Kevin Luo

Lightweight CNN Model for Pneumothorax Detection using Chest X-ray Images

Oxford Brookes University in collaboration with Chengdu University of Technology
Supervised by Dr Happy Nkanta Monday

Abstract

Pneumothorax is a life-threatening condition caused by air in the pleural cavity. Accurate and timely diagnosis is essential but often limited by radiologist experience. This study presents a deep learning model combining EfficientNet and U-Net to detect pneumothorax from chest X-rays. Trained on public datasets, the model achieved strong performance with a ROC_AUC of 0.86, mIoU of 0.20, and mDice of 0.30. The results indicate that the proposed model can support automated, accurate diagnosis, aiding clinical decision-making and improving patient outcomes.

Dataset

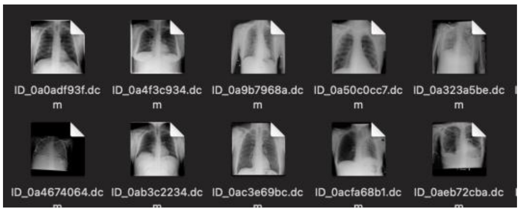

Figure 1 Examples of Pneumothorax Images

In the data preprocessing stage, all input chest X-ray images are first resized to a fixed resolution to ensure consistency across the dataset. For the training set, random transformations such as cropping, flipping, and scaling are applied to enhance data diversity and improve generalization. Image enhancement techniques, including contrast adjustment and noise injection, are optionally used to simulate real-world variability. The images are then converted into PyTorch tensors and normalized using the mean and standard deviation of the dataset to match the input requirements of the neural network.

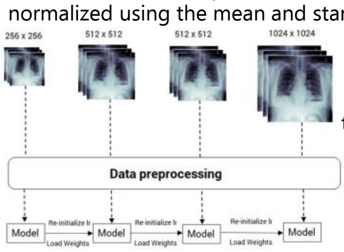

Figure 2 Data preprocessing of training process with model of UNet and EfficientNet-B0.

Interpretability and significance analysis of the model

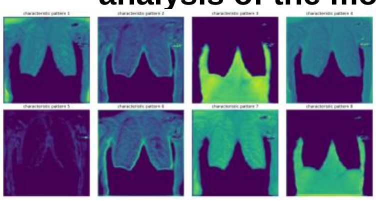

Figure 10 Intermediate feature maps from the final encoder block

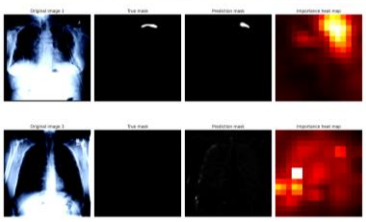

Figure 11 Occlusion sensitivity map visualization for representative samples

The Figure 10 visualization of feature attribution across layers. Different Grad-CAM patterns highlight attention areas across encoder stages, showing how the model captures pneumothorax-related regions progressively. The Figure 11 comparison of input image, ground-truth mask, predicted output, and corresponding importance heatmap, illustrating the model's focus and segmentation quality on pneumothorax cases and normal lungs.

Model Architecture

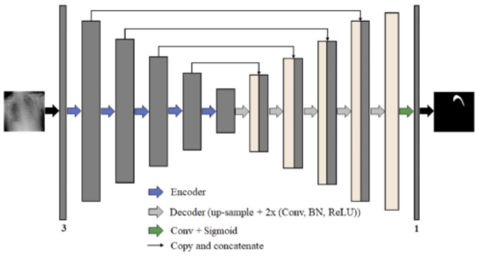

Figure 3 Architecture of U-Net

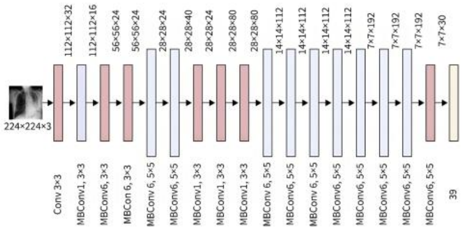

Figure 4 Common baseline architecture .

The proposed model adopts a U-Net architecture, known for its encoder decoder structure with skip connections, well suited for medical image segmentation. EfficientNetB0 serves as the encoder to extract rich, multi scale features with high efficiency, initialized using ImageNet pretrained weights to accelerate convergence.
The decoder restores spatial resolution through upsampling and convolutional layers, while skip connections ensure precise boundary localization and context retention critical for accurate pneumothorax segmentation.
To address class imbalance and small lesion challenges, a hybrid loss function combining Binary Cross Entropy (BCE) and Dice loss is used. The model produces a single-channel probability mask, with sigmoid activation applied during inference to yield interpretable outputs.

Model evaluation

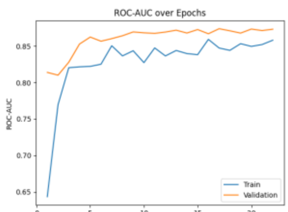

Figure 5 Training and validation ROC-AUC curves for U-Net with ResNet34 encoder

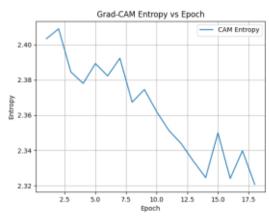

Figure 8 Grad-CAM over epochs for the optimized training setup with reduced layers

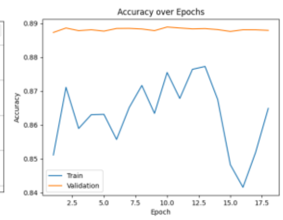

Figure 9 Accuracy over epochs for the optimized training setup with reduced layers.

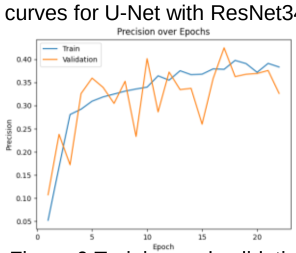

Figure 6 Training and validation precision curves for U-Net with ResNet34 encoder

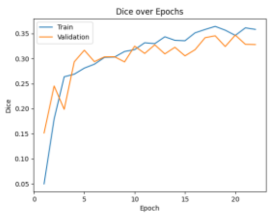

Figure 7 Training and validation Dice coefficient curves for U-Net with ResNet34 encoder.

The figure above shows the feedback obtained from the weighted integration of the three models using transfer learning, which was used to classify the fruit dataset for this experiment.

WEB APPLICATION DEVELOPMENT

Figure12 Initial view of the web-based pneumothorax prediction system, showing the file upload interface for chest X-ray images

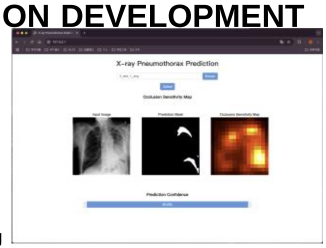

Figure13 Result visualization interface displaying the input chest X-ray, predicted pneumothorax segmentation mask, and prediction confidence score.

The browser handles UI only; model inference runs locally via a Flask backend for fast, offline predictions.

Future work

1.The model is prone to over-sensitivity, causing false positives that may burden clinicians; future work will optimize loss functions and apply post-processing to reduce unnecessary alerts. 2. Limited generalization to rare or complex cases restricts clinical reliability; improvements will include data diversification, domain adaptation, and uncertainty modeling for enhanced robustness.