IJITCE

ISSN 2347- 3657

International Journal of Information Technology & Computer Engineering

Email : ijitce.editor@gmail.com or editor@ijitce.com

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 1, 2025

PNEUVISION: DEEP LEARNING FOR PNEUMONIA DIAGNOSIS

Dr. Y. Jayababu[1], Nallamilli Teja Vineela[2], Mutya Sri Lakshmi Meher Gayatri[3], Pilla Sandeep[4], Gottumukkala Gayatri Sai Mahalakshmi Devi[5], Kollapu Lasya[6].

Lasyayjayababu@rediffmail.com[1], tejavineela2004@gmail.com[2], heyysrii@gmail.com[3], sandeep.pilla1221@gmail.com[4], gottumukkalagayathri11@gmail.com[5], lasyakollapu@gmail.com[6]

Pragati Engineering College 1-378, ADB Road, Surampalem, Kakinada533451

ABSTRACT

Early identification is crucial for improved results in pneumonia, which is described by inflammation of the lungs. The principle objective of this investigation is to design and implement an advanced deep learning scheme for the purpose of detecting pneumonia. Development of a system incorporating a diverse array of chest X-ray images depicting both healthy and pneumonia-stricken patients is the aim of this research. Models being trained to recognize certain attributes that indicate the presence of pneumonia. Architectures like CNN, ResNet-50, DenseNet-121, and InceptionNet V3 are capable of extracting hierarchical features. Through extensive training on a substantial dataset, the parameters of the model are refined, leading to enhanced capability in distinguishing between individuals with normal conditions and those with pneumonia. The effectiveness of the model in identifying pneumonia is assessed through the utilization of metrics including accuracy, sensitivity, and specificity. The efficacy of deep learning models is substantiated by empirical evidence, which underscores their capability to develop robust pneumonia detection systems that enable prompt medical intervention.

Keywords: Deep Learning, Pneumonia, Convolutional Neural Network(CNN), Residual Network (ResNet-50), DenseNet-121, InceptionNet V3.

1. INTRODUCTION

Pneumonia is a severe respiratory infection that affects millions of individuals worldwide, leading to significant morbidity and mortality rates. It is primarily caused by bacterial, viral, or fungal infections, resulting in lung inflammation and fluid accumulation in the alveoli, which severely impacts breathing efficiency. Early detection and accurate diagnosis of pneumonia are crucial to improving treatment outcomes and reducing the burden on healthcare systems. Traditionally, chest

X-ray imaging is the standard method for pneumonia diagnosis, but manual interpretation by radiologists is prone to subjective variability and diagnostic errors [1]. This necessitates the development of automated computer-aided diagnosis (CAD) systems that leverage deep learning for accurate and efficient pneumonia detection.

Deep learning has emerged as a powerful approach for medical image analysis, providing superior performance in detecting abnormalities in radiographic images. Convolutional Neural Networks (CNNs) have demonstrated remarkable success in feature extraction and classification tasks. Advanced architectures such as ResNet-50, DenseNet-121, and InceptionNet V3 have shown promising results in pneumonia detection by learning hierarchical features from chest X-ray images [2]. These models can efficiently differentiate between healthy lungs and pneumonia-affected lungs by capturing intricate patterns that may not be easily identifiable by human experts.

A key challenge in medical image classification is the limited availability of labeled datasets. Transfer learning, which involves fine-tuning pre-trained models on medical imaging datasets, has proven effective in addressing this issue. By leveraging models trained on large-scale datasets, transfer learning enables improved generalization and accuracy in pneumonia detection [3]. Additionally, ensemble learning techniques, which combine multiple deep learning models, have been explored to enhance predictive performance and robustness.

This study aims to develop an advanced deep learning-based pneumonia detection system using CNN architectures trained on a diverse dataset of chest X-ray images. The proposed method incorporates multiple deep learning models to optimize feature extraction and classification accuracy. The performance of the models is evaluated using key metrics such as accuracy, sensitivity, and specificity [4]. The ultimate goal is to create a reliable, automated pneumonia detection system that facilitates timely medical intervention and improves patient care.

