

CoronaNet: A Novel Deep Learning Model for COVID-19 Detection in CT Scans

Rohan Bhansali¹ and Rahul Kumar¹

Duke Writer, Loudon Academy of Science, Leesburg, VA, USA

ABSTRACT

Coronavirus disease (COVID-19) is currently the cause of a global pandemic that is affecting millions of people around the world. Inadequate testing resources have resulted in several people going undiagnosed and consequently untreated; however, using computerized tomography (CT) scans for diagnosis is an alternative to bypass this limitation. Unfortunately, CT scan analysis is time-consuming and labor intensive and rendering is generally infeasible in most diagnosis situations. In order to alleviate this problem, previous studies have utilized multiple deep learning techniques to analyze biomedical images such as CT scans. Specifically, convolutional neural networks (CNNs) have been shown to provide medical diagnosis with a high degree of accuracy. A common issue in the training of CNNs for biomedical applications is the requirement of large datasets. In this paper, we propose the use of affine transformations to artificially magnify the size of our dataset. Additionally, we propose the use of the Laplace filter to increase feature detection in CT scan analysis. We then feed the preprocessed images to a novel deep CNN structure: CoronaNet. We find that the use of the Laplace filter significantly increases the performance of CoronaNet across all metrics. Additionally, we find that affine transformations successfully magnify the dataset without resulting in high degrees of overfitting. Specifically, we achieved an accuracy of 92% and an F1 of 0.8735. Our novel research describes the potential of the Laplace filter to significantly increase deep CNN performance in biomedical applications such as COVID-19 diagnosis.

Introduction

Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) has been identified as the virus causing coronavirus disease (COVID-19) (Lai, C. C., Shih, T. P., Ko, W. C., Tang, H. J., & Hsueh, P. R., 2020). It is thought to have originated in the city of Wuhan, China. The disease has caused the World Health Organization to declare a global public health emergency and officially classify the disease as the cause of a pandemic. Similar to previous coronaviruses, SARS-CoV-2 is originally found in bats, and is generally consistent across different regions. COVID-19 can cause mild to severe illness that if becomes more severe, can cause pneumonia, organ failure, and death. Symptoms of the disease include fever, dry cough, shortness of breath, and fatigue. The disease spreads through respiratory particles which are produced when an infected person either sneezes or coughs. These particles must be taken into the body through the nose, mouth, or eyes. COVID-19 has spread globally to affect at least 2 million people in 184 countries at the time of publication. The worldwide extent of COVID-19 has resulted in over 130,000 deaths. Those at risk belong primarily to older demographics with historically compromised immune systems or pre-existing medical conditions such as asthma, diabetes, and heart disease, as well as those with compromised lungs.

Although many countries began preparations for the spread of the disease relatively early, the effects of the pandemic are still widespread and proliferating. In many parts of the world, a massive shortage of Reverse Transcription Polymerase Chain Reaction (RT-PCR) tests is contributing to an inability to properly combat the pandemic. One

way in which this problem could be alleviated is through the use of computed tomography (CT) scans for diagnosis. In CT scans, COVID-19 manifests itself as ground-glass opacities (GGOs) and consolidation with or without vascular enlargement, interlobular septal thickening, and air bronchogram sign. CT scans are much faster and more accessible than RT-PCR tests, allowing for easier and more quickly acting diagnostics.

More importantly, using CT scans has been shown to be an effective and accurate method for diagnosing COVID-19. RT-PCR tests, in addition to being scarce, take a long time to return results, and even then are prone to frequent false negatives; this results in many people with COVID-19 going undiagnosed. Unfortunately, as with many other medical classification problems, there is a lack of publicly available COVID-19 CT scans. A lack of training data can significantly restrict the performance of deep neural models and lead to overfitting. One potential solution that has been explored to this problem is the use of generic data augmentation techniques. Generic data augmentation involves manipulating the original images in the dataset through various methods of cropping, rotating, and zooming, in order to artificially grow the dataset while preserving the distinguishing features that are present in the images. Generic data augmentation has been shown to be especially useful in fine-grained datasets, or datasets that have low sample sizes and high degrees of similarities between images. For example, extensive use of rotating training images to increase CNN task performance for galaxy morphology classification using a fine-grained data-set [Dieleman *et al.*, 2015]. The primary concern in the application of data augmentation is that of over-fitting occurring. Thus, it is important that studies which employ data augmentation analyze differences in training and validation metrics over the runtime in order to rule out overfitting. In our study, we utilized multiple rotations in order to artificially grow our fine-grained dataset due to their success in previous classifier models (Dieleman *et al.*, 2015) (Taylor and Nitschke, 2017).

