PROJECT ABSTRACT

Pneumonia Detection using Chest X-Rays

A DenseNet-121 classifier that screens frontal chest X-rays for pneumonia-related lung opacities on the RSNA dataset, with Grad-CAM overlays — delivered as a built-to-order academic prototype (not a diagnostic device) with code, viewer app, report, PPT and viva kit.

01. Title

Pneumonia Detection using Chest X-Rays

02. Introduction / Background

Pneumonia remains a major global cause of mortality, and the chest X-ray is the frontline diagnostic — yet expert reading takes a trained radiologist, and imaging volumes keep rising faster than radiologist availability in many regions. Deep convolutional networks have shown strong results on radiograph classification benchmarks, and the 2018 RSNA Pneumonia Detection Challenge released 26,684 expertly annotated frontal-view radiographs that make this a tractable student project. This build fine-tunes a DenseNet-121 classifier on that dataset and wraps it in a study-viewer demo with Grad-CAM explainability.

03. Problem Statement

Radiologist shortages and rising imaging volumes delay chest X-ray interpretation, while manual triage does not scale. The project addresses this by building an academic screening prototype that classifies frontal radiographs as Normal, No Lung Opacity or Lung Opacity, localizes suspected opacities with bounding boxes, and visualizes its evidence with Grad-CAM heatmaps — explicitly positioned as a triage aid requiring radiologist review, never as a diagnostic device.

04. Objectives

- Build the RSNA data pipeline: DICOM-to-PNG conversion, stratified splits, augmentation, and the three-class taxonomy with radiologist bounding boxes.
- Fine-tune DenseNet-121 (ImageNet transfer learning) for 3-class classification with a design target of macro-F1 approximately 0.80+ on the validation split.
- Implement Grad-CAM heatmaps and opacity bounding-box overlays in a study-viewer demo with a batch review queue and metrics dashboard.
- Deliver the complete student kit with an explicit clinical-safety statement: code, training notebook, report, PPT and viva Q&A.

05. Proposed Methodology

Frontal-view DICOM radiographs are converted to PNG, resized to 224x224 and normalized with ImageNet statistics. The set is split into stratified train/validation/test partitions; training images receive rotation, flip and contrast augmentation. DenseNet-121 pretrained on ImageNet has its classifier head replaced with a 3-way output and is fine-tuned with cross-entropy loss. Every epoch logs training/validation loss, per-class precision, recall, F1 and ROC-AUC; the best validation checkpoint is kept. Grad-CAM backpropagates the predicted class score to the final convolutional block for heatmap overlays.

06. System Architecture / Working

The system has three stages. (1) Data pipeline: RSNA DICOMs become normalized 224x224 tensors through torchvision transforms with augmentation on the training split. (2) Model: DenseNet-121 fine-tuning with per-epoch metric logging and checkpointing. (3) Demo application: a Flask viewer loads the best checkpoint, renders class probabilities per study, and offers overlay modes — original, Grad-CAM heatmap, opacity boxes, combined — plus a batch review queue across the three classes and a metrics dashboard with the confusion matrix and one-vs-rest ROC curves computed on the validation split.

07. Hardware / Software Requirements

- A GPU (CUDA) is strongly recommended for training; CPU-only training is possible but slow
- Python 3.10, PyTorch, torchvision (transforms, DataLoaders), pydicom, PIL
- scikit-learn (per-class metrics, confusion matrix, ROC-AUC); Matplotlib, Seaborn for plots
- Flask for the study-viewer demo
- RSNA Pneumonia Detection Challenge dataset from Kaggle (26,684 radiographs; free account required)
- Jupyter notebook for the buyer-run training experiment

08. Expected Outcomes

A fine-tuned 3-class radiograph classifier with per-class precision/recall/F1 and ROC-AUC computed on the buyer's own validation split during their build; a viewer demo that makes predictions, heatmaps and opacity boxes inspectable study-by-study; and a report with error analysis and an explicit clinical-safety statement. The design target is macro-F1 approximately 0.80+ — a target, not a pre-measured claim.

09. Applications

- Academic demonstration of medical-image classification and transfer learning

- Teaching material for CNN coursework: DenseNet, Grad-CAM, ROC-AUC, handling label noise
- Prototype triage-queue interface for radiology workflow demos (non-clinical)
- Reference build for other radiograph classification tasks

10. Future Scope

- Object-detection heads (e.g. RetinaNet) trained directly on the radiologist bounding boxes for tighter localization
- External validation on the NIH ChestX-ray14 or pediatric sets with domain-shift analysis
- Uncertainty estimation (e.g. Monte-Carlo dropout) to flag low-confidence studies for priority human review
- Multi-view and clinical-context fusion — strictly as research, not as a device claim

11. References

- RSNA Pneumonia Detection Challenge 2018 (Kaggle) — 26,684 frontal-view chest radiographs with radiologist annotations.
- Shih, G. et al. — Augmenting the NIH Chest Radiograph Dataset with Expert Annotations of Possible Pneumonia, Radiology: Artificial Intelligence, 2019.
- Huang, G. et al. — Densely Connected Convolutional Networks (DenseNet), CVPR 2017.
- Selvaraju, R. R. et al. — Grad-CAM: Visual Explanations from Deep Networks, ICCV 2017.