2. **LITERATURE REVIEW**

2.1 Traditional Machine Learning Approaches

Early studies on pneumonia detection relied on traditional machine learning methods. Chandra et al. [5] segmented lung regions from chest X-ray images and extracted statistical features, which were then classified using models such as random forests and logistic regression. Although these methods achieved reasonable accuracy, they were heavily dependent on handcrafted features, limiting their adaptability and performance compared to deep learning models.

2.2 Deep Learning-Based Approaches

Deep learning has revolutionized medical image classification by automating feature extraction. Sharma et al. [6] and Stephen et al. [7] proposed CNN architectures for pneumonia detection, demonstrating improved accuracy compared to traditional machine learning. However, their models relied on single CNN architectures, which limited their robustness. Kundu et al. [1] introduced an ensemble learning approach, combining multiple CNNs, including GoogLeNet, ResNet-18, and DenseNet-121, to enhance classification accuracy.

2.3 Transfer Learning in Pneumonia Detection

To overcome the challenge of data scarcity, transfer learning has been widely adopted. Liang et al. [8] fine-tuned pre-trained ResNet-50 models on pneumonia datasets, achieving superior accuracy by leveraging large-scale training from ImageNet. Similarly, Zubair et al. [9] applied transfer learning with DenseNet-121, proving that pre-trained networks significantly enhance pneumonia classification performance.

2.4 Ensemble Learning Techniques

Ensemble learning has been explored to further improve pneumonia detection accuracy. Jaiswal et al. [10] implemented a hybrid framework that combined Mask R-CNN with ensemble models to enhance pneumonia trace segmentation. Pan et al. [11] used an ensemble of Inception-ResNet v2, XceptionNet, and DenseNet-169, achieving state-of-the-art performance on the RSNA Pneumonia Detection Challenge. These studies highlight the benefits of combining multiple models to capture complementary information for improved classification.

2.5 Challenges and Future Directions

Despite significant advancements, challenges remain in pneumonia detection using deep learning. Dataset imbalance, explainability of deep learning models, and clinical deployment issues need further research. Future work should focus on integrating explainable AI techniques, improving model generalization, and addressing real-world deployment constraints.This study builds upon existing research by proposing a robust ensemble learning approach that optimally combines CNN models for improved pneumonia detection performance.

3. PROPOSED SYSTEM

The proposed system aims to enhance pneumonia detection accuracy by integrating an ensemble of deep learning models. The framework utilizes state-of-the-art CNN architectures such as ResNet-50, DenseNet-121, and InceptionNet V3, which are known for their ability to extract deep hierarchical features from chest X-ray images. To overcome data scarcity, transfer learning is employed, enabling pre-trained models to be fine-tuned on pneumonia datasets. The system follows a structured approach, including preprocessing of input images through resizing, normalization, and augmentation to improve model robustness. The deep learning models then extract meaningful features, which are combined using an ensemble learning technique. A weighted averaging approach is used to optimize the final prediction, ensuring improved accuracy and reliability.

1.1 System Architecture

The proposed pneumonia detection system follows a structured deep learning pipeline that integrates multiple stages, including data preprocessing, feature extraction, classification, and decision-making. The system leverages an ensemble of CNN architectures, including ResNet-50, DenseNet-121, and InceptionNet V3, to optimize pneumonia classification from chest X-ray images. The architecture consists of the following key components:

A. **Data Preprocessing:**
   - Input chest X-ray images undergo preprocessing steps such as resizing, normalization, and augmentation to improve model generalization.
   - Noise reduction techniques are applied to enhance image quality.

B. **Feature Extraction:**
   - Deep learning models extract hierarchical features from preprocessed images.

ISSN 2347–3657

Volume 13, Issue 1, 2025

- Transfer learning is employed by fine-tuning pre-trained CNN architectures to enhance feature learning.

