE II

Per-Class Classification Performance of the RabiesScan CNN-LSTM Model

The ROC curve obtained by scanning the decision threshold values between [0,1] while plotting the TPR against FPR results in an Area Under the Curve (AUC) value of 0.93. If a model is said to have excellent discrimination power, then the AUC must be within the range [0.90, 1.00]. In this context, the sharp slope observed on the ROC graph signifies that, at any decision threshold value, the model has a very high sensitivity with a corresponding low false alarm rate; this is especially important when designing diagnostic tools.

B. Analysis of Error Classes

The six instances of false positives, or instances where the dog was wrongly identified as rabid, were found through visual inspection to be dogs who behaved normally in a contextually abnormal manner (such as playfully biting, or extreme panting during exercise). This suggests the existence of an error-causing factor and explains the motivation for designing a pathway that uses the symptom questionnaire; in such instances, a human evaluator completing the symptom questionnaire will respond 'No' to all three questions, serving as a correction for the prediction made based on the visual analysis.

The twelve instances of false negatives were all from the prodromal stage of the disease, where

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

behavioral symptoms were still latent and hence less visually prominent. Given what we know about the symptoms of rabies, this should not come as a surprise, and represents a limitation intrinsic to the task of visual detection. The resolution of this limitation is the main reason for considering a multimodal approach for future iterations.

C. Functional and System Validation

System testing was performed at four different test levels. Unit testing ensured the correctness of individual modules like OpenCV frame extractor, InceptionV3 feature computing pipeline, and quiz-response to text mapping module. Integration testing involved testing bidirectional data integrity between React Native frontend and all five Flask API endpoints (/register, /login, /predict, /save_report, /get_reports) for both valid and invalid input data. System testing involved execution of nine functional test cases for registration of users, uploading images, uploading videos, submitting quiz, unimodal image analysis, video analysis, multimodal risk report generation, normal-dog risk report presentation, and quiz-based symptom detection. Results of all nine test cases were consistent with expected output. Acceptance testing evaluated inference time on target Android mobile device (Snapdragon 730G with 8 GB RAM) and established video analysis time at 10-15 seconds, meeting the system's non-functional performance requirement.

D. Comparative Discussion

In comparison with the most related system, the multi-CNN-LSTM approach presented by Chen et al. [5], which achieves an accuracy of only 76–82%, RabiesScan demonstrates a higher level of accuracy at 92% while using a much simpler architecture that includes just two LSTM layers as opposed to a more complex structure composed of four components for detection, segmentation, tracking, and recognition tasks. As opposed to Tong’s [7] two-stream CNN-LSTM architecture, which was able to achieve only a mere accuracy of 45.5% in a general canine behaviour recognition task, the fact that RabiesScan is able to achieve 92% accuracy on the binary Normal/Rabies classification task indicates that symptom stratification and task-oriented training significantly surpass the value of optical flow addition. Inclusion of the symptom questionnaire adds another quality diagnostic layer to the system that other considered systems lack.

V. CONCLUSION

RabiesScan, described in this document, is a multi-modal deep-learning based platform for the ante-mortem detection of rabies in canines, which enables users to employ a readily available smartphone device running the Android operating system. The key technical innovation of the proposed method is its use of a Masked InceptionV3-LSTM network which is capable of extracting 2,048-dimensional spatial representations from up to 20 consecutive video frames, and classifying them as being either in the Normal or Rabies category with an accuracy of 92%, and a Precision of 0.99 for the latter, with an AUC of 0.93. The deterministic fusion of the visual result with a clinically validated three-item questionnaire yields a stratified risk report along with clinical interpretation for each of the symptoms, a feature unavailable to any prior work.

RabiesScan falls into a distinct yet unexplored category of clinical applications in which it is the first portable ante-mortem diagnostic for rabies that utilizes deep learning technology in combination with structured behavioral self-reporting. Its almost perfect classification of Rabies provides a reliable indicator for HIGH RISK cases, which helps to establish community trust. Its 14 percent false negatives reflect reality in an honest manner, representing the limitations of the system as a screening device, not a diagnostic one, which is key to responsible application of AI technology in clinical settings. This innovation helps achieve the objective of WHO’s “Zero by 30” program for eliminating dog-mediated rabies deaths by 2030.

