

The Effectiveness of Deep Learning Models on the Detection of Pneumonia

Abstract

The purpose of this experiment was to find out if deep learning models accurately diagnose patients with pneumonia. If deep learning models were used with Convolution Neural Network layers and pooling techniques, then the efficiency of detection would increase with max pooling technique and the most amount of CNN layers. If deep learning models were used with Convolution Neural Network layers, then the amount of time the model takes to run would increase with the increase of CNN layers. A model was created with three convolution neural network (CNN) layers and max pooling technique. The number of CNN layers was increased and the pooling technique changed to average pooling technique. The models were trained, validated, and tested using normal and pneumonia x-ray images. Overall, the experiment supported the hypothesis. For both pooling techniques, the increase of CNN layers increased the accuracy and the time it took for the model to run. In addition, the most layers and max pooling lead to a higher accuracy of around 74 percent while most layers and average pooling lead to an accuracy of around 49 percent. Different deep learning models using max pooling could be applied to other lung conditions where patients come in with bronchitis or pulmonary edema. Validated models could also be used in other areas of the body to help diagnose patients that have pain.

Introduction

The purpose of this experiment is to find out if deep learning models accurately diagnose patients with pneumonia.

Pneumonia is an infection that inflames the air sacs in one or both lungs. The air sacs are filled with fluid, causing coughing with phlegm, fever, chills, and difficulty breathing. A variety of organisms, including bacteria, viruses and fungi, can cause pneumonia (Staff M.C., 2020). Viral and bacterial pneumonia can spread from one person to another from inhaling, sneezing, and coughing. In addition, people can come into contact with pneumonia by surfaces or objects. On the other hand, fungal pneumonia, not able to be spread from person to person, can be contacted from the environment (Pneumonia, 2020). Around 1.3 million people in the United States were admitted to the hospital, and around 50,000 deaths occurred. It can range in seriousness from mild to life-threatening. It is most serious for infants and young children, people older than age 65, and people with health problems or weakened immune systems (National Center for Health Statistics, 2020).

X-rays are a form of radiation-like light or radio waves-that can be directed at the body. Different body parts absorb the X-rays to varying degrees. In a conventional X-ray exam, a small burst of radiation is aimed at and passes through the body, recording an image on photographic film or a special image recording plate. As the radiation wave passes through the body, the tissues and bones either block or absorb the radiation depending on the density. X-ray images are black, grey, and white where the lighter color represents a denser material. On a normal chest X-ray, the bones will appear white because it is hard and mineralized while air appears to be black. Soft tissue and muscles will be grey on the image. X-ray images are very helpful for Pneumonia patients and detecting fluid in the lungs (X-ray (roentgenogram), 2018). Chest X-rays

of Pneumonia patients are first used to put critical patients in the emergency rooms. Then, using the X-rays, the different portions of the lungs are examined to figure out which treatment would be best for the patients. In the image, the lungs should be black while the ribs are white. Pneumonia will cause patients to build up fluid in their lungs making the black lungs to have white streaks (Draelos, 2019) .

Machine learning refers to when a computer learns through a large number of data sets instead of coded rules. Deep learning is a type of machine learning where an Artificial Intelligence (AI) is trained to give an output for an input. It mimics the workings of the neurons of the human brain. It uses observations such as images, texts, or sounds as an input (Bonner, 2019). The machine will process the data using Convolutional Neural Networks (CNN). CNN is made up of different neural layers where each layer has a different function. The layers represent the neurons in the human brain. Like the brain, the information will have to go through all the layers to be processed. As the information is processed, the layers extract different features that will help give an accurate conclusion without the help of a human. The hierarchical AI layers of deep learning enables it to process data in a non linear manner. The output will be the object that the machine is supposed to detect (Chatterjee, 2019). Deep learning takes up alot of memory because the image needs to be broken down to give an accurate result. However, the more the image is broken down, the more memory is needed. This is a huge problem because it restricts research to be done on neural networks (Hanlon, 2016).