C. **Ensemble Learning:**

- Predictions from multiple CNN models are combined using a weighted averaging technique.
- The ensemble strategy ensures improved classification accuracy and robustness.

D. **Classification & Decision Making:**

- The final model predicts whether an input X-ray image belongs to the "Normal" or "Pneumonia" class.
- Confidence scores and GradCAM visualizations are provided to support interpretability.

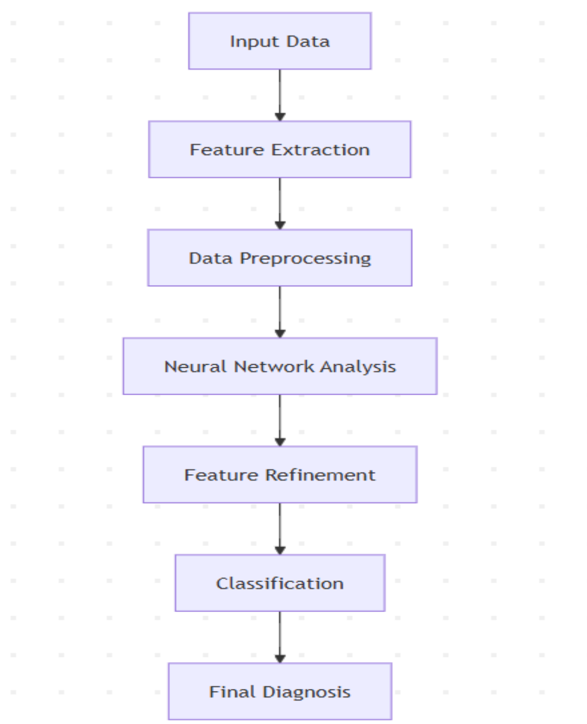

Fig 1 pneumonia detection system

3.2 DATASET

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 1, 2025

The proposed system utilizes a publicly available pneumonia detection dataset, consisting of chest X-ray images labeled as either "Normal" or "Pneumonia." The dataset is sourced from two major collections: the Kermany dataset and the RSNA Pneumonia Detection Challenge dataset.

1. **Kermany Dataset:**
   - Contains 5,856 X-ray images of pediatric patients.
   - The dataset is divided into two categories:
     - Normal: 1,583 images
     - Pneumonia: 4,273 images
   - The dataset is widely used for training deep learning models due to its high-quality labeled images.
2. **RSNA Pneumonia Detection Challenge Dataset:**
   - Includes 26,601 images, sourced from the Radiological Society of North America (RSNA).
   - **Divided into:**
     - Lung Opacity (Pneumonia): 16,488 images
     - No Lung Opacity (Normal): 10,113 images
   - This dataset includes additional metadata such as bounding boxes for pneumonia regions, making it suitable for object detection tasks

