Chest X-Ray Classification (Alveoli)

Workflow

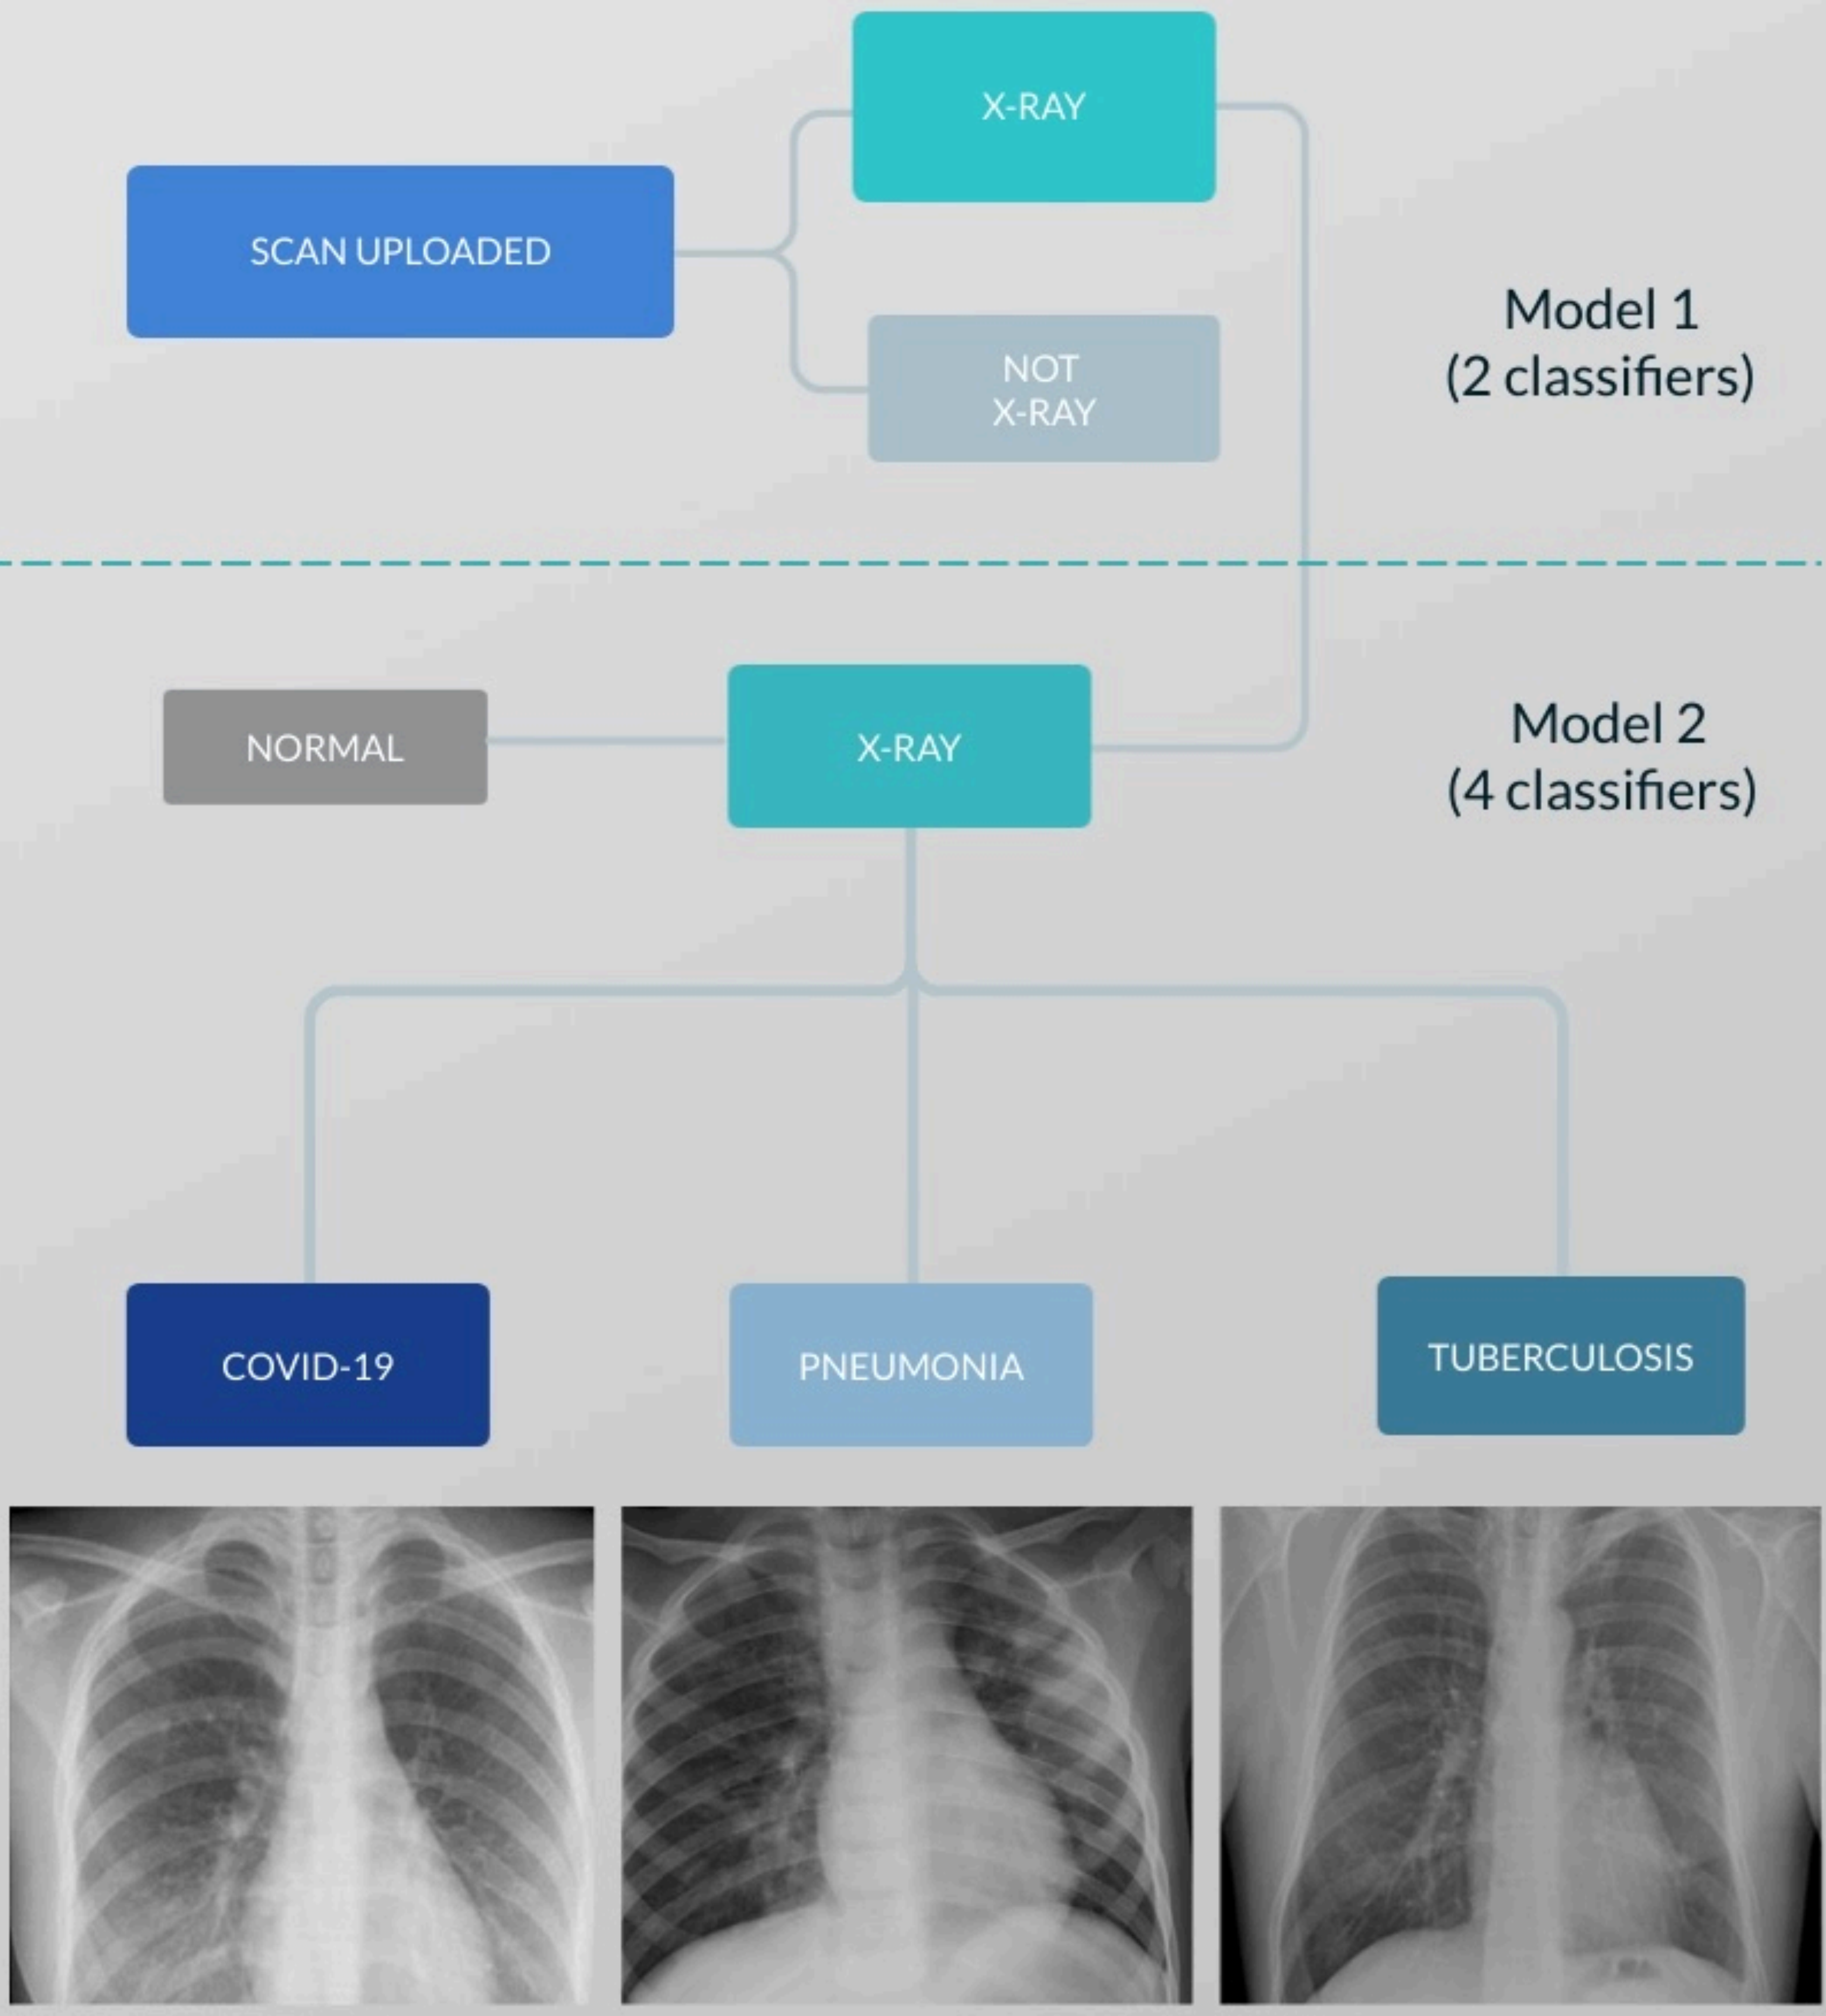
SCAN UPLOADED
X-RAY
NOT
X-RAY
Model 1
(2 classifiers)
NORMAL
X-RAY
Model 2
(4 classifiers)
COVID-19
PNEUMONIA
TUBERCULOSIS

