


Pneumonia Detection Using CNN



Class: Applied Machine Learning (W207)

Team: Victor Inyang, Kodzai Nyakurimwa, Naveen Purushotham

WHAT IS PNEUMONIA?



- Pneumonia is an infection that inflames the air sacs in one or both lungs. The air sacs may fill with fluid or pus (purulent material)
- How do you confirm if you have pneumonia?
 - **Blood tests**, are used to confirm an infection and to try to identify the type of organism causing the infection.
 - **Chest X-ray**, to determine the extent and location of the infection.
 - **Pulse oximetry**, this measure of oxygen level in the blood
 - **Sputum test**, using a sample of fluid from the lungs (sputum)



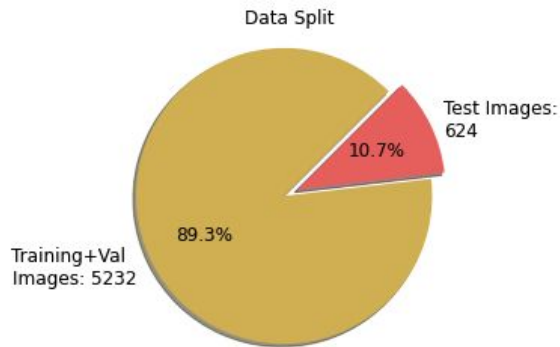

RESEARCH QUESTION

Can we accurately diagnose Pneumonia using CNNs given chest X-ray data?

- This is an exciting question because if we are able to detect or diagnose pneumonia from X-rays with acceptable accuracy, we see the endless application of machine learning being used in the medical field to assist doctors in accurately determining what's wrong with patients.
- These algorithms can be used in remote areas where medical specialists are not available to diagnose patients by reviewing their images through AI.
- Benefits to the patients include the following;
 - Consistent and faster diagnosis of their health
 - Reduced cost (eliminates specialist fees and travel to big cities with specialists)


DATASET

- Our dataset is from Kaggle which includes chest X- rays of normal and pneumonia-infected lungs
- The dataset has 5,856 chest X-ray images and is divided into train, val, and test
- Combined train and val, then split the combined total to 80:20 (train:val)
- We have observed some data imbalance for labeled dataset, normal 27% and pneumonia 73%
- Training data was resampling to address imbalance by matching the number of pneumonia labels with normal labels




DATASET AUGMENTATION

- We are augmenting the images using *imagedatagenerator*, keras image preprocessor
- We are using a different transformation of the original image in each epoch
- Total number of images does not change after Augmentation
- This is meant to reduce bias and bring fairness, so as to accommodate images that might not have been perfectly straight, imperfect zoom, incorrect brightness setting etc.