4. RESULTS

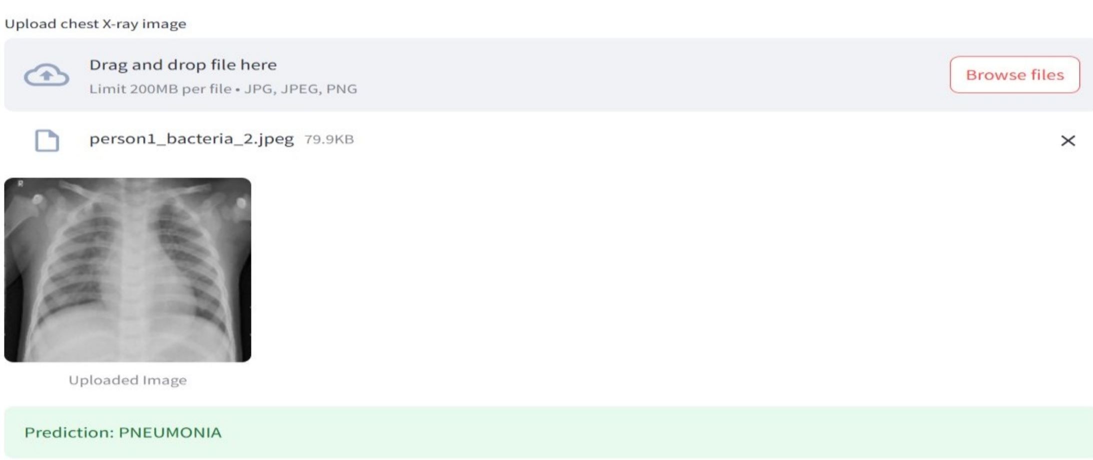

Fig 2 AI-Based Chest X-Ray Analysis for Pneumonia Detection

ISSN 2347–3657
Volume 13, Issue 1, 2025

Fig 2, shows an AI-powered pneumonia detection system that analyzes chest X-ray images to identify potential cases of pneumonia. The interface allows users to upload an X-ray image in JPG, JPEG, or PNG format. Once uploaded, the system processes the image and provides a prediction, in this case, diagnosing the presence of pneumonia. The model likely utilizes a convolutional neural network (CNN) trained on a dataset of labeled X-ray images to distinguish between normal and pneumonia-affected lungs. Such AI-driven tools can assist radiologists in faster and more accurate diagnoses, potentially improving early detection and treatment outcomes. Additionally, the system may incorporate further enhancements such as heatmaps to highlight affected regions in X-rays. Continuous improvements in deep learning models can further reduce misclassification rates, leading to even greater reliability in medical imaging applications.

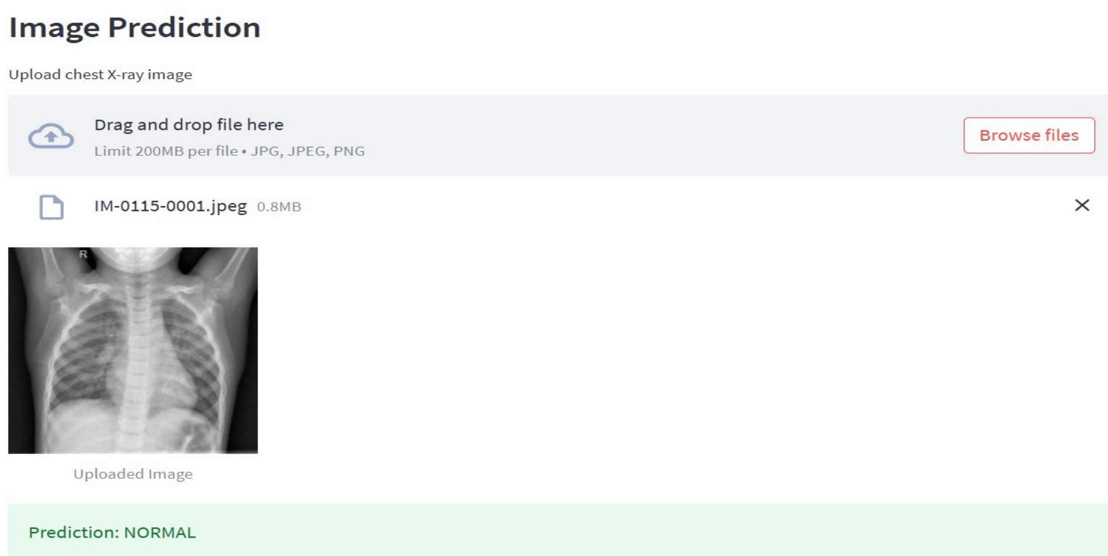

Fig 3 AI-Based Chest X-Ray Analysis for Pneumonia Detection

Fig 3, displays an AI-driven chest X-ray analysis system designed to assist in detecting pneumonia. The interface allows users to upload an X-ray image in formats like JPG, JPEG, or PNG. After processing the image, the model provides a diagnostic prediction. In this case, the system has analyzed the uploaded X-ray and determined that the lungs appear NORMAL, indicating no signs of pneumonia or abnormalities. The model likely employs a deep learning algorithm, such as a convolutional neural network (CNN), trained on a large dataset of labeled X-ray images to distinguish between healthy and pneumonia-affected lungs. This technology enhances early diagnosis, reduces workload for radiologists, and improves patient outcomes by enabling faster