Dataset

This project uses the Kaggle Chest X-Ray (Pneumonia,Covid-19,Tuberculosis) for the classification process in Model 2. The Kaggle Chest X-Ray (Pneumonia,Covid-19,Tuberculosis)dataset consists of 7135 chest X-ray images and is divided into four classes namely Covid-19,Normal,Pneumonia,Tuberculosis.The dataset is organised into 3 folders (train, test, val) and contains subfolders for each image category (Normal/Pneumonia/Covid-19/Tuberculosis).The model 1 to detect if an image is a chest X-Ray or not uses a customised dataset from Kaggle including pictures of Chest X-Ray and random items and screenshohts which could be stored in a phone.

Transfer Learning

Transfer learning is a widely used method where the knowledge acquired from one problem domain is transferred to another domain, with the aim of reducing the learning cost.Transfer learning can be used for fine-tuning the models where the initial layers of the model are frozen and the top-level layers are retrained. Hence, the generic features are extracted from the initial layers while more specific features are learned from the higher layers. In this project, we trained the selected MobileNetV2 model with transfer learning. In applying transfer learning to fine-tune the model, we first froze the entire model apart from the newly added top layers and trained it, and then we unfroze a few layers from the top of the frozen model and trained them together with the newly added layers. This helped in fine-tuning the generic features of the model to make them more applicable.