Rectified Linear Activation Function (ReLU) is a piecewise nonlinear function that is the default activation when working with CNN which allows for deep neural networks. ReLU acts like a linear piecewise function even though it is nonlinear. If ReLU was not used in the code, then the outcome will end up with many data numbers that would not make sense in the pooling

technique. The model will exit the code giving the user an error message. After the input runs through the CNN the ReLu will then sort the numbers. If the number is negative or zero, the output will be zero. However, if the number is positive, the number will stay the same. The output numbers will then be used by the pooling techniques. Relu helps make training for the models easier and quicker. Using ReLu will result in a better outcome (Brownlee, 2019).

The CNN layers record precise features in the input that the model is looking for which result in a different feature map. However, features are not useful to determine the output are put aside which can interfere with the signaling process. Down sampling is used where the large or important structural elements are kept for higher accuracy while the smaller and less important features are sorted out of the process to keep the programming running quickly and smoothly. A common approach is pooling techniques. A pooling layer is a new layer added after the convolutional layer; specifically, after ReLu has been applied to the feature maps output by a convolutional layer. Two common pooling methods are average pooling and max pooling. In both pooling techniques, the image pixel numbers are put into different matrices. Average pooling technique takes the average of each matrix while max pooling technique takes the maximum number. After, the image is reduced to be run again through a CNN where there will be a smaller amount of pixel numbers to be sorted by pooling (Brownlee, 2020).

If deep learning models are used with Convolution Neural Network layers and pooling techniques, then the efficiency of detection will increase with max pooling technique and the most amount of CNN layers. If deep learning models are used with Convolution Neural Network layers, then the amount of time the model takes to run will increase with the increase of CNN layers.

Methods

Materials

- Python version 3.7
- MacBook pro with 32GB of memory
- Pycharm Integrated Development Environment version 2020.2
- Google tensorflow library version 2.3
- 500 x-ray images of normal lungs from Kaggle data set
- 500 x-ray images of pneumonia lungs from Kaggle data set
- 20 additional x-ray images of normal lung for negative control
- 20 additional x-ray images of pneumonia lungs for positive control
- Microsoft Excel from Office 365
- Jmp Statistical Package

Procedures

1. Download all normal and pneumonia x-ray images from Kaggle.
2. Create a model using 3 convolution neural networks (CNN) layer and ReLu using python and Google tensorflow in pycharm.
3. Add max pooling technique to the CNN.
4. Train the model by inputting 350 pneumonia x-ray images and 350 normal x-ray images.
5. Validate the model by inserting the rest of the 150 pneumonia and 150 normal x-ray images into the model.
6. Record the validation accuracy of the model and the time it took for the model to run.
7. Do steps 1-7 using max pooling and ReLu, for 5, 7, and 9 layers in the CNN.
8. Download all normal and pneumonia x-ray images from Kaggle.

9. Create another model using 3 convolution neural networks (CNN) layer and ReLu using python and google tensorflow in pycharm.
10. Add average pooling technique to CNN.
11. Train the model by inputting 350 pneumonia x-ray images and 350 normal x-ray images.
12. Validate the model by inserting the rest of the 150 pneumonia and 150 normal x-ray images into the model.
13. Record the validation accuracy of the model and the time it took for the model to run.
14. Do steps 9-15, for average technique pooling and ReLu, for 5, 7, and 9 layers in the CNN.
15. Test the model using the 20 positive controls and the 20 negative controls on the best deep learning model.
16. Record the binary outcome.
17. Analyse the data in Excel and Jmp Statistical Package.

Results

Table

Validating Accuracy and Time for Max and Average Pooling					
Trials	layer	Max pooling		Average pooling	
		Time(s)	validation accuracy	Time(s)	validation accuracy
1	3	251	0.5303	386	0.3687
	5	343	0.5454	457	0.3945
	7	446	0.6363	548	0.4523
	9	649	0.7424	754	0.4848
2	3	258	0.5313	388	0.36
	5	350	0.546	462	0.3989
	7	500	0.637	549	0.46
	9	750	0.7432	755	0.489
3	3	257	0.5308	389	0.37
	5	349	0.5469	461	0.396
	7	501	0.6378	529	0.455
	9	740	0.74	759	0.4888
4	3	256	0.5311	389	0.369
	5	361	0.5472	467	0.4
	7	498	0.6372	500	0.455
	9	721	0.7414	753	0.49
5	3	255	0.5303	390	0.3689
	5	369	0.5454	451	0.3945
	7	500	0.6363	501	0.4523
	9	714	0.7424	754	0.4848