In mathematical terms, the Laplace filter is a filter that is defined by the divergence of a scalar field's gradient. In image processing, it is used for enhancing an image's edges to help in its detection. Because derivative filters, among them the Laplace filter, are sensitive to noisy images, they are often performed in association with a smoothing filter to remove noise. The filter operates by calculating a sum of differences across multiple neighboring pixels to replace the magnitude of each individual pixel and is defined as $Laplace(f) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$. This is how it effectively locates edges and simultaneously removes noise from images. In the presence of a bright spot located in a dark region of the image, the Laplace filter will return an even brighter spot to highlight the disparity.

Methods

3.1 Dataset

The dataset that was used in this study was produced by researchers at the University of California at San Diego (He, Xie, Zhao, Zhang, 2020). Currently, the database contains 350 images of CT scans of patients that are positive for COVID-19, and 398 images of CT scans of patients that tested negative.

3.2 Training and Test Set

First, we split the dataset into a training and testing subset; 80% of images were used for training, with the remainder being used for testing. Examples of infected and non-infected CT scans found in our dataset are shown in **Figures 1 and 2**.




Figure 1. CT Scan of Infected Thoracic Region Marked With COVID-19 Indicator




Figure 2. CT Scan of Non-infected Thoracic Region

3.3 Application of the Laplace Filter

The creators of the dataset at UCSD utilized Densenet-169 in order to produce baseline metrics for COVID-19 diagnosis using CT scans. They achieved relatively high performance in all metrics aside from recall, which was a mere 0.762 (Zhao et al., 2020). Thus, we strove to improve the recall of the network while simultaneously improving or preserving the quality of the other metrics. In order to do this, the Laplace filter was utilized in Image Preprocessing below, since it has shown the ability to improve the recall rate of CNNs when tested on X-ray images (Chen, 2019).

3.4 Dataset Split

Using the original UCSD dataset, we created two new datasets. The first set of images did not undergo the Laplace Operator, while the second set did; both sets had the same number of affine transformations performed on them.

3.5 Preprocessing: Image Resizing and Data Augmentation and Laplacian Filter

The images in both datasets were resized to 512x512 pixels to maintain a uniform image size; this was done via the bicubic interpolation technique, which downscaled the images to a uniform size; this was necessary as the images were previously of variable size and scale. Applying bicubic interpolation minimized image distortion and retained information better than less sophisticated methods such as resizing or cropping and resulted in a smoother image than nearest-neighbors or bilinear interpolation.

In addition, we applied affine transformations of 5°, 10°, 15°, 20°, and 25° to the images in order to increase the size of the dataset by 6 times.

Lastly, we applied the Laplace Operator to the second dataset in addition to the aforementioned transformations. A sample image with the Laplace Operator applied to it is shown in **Figure 3**.