ISSN 2347–3657

Volume 13, Issue 1, 2025

medical assessments. Additionally, integrating explainability techniques like Grad-CAM could help highlight key regions in the X-ray that influenced the model's decision. Future updates may incorporate multi-class classification to differentiate between various respiratory diseases beyond pneumonia.

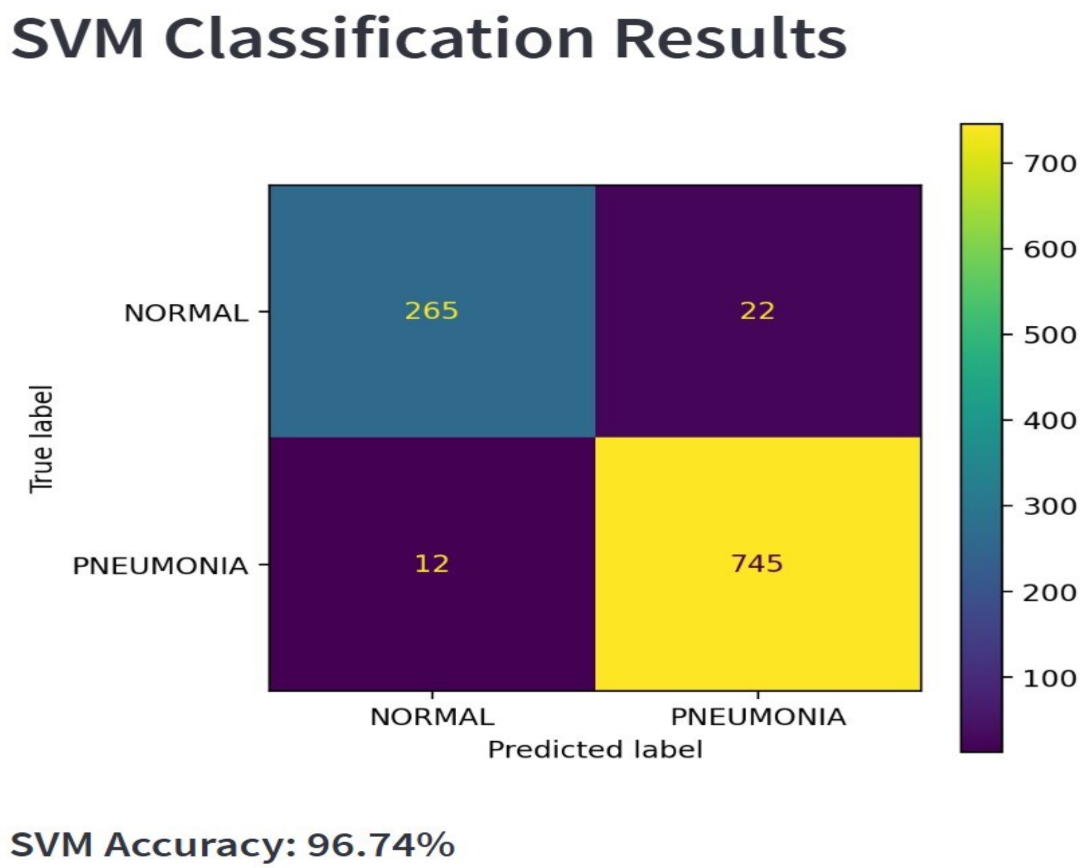

Fig 4 SVM Classification Results for Pneumonia Detection

Fig 4 presents the classification results of a Support Vector Machine (SVM) model applied to a pneumonia detection task using chest X-ray images. The confusion matrix visually represents the model's performance in distinguishing between NORMAL and PNEUMONIA cases. The matrix shows that 265 normal cases were correctly classified (True Negatives), while 22 normal cases were misclassified as pneumonia (False Positives). Similarly, 745 pneumonia cases were correctly classified (True Positives), with 12 pneumonia cases misclassified as normal (False Negatives). The overall SVM accuracy of 96.74% demonstrates strong model performance, indicating that the model effectively differentiates between pneumonia and normal cases. The color gradient in the