Original Image




Epoch 1
Transformation




Epoch 2
Transformation




Epoch 3
Transformation



Epoch 4
Transformation




Epoch 5
Transformation


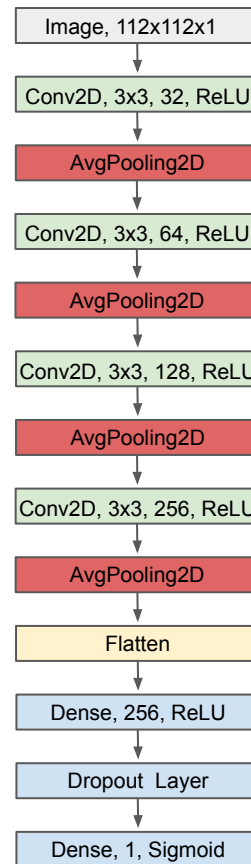




Epoch 6
Transformation

Model Exploration



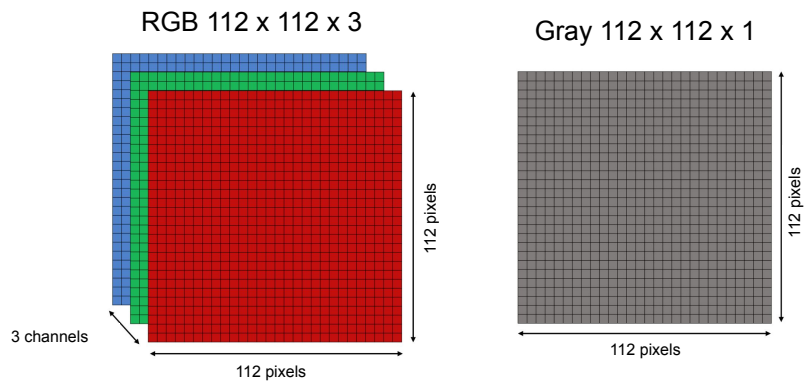
FFNN



CNNs


Experiments

- Different (input shapes - rgb and grayscale)




- Different methods to treat imbalance

- Random Sampling




- Class weights



Weight for class Pneumonia: 0.67
Weight for class Normal: 1.94


Experiments

- Max Pooling vs Average Pooling




Original Image

Max Pooling





Average Pooling





Experiments

- Raw data (resized)




- Augmented data




- Different Fill modes for Augmentation

- Nearest




- Constant




Experiments

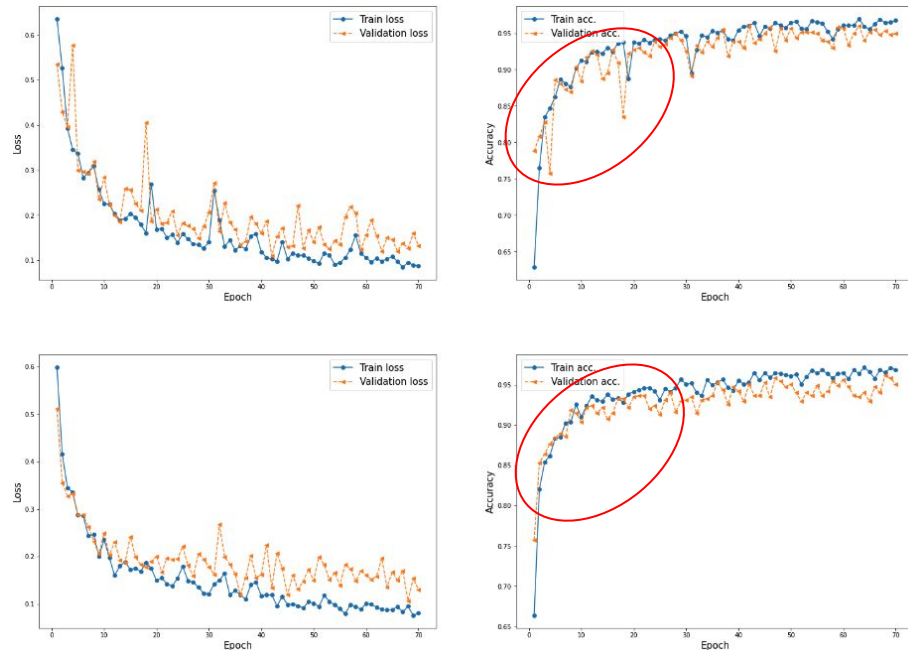
- Epochs




- No Validation Data Augmentation



- Image shuffling after each batch





CNN Baseline Before Tuning



Model: "sequential"

Layer (type)	Output Shape	Param #
conv_1 (Conv2D)	(2146, 112, 112, 32)	320
pool_1 (MaxPooling2D)	(2146, 56, 56, 32)	0
conv_2 (Conv2D)	(2146, 56, 56, 64)	18496
pool_2 (MaxPooling2D)	(2146, 28, 28, 64)	0
conv_3 (Conv2D)	(2146, 28, 28, 128)	73856
pool_3 (MaxPooling2D)	(2146, 14, 14, 128)	0
conv_4 (Conv2D)	(2146, 14, 14, 256)	295168
pool_4 (MaxPooling2D)	(2146, 7, 7, 256)	0
flatten (Flatten)	(2146, 12544)	0
fc_1 (Dense)	(2146, 256)	3211520
dropout (Dropout)	(2146, 256)	0
fc_2 (Dense)	(2146, 1)	257

=====
Total params: 3,599,617
Trainable params: 3,599,617
Non-trainable params: 0



Optimizer and Loss function

- We used Adam optimizer for the final implementation.
- The loss (cost) function suitable for our binary classification model is `binary_crossentropy`.