Figure 3. CT Scan with Laplacian filter Applied

3.6 CoronaNet

The COVID-19 detection task is a binary classification problem, where the input is an image of a CT scan X and the output is a binary label $y \in \{0,1\}$ indicating the absence or presence of COVID-19, respectively. To accomplish this task, a seven-layer CNN was utilized that we dubbed CoronaNet. This model contained two convolution layers, two pooling layers, a flattening layer, and two fully connected dense layers. The dense layers alleviate the vanishing gradient problem, strengthen feature propagation, encourage feature reuse, and substantially reduce the number of parameters. The network structure can be seen in the CoronaNet Architecture. Within our model, we utilized Adaptive Moment Estimation (Adam), an adaptive learning rate optimization algorithm, with an initial learning rate of 0.00001. Our activation functions for the intermediate layers was the rectifier function, with the final layer using a sigmoid nonlinearity. A batch size of five was used over 50 epochs, after each of which the model output accuracy, performance metrics (precision, recall, and F1 score), and loss, given by the binary cross-entropy loss function: $L(X, y) = -w_+ y \log p(Y = 1|X) - w_-(1 - y) \log p(Y = 0|X)$, where $p(Y = i|X)$ is the probability that the network assigns to the label i , $w_+ = |N|/(|P| + |N|)$, and $w_- = |P|/(|P| + |N|)$ with $|P|$ and $|N|$ the number of COVID-19 and Non-COVID-19 CT scans, respectively. In order to prevent the networks from overfitting, early stopping was performed by saving the network after every epoch and choosing the saved network with the lowest loss on the tuning set. Overall, 2,177,185 parameters were trained and optimized for this task. The architecture of the network can be seen in **Table 1**.

Table 1. CoronaNet Architecture

Layer	Output Shape	Parameters
conv2d_1 (Conv2D)	(None, 98, 98, 32)	896
Max_pooling2d_1 (MaxPooling2)	(None, 49, 49, 32)	0
conv2d_2 (Conv2D)	(None, 47, 47, 32)	9,248
max_pooling2d_2 (MaxPooling2)	(None, 23, 23, 32)	0
flatten_1 (Flatten)	(None, 16928)	0
Dense_1 (Dense)	(None, 128)	2,166,912

Results

We trained CoronaNet on both datasets and acquired the results shown in **Table 2**.

Table 2. CoronaNet Results

Dataset	Accuracy	F1	Precision	Recall
#1	57.82%	0.0000	0.0000	0.0000
#2	92.14%	0.8735	0.9107	0.8938

The following figures depict CoronaNet's progression as it was trained on the first dataset.

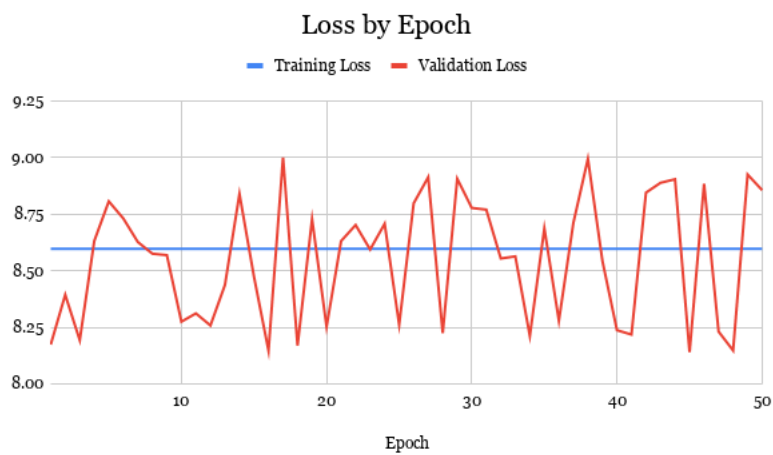


Figure 4. Loss by Epoch




Figure 5. Accuracy by Epoch

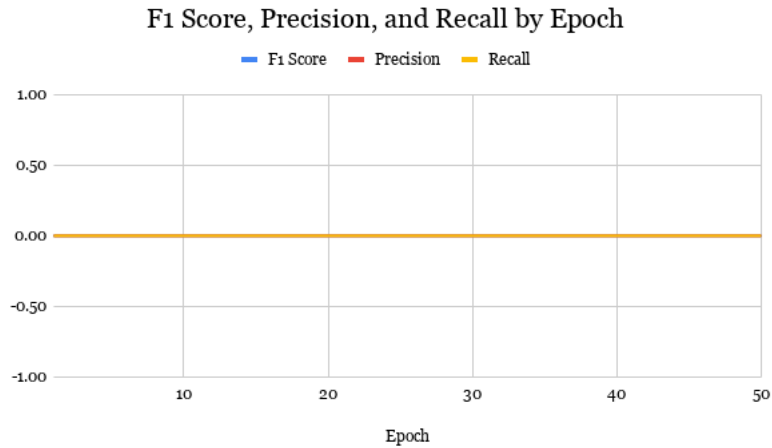


Figure 6. F1 Score, Precision, and Recall by Epoch

The following figures depict CoronaNet’s progression as it was trained on the second dataset.