6	3	259	0.5356	350	0.37
	5	342	0.5479	452	0.45
	7	501	0.639	541	0.49
	9	722	0.75	752	0.51
7	3	261	0.539	350	0.3691
	5	341	0.549	457	0.39
	7	489	0.639	542	0.4234
	9	702	0.749	755	0.4856
8	3	278	0.53	380	0.3695
	5	380	0.5412	450	0.3989
	7	520	0.632	540	0.45
	9	735	0.7481	750	0.48
9	3	257	0.53	390	0.3687
	5	368	0.5412	412	0.3945
	7	511	0.6382	550	0.4523
	9	723	0.7412	762	0.4848
10	3	259	0.5389	387	0.3693
	5	378	0.55	456	0.399
	7	501	0.6401	549	0.4578
	9	801	0.7498	754	0.489

Table 1. Showing all the simulation data for the different CNN layers and pooling techniques.

Binary Output for X-Ray Images on Max Pooling Model			
Normal Chest X-Ray		Pneumonia Chest X-Ray	
Image	Output (Y/N)	Image	Output (Y/N)
1	Y	1	Y
2	N	2	N
3	Y	3	Y
4	Y	4	Y
5	Y	5	N
6	N	6	N
7	N	7	Y
8	Y	8	N
9	N	9	Y
10	N	10	Y
11	N	11	N
12	N	12	Y
13	Y	13	N
14	N	14	Y
15	Y	15	Y
16	Y	16	Y
17	Y	17	Y
18	N	18	N
19	Y	19	Y
20	N	20	Y

Of the 20 pneumonia chest x ray only 13 showed positive (65%), and out of the 20 normal images only 10 (50%) showed normal.

Table 2. Showing the binary outcome, yes or no, for both the normal and pneumonia on the max pooling model with 9 layers.

Graphs

Variability graph for Validation accuracy




Fig 3. Variability graph showing how the validation accuracy increases with the number of CNN layers and with max pooling. The validation accuracy is better for max pooling than for average pooling.

Variability graph for Validation accuracy




Fig 4. Standard deviation of variability graph showing that for max pooling the standard deviation value is pretty constant at about .005, whereas for average pooling it is erratic and varies from less than .005 to about .0155.

Variability graph for Pooling-time(s)




Fig 1. Variability graph showing how the computation time increases with the number of CNN layers and with max and average pooling both.

Variability graph for Pooling-time(s)




Fig 2. Variability graph showing standard deviation of pooling times with respect to CNN layers and pooling techniques. Average pooling seems to have a tighter standard deviation than max pooling.

Photographs




Fig. 5 This picture represents a normal chest x-ray.

It was taken from a website called kaggle.



Fig. 6 This picture represents a pneumonia chest

x-ray. It was taken from a website called kaggle.

```
Image_Width = 150
Image_Height = 150
Channels = 3

INPUT_SHAPE = (Image_Width, Image_Height, Channels)
NB_CLASSES = 2
EPOCHS = 10
BATCH_SIZE = 6

t1 = datetime.datetime.now()
model = Sequential()

model.add(Conv2D(32, (3, 3), input_shape=INPUT_SHAPE))
model.add(Activation('relu'))
model.add(AvgPool2D(pool_size=(2, 2)))

model.add(Conv2D(64, (3, 3)))
model.add(Activation("relu"))

model.add(Conv2D(250, (3, 3)))
model.add(Activation("relu"))
model.add(AvgPool2D(2, 2))
```