Evaluation

- **Test Accuracy 90%**

Hyperparameter Tuning

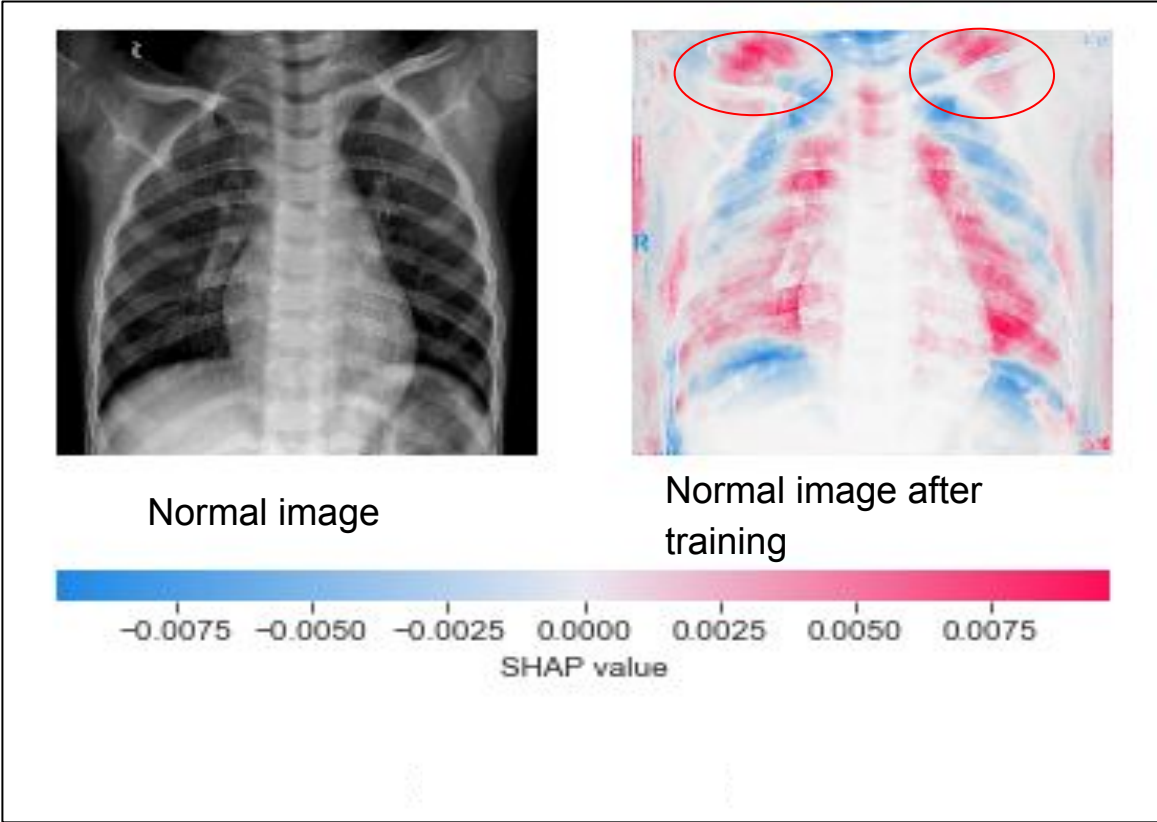
	Model	Training accuracy	Validation accuracy	Test accuracy	kernel size	strides	pool size	learning rate	optimizer	brightness_range	zoom_range	horizontal_flip
CNN baseline	0	0.96	0.95	0.90	3,3	1,1	2,2	0.001	Adam	None	0.2	True
	1	0.90	0.89	0.89	3,3	1,1	2,2	0.01	Adam	None	0.2	True
	2	0.93	0.92	0.89	3,3	2,2	2,2	0.001	Adam	None	0.2	True
	3	0.95	0.94	0.90	5,5	1,1	2,2	0.001	Adam	None	0.2	True
	4	0.96	0.95	0.87	3,3	1,1	3,3	0.001	Adam	None	0.2	True
	5	0.96	0.95	0.88	3,3	1,1	4,4	0.001	Adam	None	0.2	True
	6	0.69	0.50	0.37	3,3	1,1	2,2	0.001	Adam	(0.1,0.3)	0.2	True
	7	0.95	0.94	0.89	3,3	1,1	2,2	0.001	Adam	None	0.4	True
★	8	0.97	0.95	0.91	3,3	1,1	2,2	0.001	Adam	None	0.2	False
	9	0.92	0.92	0.85	3,3	1,1	2,2	0.001	SGD	None	0.2	True