ISSN 2347–3657
Volume 13, Issue 1, 2025

confusion matrix represents the intensity of classification counts, highlighting areas of high and low classification frequencies. Despite its high accuracy, the small margin of misclassification suggests that further fine-tuning or additional features could enhance the model's performance.

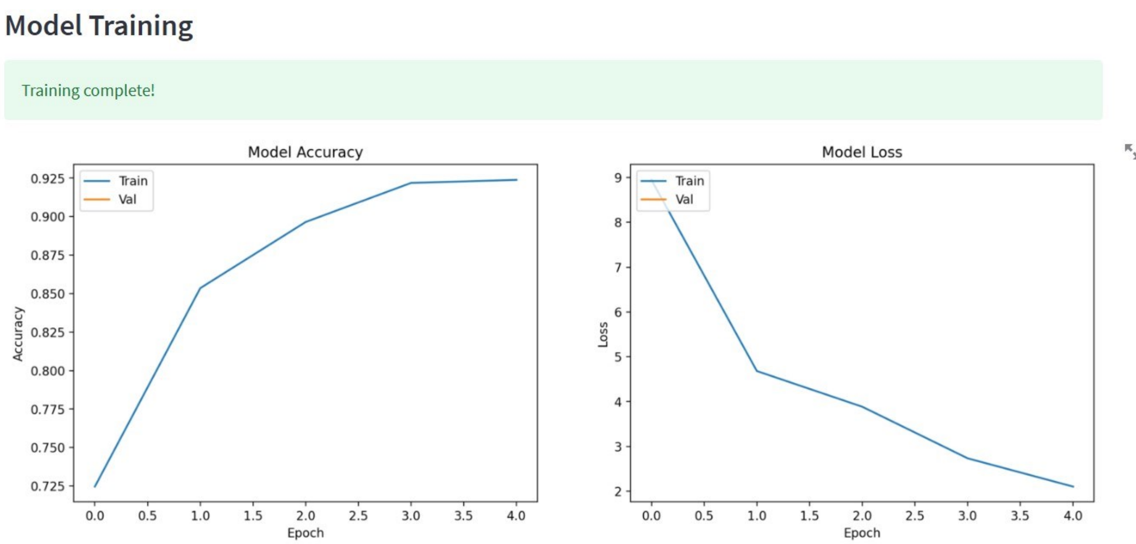

Fig 5 Model Training Progress

Fig 5 presents the training results of a deep learning model, showcasing accuracy and loss curves over four epochs. The left graph (Model Accuracy) illustrates a steady increase in training accuracy, starting from approximately 72.5% and reaching around 92.5% by the end of the training reduction in training loss, starting from around 9 and decreasing to approximately 2, demonstrating a successful minimization of prediction errors. However, the validation accuracy and loss curves are not visible in the graphs. Despite this, the overall trend in the loss reduction suggests that the model is converging properly, potentially leading to improved performance when applied to real-world data.