Fig 7. This picture represents the average pooling

model with three layers. It was taken on December 2, 2020.

```

t1 = datetime.datetime.now()
model = Sequential()

model.add(Conv2D(32, (3, 3), input_shape=INPUT_SHAPE))
model.add(Activation('relu'))
model.add(AvgPool2D(pool_size=(2, 2)))

model.add(Conv2D(64, (3, 3)))
model.add(Activation('relu'))

model.add(Conv2D(256, (3, 3)))
model.add(Activation('relu'))
model.add(AvgPool2D(2, 2))

model.add(Conv2D(128, (3, 3)))
model.add(Activation('relu'))

model.add(Conv2D(64, (3, 3)))
model.add(Activation('relu'))
model.add(AvgPool2D(2, 2))

model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
model.add(AvgPool2D(2, 2))
#
model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
model.add(AvgPool2D(2, 2))
#
model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
# # model.add(AvgPool2D(2, 2))
#
model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
model.add(AvgPool2D(2, 2))

```

Fig 8. This picture represents the average pooling model with nine layers. It was taken on December 10, 2020.

```

Image_Width = 150
Image_Height = 150
Channels = 3

INPUT_SHAPE = (Image_Width, Image_Height, Channels)
NB_CLASSES = 2
EPOCHS = 10
BATCH_SIZE = 6

t1 = datetime.datetime.now()
model = Sequential()

model.add(Conv2D(32, (3, 3), input_shape=INPUT_SHAPE))
model.add(Activation('relu'))
model.add(MaxPool2D(pool_size=(2, 2)))

model.add(Conv2D(64, (3, 3)))
model.add(Activation("relu"))

model.add(Conv2D(256, (3, 3)))
model.add(Activation("relu"))
#model.add(MaxPool2D(2, 2))

```

Fig 9. This picture represents the max pooling model with three layers. It was taken on December 10, 2020.

```

model.add(Conv2D(32, (3, 3), input_shape=INPUT_SHAPE))
model.add(Activation('relu'))
model.add(MaxPool2D(pool_size=(2, 2)))

model.add(Conv2D(64, (3, 3)))
model.add(Activation('relu'))

model.add(Conv2D(256, (3, 3)))
model.add(Activation('relu'))
#model.add(MaxPool2D(2, 2))

model.add(Conv2D(128, (3, 3)))
model.add(Activation('relu'))

model.add(Conv2D(64, (3, 3)))
model.add(Activation('relu'))
model.add(MaxPool2D(2, 2))

model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
#model.add(MaxPool2D(2, 2))
#
model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
model.add(MaxPool2D(2, 2))
#
model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
# # model.add(MaxPool2D(2, 2))
#
model.add(Conv2D(256, (2, 2)))
model.add(Activation('relu'))
model.add(MaxPool2D(2, 2))

```

Fig 10. This picture represents the max pooling model with nine layers. It was taken on December 10, 2020.

Discussion

Conclusion

Table 1 shows the data for all the simulations using 3,5,7 and 9 layers of CNN with max pooling and average pooling.

Hypothesis 1, which states that if the number of CNN layers is increased in the model with max pooling then the validation accuracy will increase is clearly supported by the data in table 1 and figures 3 and 4. Figure 3 shows that as the number of layers increases, the validation accuracy increases both for max pooling and for average pooling. However the validation accuracy for max pooling increases more than that for average pooling. Figure 4 which shows the standard deviation of the data reveals that the variability in max pooling is less than that of average pooling and pretty constant.

Hypothesis 2, which states that if the number of CNN layers would increase then the time of computation would increase, both for max pooling and average pooling. This conclusion is clearly shown in figures 1 and 2. Pooling layer generally operates upon each feature map separately to create a new set of the same number of pooled feature maps. Pooling involves selecting a pooling operation, much like a filter to be applied to feature maps. The size of the pooling operation or filters is smaller than the size of the feature map, specifically it is always 2x2 pixels applied with a stride of 2 pixels. Average Pooling calculates the average value for each patch of the feature map, while Maximum Pooling calculates the maximum value for each patch of the feature map. Calculating the average value for each patch seems to take longer than that of calculating the maximum patch size.

Table 2 shows the simulation of a binary outcome on chest x-ray images for the max pooling model with 9 CNN layers.

