PROJECT ABSTRACT

Pneumonia Detection from Chest X-ray using CNN

A medical-imaging AI that classifies chest X-rays as normal or pneumonia using a fine-tuned CNN (ResNet/MobileNet). The app shows the prediction with confidence and Grad-CAM heatmaps over lung regions, plus a batch-screening dashboard — a high-impact healthcare AI project with real clinical datasets.

01 Title

Pneumonia Detection from Chest X-ray using CNN

02 Introduction / Background

Pneumonia is a leading killer of children under five, and early radiological detection dramatically improves outcomes. But radiologist coverage is thin in many regions, and reading scans is slow and fatiguing. A CNN that spots pneumonia in chest X-rays can triage cases for human review. This project builds exactly that: a transfer-learned classifier with visual explanations — an educational screening demonstrator, not a certified diagnostic device.

03 Problem Statement

Manual X-ray interpretation is slow and expertise is scarce. Classifiers trained from scratch overfit on limited medical data, while black-box models give clinicians no reason to trust them. The project must deliver high accuracy with every prediction auditable — proving the model reads lung pathology, not artifacts.

04 Objectives

- Classify chest X-rays as normal or pneumonia-positive.
- Use transfer learning from pre-trained ResNet50/MobileNetV2 backbones.
- Preprocess with CLAHE contrast enhancement at 224×224.
- Handle class imbalance with weighted loss and plausible augmentation.
- Explain every prediction with Grad-CAM heatmaps over lung fields.
- Evaluate with sensitivity, specificity, F1 and AUC; ship the LungScan app.

05 Proposed System / Working

Input: a chest X-ray image uploaded to the LungScan Streamlit app. Processing: the image is resized to 224×224, normalized and contrast-enhanced with CLAHE, then passed through a fine-tuned ResNet50/MobileNetV2 backbone — ImageNet pre-trained, with early layers frozen while

deeper layers learn pneumonia patterns such as consolidations. Training uses the NIH Chest X-ray dataset (112,120 frontal-view images from 30,805 patients), with weighted loss and augmentation (rotations, flips, brightness shifts) countering class imbalance.

Decision and output: the model returns a normal/pneumonia verdict with a confidence score, and Grad-CAM highlights the lung regions that drove the prediction — proving the model reads anatomy, not text markers or borders. Batch mode screens multiple X-rays with CSV export, and the training notebooks log ROC curves, the confusion matrix and per-class metrics. Inference runs in about a second per image on CPU.

06 Technologies / Components Used

Python · TensorFlow/Keras · ResNet50 · MobileNetV2 · OpenCV (CLAHE) · Streamlit · NIH Chest X-ray dataset · scikit-learn

07 Key Features

- Normal vs. pneumonia classification with confidence score.
- Transfer learning on ResNet50/MobileNetV2 backbones.
- Grad-CAM heatmaps over lung regions.
- LungScan app — upload, predict, visualize in one click.
- Batch screening with CSV export.
- Class-imbalance handling (weighted loss, augmentation).
- Training notebooks: ROC curves, confusion matrix, per-class metrics.

08 Applications / Use Cases

- Radiology triage in understaffed hospitals and rural screening camps.
- Telemedicine pre-screening and medical-education demos.
- Explainable-AI research — always a screening aid, never an autonomous diagnostic device.

09 Future Scope

- Multi-disease classification across all 14 NIH pathology labels.
- DICOM/PACS integration for direct hospital deployment.
- Lung-segmentation preprocessing to focus the model on anatomy.
- Federated learning across hospitals without sharing patient data.
- Mobile screening app for rural health camps.

10 Conclusion

This project combines social impact with examiner-favourite topics — transfer learning, class imbalance, medical preprocessing, explainability and clinical metrics. Grad-CAM heatmaps prove the model reads anatomy, not artifacts.