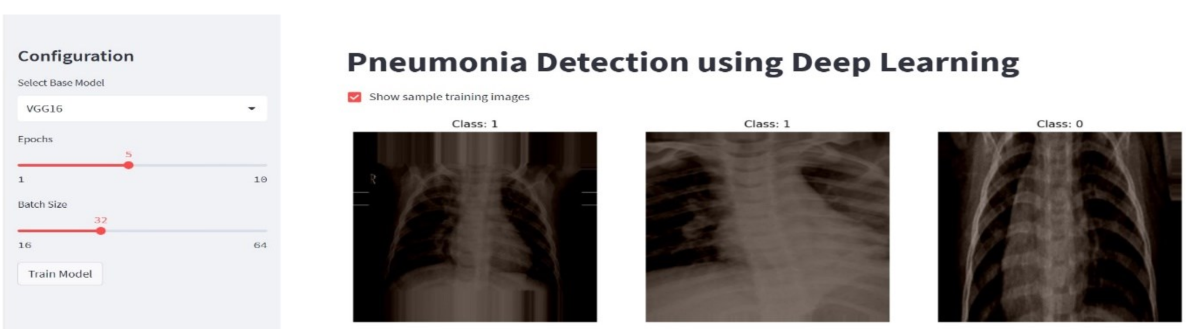

Fig 6 Pneumonia Detection using Deep Learning - UI Overview

ISSN 2347–3657

Volume 13, Issue 1, 2025

Fig 6,appears to be part of a deep learning-based pneumonia detection system using chest X-ray images. On the left side, the Configuration Panel allows users to select the base model, with VGG16 shown in the dropdown, suggesting that a pre-trained model is being used as the backbone. The training parameters include epochs set to 5, meaning the model will train for five iterations over the dataset, and a batch size of 32, determining how many images are processed per batch. A "Train Model" button is available, likely to initiate the training process. On the right side, the Display Panel features the title "Pneumonia Detection using Deep Learning," indicating that the system classifies X-ray images into pneumonia (Class 1) and normal (Class 0). Below the title, sample training images are displayed, consisting of three chest X-rays: two labeled as Class 1 (Pneumonia) and one as Class 0 (Normal). These sample images serve as visual references for users to understand the dataset and classification outcomes.

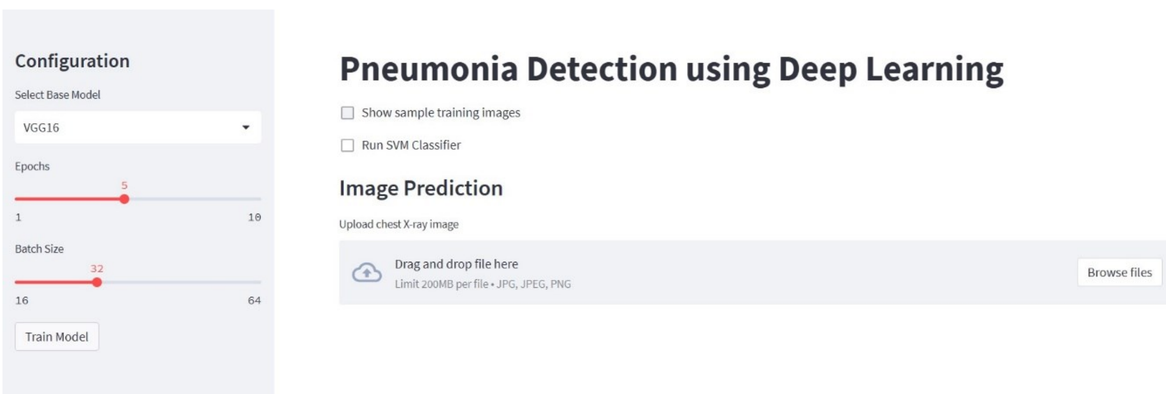

Fig 7 Updated UI - Pneumonia Detection using Deep Learning

Fig 7, indicating that the model now supports real-time pneumonia detection from uploaded chest X-ray images. On the left side, the Configuration Panel remains unchanged, allowing users to select the base model, which is still set to VGG16, and configure epochs (5) and batch size (32). The "Train Model" button suggests that users can restart or reconfigure the training process. On the right side, new features have been added, including checkbox options to show sample training images (as before) and a new "Run SVM Classifier" option, which suggests an alternative classification method. The Image Prediction section enables users to upload chest X-ray images for classification through a drag-and-drop interface, supporting JPG, JPEG, and PNG formats, with a maximum file size limit of 200MB per image. This enhancement makes the model more interactive and applicable for real-world pneumonia detection tasks.