- Strides
- Learning rate
- Optimizer

- Brightness_range
- Zoom_range
- Horizontal_flip

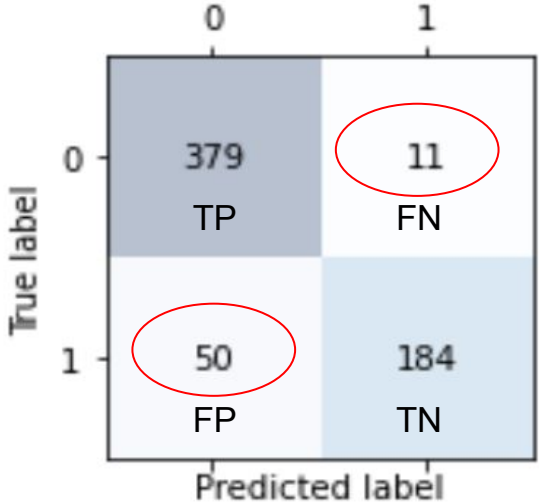
- Pool size

Model Explainability



Key Metrics

No of Misclassified label	61
Accuracy	0.91
Precision	0.88
Recall	0.97



Confusion Matrix

Conclusions

- With a 97% recall rate, we believe that ML algorithms (CNNs) can accurately diagnose pneumonia.
- 9.7% of test data was misclassified, given more time we'd like to tweak and tune the model to increase the accuracy
- Data was only collected on children with age below 5, out of the children we do not know what percent of boys or girls were included in the data
- Moving forward we'll make sure different training data accounts for amount for each gender, age group and other geographical areas are equally represented and be part of the data collection process with radiologist resources.
- Moving forward we'll use a GPU system to train our model and fine tune on our hyper-parameter we'll also incorporate transfer learning

References

- Kaggle.com
 - <https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia>
 - <https://www.youtube.com/watch?v=5IBnuauBvkQ>
 - <https://www.kaggle.com/code/madz2000/pneumonia-detection-using-cnn-92-6-accuracy>
- [Gender Rib Cage](https://pubmed.ncbi.nlm.nih.gov/16476656/) - <https://pubmed.ncbi.nlm.nih.gov/16476656/>
- [Misdiagnosis](https://pubmed.ncbi.nlm.nih.gov/32723999/) - <https://pubmed.ncbi.nlm.nih.gov/32723999/>

Thank You

Any Questions or Feedback?

Contributions/Primary Areas of Focus

	Victor Inyang	Kodzai Nyakurimwa	Naveen Purushotham
Theoretical Research	X	X (Lead)	X
Environment set-up	X	X	X
Data Cleaning	X	X	X
Data splitting	X	X	X
Hyperparameter tuning / Metrics	X (Lead)	X	X
Augmentation	X	X	X (Lead)
Presentation slides	X	X	X