Using the best model created with 9 CNN layers and max pooling, a binary test was done to evaluate if the model could do an accurate prediction. 20 pneumonia and 20 normal chest x ray was fed to the model and the result obtained was not about to the same accuracy as the model validation, about 70%. Of the 20 pneumonia chest x ray only 13 showed positive (65%), and out of the 20 normal images only 10 (50%) showed normal. Even though this experiment has shown that deep learning models can make predictions in diagnosis of diseases in patients, the model in this experiment is not recommended to be used clinically. As seen the best accuracy of 9 layers of CNN with an epoch of 10 and a batch size of 6 yielded an accuracy of only 74%, this is suboptimal in clinical use. For better prediction accuracy closer to 100%, much better datasets and more robust computers should be used to train and validate the models. Generally about 100 to 500 epochs are used in real deep learning models and for this type of computation an ordinary computer would not be able to train and validate the models.

Application

This research has tremendous applications in all subfields of medicine starting from diagnosis to treatment follow up and even prevention. When diagnosis is concerned in radiology and pathology, physicians can use the deep learning models to help them clinically on a daily basis. Some of these models could possibly find pneumonia and many other diseases in people in much shorter time than it takes today with only the physician reading all the images and interpreting it subjectively. Without deep learning, doctors would have to use their best judgment to diagnose patients with pneumonia and or other conditions where images are used in diagnosis. However, this could lead to a misdiagnosis that could cost a patient his or her life. Deep learning models can certainly help doctors by doing an early triage and hence help reduce the load of physicians. Deep learning models, not only pneumonia, can be used for other lung conditions

such as bronchitis and covid-19. If a patient comes in with pain in their chest, the models, learning from previous chest x-rays, will be able to aid doctors in diagnosing. With more experiments and research, deep learning can broaden to other parts of the human body. Chest x-rays are not the only input that can be used, MRI and ultrasound are also alternatives. X-rays from abdominal and dental sites can also be used to help doctors diagnose patients that come in with pain in their mouth or abdominal area. Thus this kind of research can expand to the field of pathology where pathologists have to spend hours looking at slides under the microscope to make a diagnosis of surgical specimens from biopsy especially in cancer diagnosis. An earlier diagnosis can lead to a higher probability of local control of the cancer and probably a cure. Patients undergoing radiation treatment generally use PET scans to see if the cancer has shrunk. Here deep learning models can provide cancer care follow up on a personal level and inform the treating physician if they need to modify their treatment plan.

Limitations

One of the main limitations encountered when running the simulation was that the program would crash once in a while because running the program would take up a lot of memory when the number of CNN layers would increase. Regular home computers generally do not have that kind of resources. This limits the usefulness of the model because there would be no point to it if it crashed when analyzing a patient's x-ray. More sophisticated computers with lots of memory and preferably with Graphical Processing Unit (GPU) are needed to train and validate these models. The x-ray images that were used to train and validate the model had different resolutions and sizes which made it harder for the model to learn from the x-rays and to give a high accuracy when validating it. Moreover it is not known whether the X-ray machine that these images were taken from were calibrated using the same protocol and used the same

technique used to acquire the images. Image quality is very much dependent upon the calibrated X-ray machine and technique of acquisition, making it a very legitimate limitation. In addition, this model only used images with no implants in the chest area. However, some pneumonia patients can come in with a pacemaker, stents or breast implant which would invalidate the use of these models. In these cases, the doctors would have to read the x-ray which would make the diagnosis subjective and may result in an improper diagnosis as the implants would cause an interference in the image. More research would have to be done to manipulate the code to read the x-ray images with different kinds of implants. The images could vary immensely which should be investigated to properly find models to accommodate the artifacts. These models are only useful to patients who have pneumonia and no implants to interfere with the reading.

Error Analysis

The computer used in this experiment could introduce computational error as this is a simulation kind of experiment where the final result is based on some kind of accuracy loss. There was not enough memory to run the program, so the result could be inaccurate. As the computer started running the same model many times and as the temperature of the CPU increased, the computer started heating up and may introduce computational errors in the computation. There was a delay in the time because the program would run through the code quickly but slow down as the images were processed many times. This resulted in fluctuations in the timing. In addition, the x-ray images were not of the same size and had rotational bias including tilt and zoom in them. The model training and validation from these x-ray images could have improperly observed the very object being measured. As the model was validated to give an accuracy of how well the model did to process the images, the accuracy could be inaccurate.