Model Architecture

The models used in this project are created using Deep Learning concepts.A convolutional neural network known as MobileNetV2 pre-trained on Image-net dataset has been used to perform classification on the CXR scans.This model improves the state of the art performance of mobile models on multiple tasks and benchmarks as well as across a spectrum of different model sizes. The MobileNetV2 architecture is based on an inverted residual structure where the input and output of the residual block are thin bottleneck layers opposite to traditional residual models which use expanded representations in the input an MobileNetV2 uses lightweight depth-wise convolutions to filter features in the intermediate expansion layer. Additionally the model removes non-linearities in the narrow layers in order to maintain representational power. This approach allows decoupling of the input, output domains from the expressiveness of the transformation, which provides a convenient framework for further analysis.

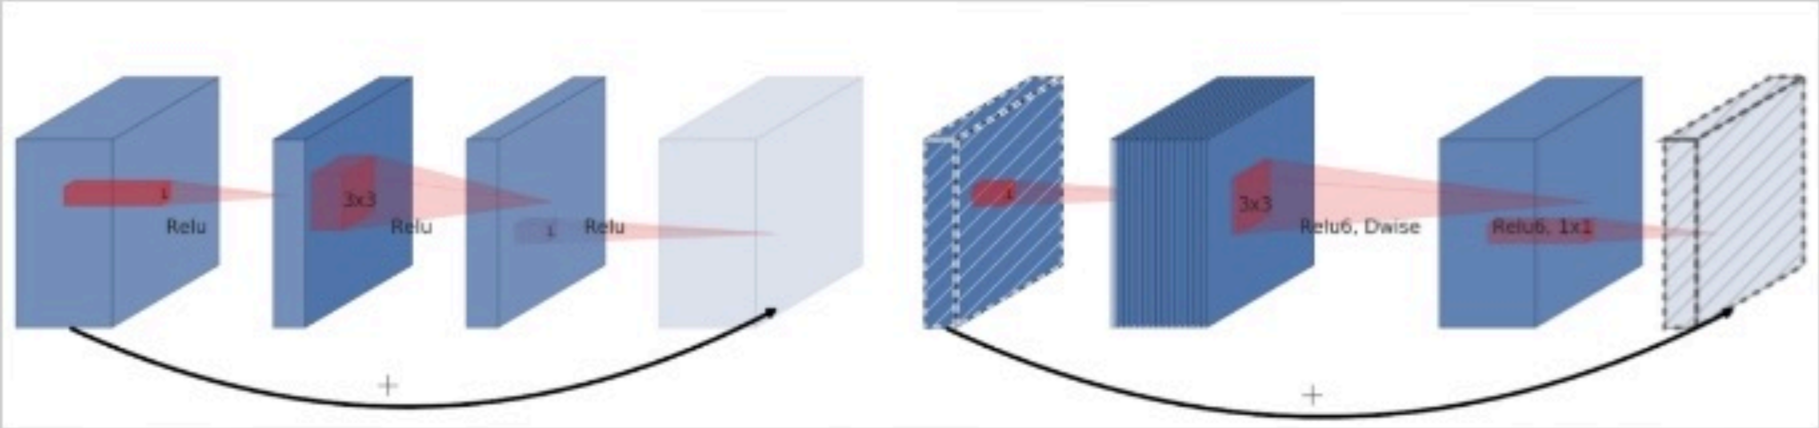

RESIDUAL BLOCK

INVERTED -RESIDUAL BLOCK

Model Training

The model was trained using transfer learning using pre-trained weights of the ImageNet dataset in the initial layers. A batch size of 64, Adam as the optimiser, and categorical cross-entropy as the loss function were selected. For the first model since binary classification is done the sigmoid activation function was used with loss set as binary_crossentropy. For the second model since there are 4 classes ,softmax activation was used in the last dense layer and loss is set as categorical_crossentropy. After setting the initial learning rate to 1e-3, for better convergence, learning rate reduction was used. The models were trained for 20 epochs. These hyper-parameters were selected by experimenting with several hyper-parameter combinations.

Results

The model gave a training accuracy of 95.26%(last epoch) and train accuracy of 94.79% and the test accuracy was around 93.22% .The F1 score for Covid-19 was 0.95, for Normal 0.83 ,for Pneumonia 0.91, for Tuberculosis 0.92 on the Test data.