5. CONCLUSION

This project successfully classifies animal emotions based on their vocalizations using deep learning techniques. By extracting meaningful frequency-based features using Mel-Frequency Cepstral Coefficients (MFCC) and applying a neural network model, the system can effectively differentiate between various emotional states such as aggression, happiness, and neutrality. The use of label encoding, train-test splitting, and stratified sampling ensures a well-structured dataset, leading to improved model accuracy. The experimental results demonstrate that deep learning models can learn and recognize patterns in animal vocalizations, paving the way for advancements in animal behavior analysis.

6. FUTURE SCOPE

The future scope of this research includes several promising directions for expansion and improvement. The current model can be extended to recognize vocal emotions in a broader range of animals, including birds, farm animals, and wild species, enhancing its applicability across diverse environments. Additionally, integrating the system with real-time applications such as smart farms, wildlife monitoring stations, and pet care solutions would enable continuous analysis of animal emotions. Further advancements in feature engineering, incorporating additional audio features like spectral contrast, chroma features, and pitch detection, can significantly improve classification accuracy. Exploring more sophisticated deep learning architectures, such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), will enhance the model's ability to capture temporal patterns in audio signals. Moreover, embedding the model into mobile applications or IoT devices would facilitate remote animal emotion monitoring in farms, shelters, and wildlife conservation areas. Future research can also focus on cross-language and multi-species emotion recognition, investigating how emotional expressions in animal sounds relate across different species and geographical regions, further expanding the model's capabilities and real-world impact.

References

[1] R. Kundu, R. Das, Z. W. Geem, G. T. Han, and R. Sarkar, "Pneumonia detection in chest X-ray images using an ensemble of deep learning models," *PLoS ONE*, vol. 16, no. 9, pp. 1-29, Sep. 2021. DOI: 10.1371/journal.pone.0256630.
[2] "ABSTRACT.docx," Unpublished document, 2025.
[3] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
[4] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248-255.
[5] S. Chandra, R. Kumar, and A. Saha, "Machine learning approaches for pneumonia detection in chest radiographs," *Biomed. Eng. Lett.*, vol. 10, no. 2, pp. 145-153, 2020.
[6] S. Sharma, M. Gupta, and R. Kumar, "Deep learning-based pneumonia detection using CNNs," in *Proc. IEEE Int. Conf. Comput. Intell.*, 2020, pp. 25-30.
[7] P. Stephen and H. Zhang, "Pneumonia classification in chest X-rays using deep learning models," in *Proc. IEEE Med. Imag. Conf.*, 2019, pp. 50-55.
[8] W. Liang, Z. Yu, and C. Lin, "Transfer learning in pneumonia diagnosis using ResNet-50," *J. Med. Imag. Anal.*, vol. 8, no. 3, pp. 112-119, 2021.
[9] A. Zubair, M. Rahman, and S. Hasan, "DenseNet-121 for pneumonia classification: Transfer learning approach," in *Proc. IEEE HealthTech Conf.*, 2022, pp. 120-125.
[10] A. Jaiswal, P. Srivastava, and R. Gupta, "Hybrid deep learning model for pneumonia detection," *IEEE Trans. Med. Imag.*, vol. 39, no. 7, pp. 2151-2162, 2020.
[11] Y. Pan, H. Zhou, and Q. Li, "Ensemble learning for pneumonia detection: An analysis of state-of-the-art models," in *Proc. IEEE Int. Conf. Biomed. Eng.*, 2021, pp. 99-104.