Figure 7. Loss by Epoch

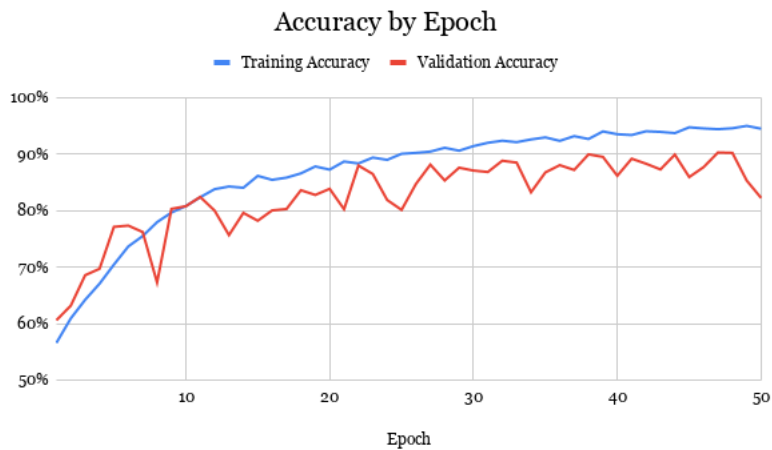


Figure 8. Accuracy by Epoch

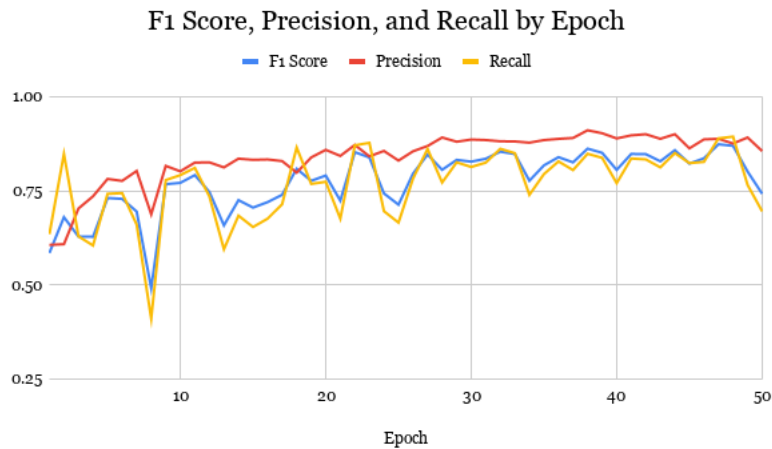


Figure 9. F1 Score, Precision, and Recall by Epoch

Discussion

We trained a novel 7-layer CNN called CoronaNet in order to classify images of thoracic CT scans that are infected by COVID-19 using a dataset provided by researchers at UCSD. We strove to specifically improve the recall metric that was subpar in their study by applying a Laplacian Operation to the images. We found that the Laplacian Operation significantly increased the recall of CoronaNet when compared to the results produced without its application. Moreover, we found that the Laplacian filter significantly improved all of the metrics, including accuracy, F1, and precision. We found that when the Laplacian Operation was applied, CoronaNet reached optimal performance with a peak accuracy of 92.14% and a peak recall of 0.8892. This is likely due to the fact that the Laplace filter highlights key edges within the CT scans that make specific COVID-19 features more distinguishable. Additionally, we noted that when the CoronaNet was trained on images that were unfiltered, the results were extremely poor, scoring 0 in F1, Precision, and Recall, with a peak accuracy 57.82% accuracy. Furthermore, our experiment displays the viability of using data augmentation in biomedical imaging projects, while avoiding high degrees of overfitting as indicated by the similarity in the training and validation metrics over epochs graph. Finally, CoronaNet produced relatively high metrics in COVID-19 diagnosis through CT scan analysis while using a small dataset. This is significant as the traditional diagnosis test that utilizes RT-PCR not only takes much longer but has been shown to be less accurate as well due to its high subjectivity to misdiagnosis by falsely providing a negative test output. We hope that our technology will be able to provide increased diagnostic capabilities in this viral pandemic and provide a foundation to be built off of in the future in terms of biomedical preprocessing techniques.