VI. FUTURE WORK

Several possible areas of research are discussed for future development. Firstly, expanding the training set by incorporating geographically varying, breed stratified, and multilighting video samples should decrease the error rate associated with atypical presentations of the symptoms. Secondly, the LSTM module may be substituted with an architecture enhanced with attention mechanisms, such as Video Vision Transformer (ViViT) or TimeSformer, which may better capture inter-frame long-term dependencies in the video sample. Thirdly, the generation of Grad-CAM visualisations in the form of overlaid heatmap in the mobile interface will facilitate interpretability and help detect artifacts responsible for classification errors, thus increasing the trust of clinicians using the application. Fourthly, the current three-point scale of the questionnaire can be augmented and validated according to veterinarians' clinical checklists and optionally supplemented by a free-text input processed via natural language processing methods. Finally, incorporating a sensor device that measures signs related to autonomic dysfunction, such as heart rate variability, skin temperature, and motor tremors, would allow diagnosing dogs in the early phase of the rabies infection when there are no obvious visible signs.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

. REFERENCES

[1] S. Balaji, B. Vijaya Kumar, D. Prabhu, T. Sakthiswarrna, M. Rithikaa Shree, and V. Kiruthiga, "Dog Bite Detection and Hybrid Recovery Mechanism to Ensure Human Safety using Deep Learning," J. Inf. Syst. Eng. Manage., vol. 10, no. 245, 2025.

[2] S. Hermawan, A. M. Siregar, S. Faisal, and T. Al Mudzakir, "Classification of Dog Emotions Using Convolutional Neural Network Method," JANAPATI, vol. 13, no. 2, pp. 97–107, Jul. 2024.

[3] L. Saad Saoud, A. Sultan, M. Elmezain, M. Heshmat, L. Seneviratne, and I. Hussain, "Beyond Observation: Deep Learning for Animal Behavior and Ecological Conservation," Ecol. Inform., vol. 84, 102896, 2024.

[4] J. Joshy, M. R. Thayil, R. Xavier, N. Nandalal, and S. Xavier P, "A Web Application for Enhancing Stray Dog Welfare: Identification and Virus Detection," in Proc. 9th Int. Conf. Smart Comput. Commun. (ICSCC), 2023, pp. 1–6.

[5] H.-Y. Chen, C.-H. Lin, J.-W. Lai, and Y.-K. Chan, "Convolutional Neural Network-Based Automated System for Dog Tracking and Emotion Recognition in Video Surveillance," Appl. Sci., vol. 13, no. 8, p. 4596, 2023.

[6] V. J. Brookes, S. Dürr, and M. P. Ward, "Rabies-Induced Behavioural Changes Are Key to Rabies Persistence in Dog Populations: Investigation Using a Network-Based Model," PLOS Negl. Trop. Dis., vol. 13, no. 9, e0007739, 2019.

[7] A. T. S. Tong, "Capturing Spatial and Temporal Context for Fine Grained Canine Behaviour Action Recognition," Undergraduate Thesis, University of Sydney, Apr. 2019.

[8] S. Chae, S. Kwon, and D. Lee, "Predicting Infectious Disease Using Deep Learning and Big Data," Int. J. Environ. Res. Public Health, vol. 15, no. 8, p. 1679, 2018.

[9] World Health Organization, "Rabies," WHO Fact Sheet, 2023. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/rabies

[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 25, 2012, pp. 1097–1105.

[11] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," in Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR), 2016, pp. 2818–2826.

[12] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. 3rd Int. Conf. Learn. Represent. (ICLR), San Diego, CA, 2015.

[13] T. Fawcett, "An Introduction to ROC Analysis," Pattern Recognit. Lett., vol. 27, no. 8, pp. 861–874, Jun. 2006.