Future Research

In the field of deep learning with convolutional neural networks there are numerous parameters called hyper parameters called pooling and activation with variable inputs to them may affect the results of the model. For this experiment, two types of pooling techniques, max and average, and activation hyper parameter called ReLu were used. However, there are many other pooling such as min, tree, and wavelet. There are also many other activation functions such as sigmoid, tanh, and softmax. This could make a difference in the outcome leading to either a higher or a lower accuracy and faster or slower time, resulting in an accurate or inaccurate diagnosis of pneumonia for the patient. In addition, this experiment mainly focused on using deep learning for pneumonia. Having deep learning for other conditions of different body parts such as cancer in the abdominal area where there are other organs can give a different outcome. Patients can have implants which can interfere with the reading of the x-ray by the model, so more research will have to be done to allow models to read the x-ray images without having a problem with the implant. Different numbers of CNN layers should also be included at the same time a variety of areas and implant interference are being tested to record any abnormal accuracies and timing that could harm the patients. Creating a model with more epochs can result in higher accuracy because the data will be looped in more thoroughly. For this research to be done, a computer with a tremendous amount of memory and computational power will have to be used to get a reliable model.

References

- Bonner, A. (2019, September 7). *Top highlight What is Deep Learning and How Does it Work?* Towards Data Science. Retrieved October 12, 2020, from <https://towardsdatascience.com/what-is-deep-learning-and-how-does-it-work-f7d02aa9d477>
- Brownlee, J. (2019, January 9). *A Gentle Introduction to the Rectified Linear Unit (ReLU)*. Machine Learning Mastery. Retrieved September 29, 2020, from <https://machinelearningmastery.com/rectified-linear-activation-function-for-deep-learning-neural-networks/>
- Brownlee, J. (2019, July 5). *A Gentle Introduction to Pooling Layers for Convolutional Neural Networks*. Machine Learning Mastery. Retrieved October 7, 2020, from <https://machinelearningmastery.com/pooling-layers-for-convolutional-neural-networks/>
- Chatterjee, C. C. (2019, July 31). *Basics of the Classic CNN*. Towards Data Science. Retrieved September 23, 2020, from <https://towardsdatascience.com/basics-of-the-classic-cnn-a3dce1225add>
- Chest X-ray Images (Pneumonia)* [Photograph]. (2017). Kaggle. <https://www.kaggle.com/paultimothymooney/chest-xray-pneumonia>
- Draelos, R. (2019, February 10). *Radiology: Normal Chest X-Rays*. Glass Box Medicine. Retrieved September 12, 2020, from <https://glassboxmedicine.com/2019/02/10/radiology-normal-chest-x-rays/#:~:text=Darker%20colors%20indicate%20less%20dense,and%20block%20x%2Drays%20effectively>
- Hanlon, J. (2016). *WHY IS SO MUCH MEMORY NEEDED FOR DEEP NEURAL NETWORKS?* Graphcore. Retrieved October 7, 2020, from

<https://www.google.com/search?q=google+slides&oq=google+slides&aqs=chrome.0.69i59j0i433l2j0l3j69i60l2.1498j0j1&sourceid=chrome&ie=UTF-8>

National Center for Health Statistics. (2020). CDC. Retrieved October 6, 2020, from

<https://www.cdc.gov/nchs/fastats/pneumonia.htm>

Pneumonia. (2020, October 1). Medicine Plus. Retrieved October 5, 2020, from

<https://medlineplus.gov/pneumonia.html>

Society for Science and the Public (2020-21). International Science and Engineering Fair

2020-21: International Rules & Guidelines. Washington, DC: Society for Science and the Public.

Staff, M. C. (2020, June 13). *Pneumonia*. Mayo Clinic. Retrieved October 2, 2020, from

<https://www.mayoclinic.org/diseases-conditions/pneumonia/symptoms-causes/syc-20354204#:~:text=Pneumonia%20is%20an%20infection%20that,and%20fungi%2C%20can%20cause%20pneumonia.>

X-ray (roentgenogram). (2018). May Field Clinic. Retrieved October 2, 2020, from

<http://www.mayfieldclinic.com/pe-x-ray.htm>