Conclusion

Based upon the results obtained in this study, it can be concluded that CoronaNet is a viable framework that is less computationally expensive and more effective in comparison to models that used similarly sized datasets. Additionally, this study establishes a precedent for the selective use of the Laplace Operator to boost feature recognition in biomedical imaging analysis using deep learning techniques. Finally, we show that data augmentation can viably be used to mitigate the accuracy limitations that smaller datasets present.

Limitations

This study has various limitations and points of improvement. Firstly, the model could have been trained using a larger dataset. This would likely have improved metrics and made the model less susceptible to overfitting. Additionally, the data that the model was trained upon was all received from the same institution. Future studies might look to use multi-institutional data to make the algorithm more robust. Our study also did not evaluate the performance enhancements that the Laplace Operator provides in other models. Other studies might look to apply our methodology to more widely used model architectures. Finally, we only utilized rotations as part of our data augmentation, whereas other studies may look to use multiple techniques such as random cropping and flipping to improve model performance.

Acknowledgements

We would like to thank Mr. Duke Writer and Dr. Julie Sohl at the Loudoun Academy of Science for their mentorship and guidance in the completion of this project.

References

- A. Krizhevsky, I. S. Nitish, H., W. Karl, M., Shao, L., J. Yosinski, J., D. Erhan, Y., . . . DJ. Drown, T. (1970, January 01). A survey on Image Data Augmentation for Deep Learning. Retrieved July 29, 2020, from <https://link.springer.com/article/10.1186/s40537-019-0197-0>
- Chen, X. (n.d.). Image enhancement effect on the performance of convolutional neural networks. Retrieved from <http://www.diva-portal.org/smash/get/diva2:1341096/FULLTEXT02.pdf>
- Hall, L., Goldgof, D., Paul, R., & Goldgof, G. M. (2020). Finding COVID-19 from Chest X-rays using Deep Learning on a Small Dataset. doi:10.36227/techrxiv.12083964
- Holshue, M. L., Doremalen, N. van, Vaduganathan, M., Fineberg, H. V., & Epidemic Intelligence Service. (2020, March 17). First Case of 2019 Novel Coronavirus in the United States: NEJM. Retrieved from <https://www.nejm.org/doi/full/10.1056/NEJMoa2001191>
- Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L., & Aerts, H. (2018, August). Artificial intelligence in radiology. Retrieved July 29, 2020, from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6268174/>
- Hurt, B., Kligerman, S., & Hsiao, A. (2020). Deep Learning Localization of Pneumonia. *Journal of Thoracic Imaging*, 1. doi: 10.1097/rti.0000000000000512
- Lai, C. C., Shih, T. P., Ko, W. C., Tang, H. J., & Hsueh, P. R. (2020). Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) and coronavirus disease-2019 (COVID-19): The epidemic and the challenges. *International journal of antimicrobial agents*, 55(3), 105924. <https://doi.org/10.1016/j.ijantimicag.2020.105924>
- Litjens, G., Kooi, T., Bejnordi, B., Setio, A., Ciompi, F., Ghafoorian, M., . . . Sánchez, C. (2017, July 26). A survey on deep learning in medical image analysis. Retrieved July 29, 2020, from <https://www.sciencedirect.com/science/article/abs/pii/S1361841517301135>

Zhao, J., Zhang, Y., He, X., & Xie, P. (n.d.). COVID-CT-Dataset: A CT Scan Dataset about COVID-19. Retrieved from <https://arxiv.org/pdf/2003.13865.pdf>