

MBZUAI

Digital.Commons@MBZUAI

Computer Vision Faculty Publications

Scholarly Works

1-26-2022

Hyperparameter Optimization for COVID-19 Chest X-Ray Classification

Ibraheem Hamdi

Mohamed bin Zayed University of Artificial Intelligence

Muhammad Ridzuan

Mohamed bin Zayed University of Artificial Intelligence

Mohammad Yaqub

Mohamed bin Zayed University of Artificial Intelligence

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/cvfp>



Part of the [Artificial Intelligence and Robotics Commons](#), [Biomedical Engineering and Bioengineering Commons](#), and the [Medicine and Health Sciences Commons](#)

Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC BY 4.0

Uploaded 25 March 2022

Recommended Citation

I. Hamdi, M. Ridzuan, and M. Yaqub, "Hyperparameter optimization for COVID-19 chest x-ray classification", 2022, arXiv:2201.10885

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Computer Vision Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Hyperparameter Optimization for COVID-19 Chest X-Ray Classification

Ibraheem Hamdi, Muhammad Ridzuan, and Mohammad Yaqub

Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, U.A.E

Abstract. Despite the introduction of vaccines, Coronavirus disease (COVID-19) remains a worldwide dilemma, continuously developing new variants such as Delta and the recent Omicron. The current standard for testing is through polymerase chain reaction (PCR). However, PCRs can be expensive, slow, and/or inaccessible to many people. X-rays on the other hand have been readily used since the early 20th century and are relatively cheaper, quicker to obtain, and typically covered by health insurance. With a careful selection of model, hyperparameters, and augmentations, we show that it is possible to develop models with 83% accuracy in binary classification and 64% in multi-class for detecting COVID-19 infections from chest x-rays.

Keywords: COVID-19 · X-Ray · Hyperparameter Optimization.

1 Introduction

COVID-19 is an infectious disease that causes mild to moderate respiratory illness, and in severe cases death. Over 2.5 million new cases were reported in the last 24 hours [1], and 60% of the US population are expected to be infected by March of 2022 [2].

PCR, a molecular test that can identify the disease in genetic material from collected swabs, is commonly used to detect the disease [3]. However, the test can cost over 400 US dollars [4], and results are taking longer and longer due to the rise in cases [5,6,7]. Therefore, it is becoming increasingly important to develop a relatively cheaper, faster, and more accessible screening procedure.

X-rays were discovered in the late 19th century and have been used in the medical field ever since [8]. Chest x-rays can be obtained within 15 minutes, including positioning time [9], making them rather quick to produce for a rapid diagnosis. To create a model that can recognize and classify COVID-19 infections from chest x-rays would be immensely beneficial, especially at times like these when early quarantine is critical to limit the spread of the virus.

For learning purposes, we made use of vanilla DenseNet-121 [10] architecture as well as TorchXRayVision [11] library's variations that are pre-trained on x-ray data from different data combinations. We investigated the effect of model, hyperparameter, and augmentation choice on the accuracy of classification. In the beginning, the program was written as a Jupyter Notebook using only vanilla

DenseNet-121 and without any augmentations. Each experiment required manual modification to the code in multiple locations. With the introduction of TorchXrayVision’s models and more augmentations, making a mistake while modifying the code became easier and tracebacks errors started happening. The solution to that was adding a new section in the beginning of the code that asked the user what augmentations, hyperparameters, and models are required. At that time, the only way to stop the code without reaching the set maximum number of epochs was to physically interrupt the program. Furthermore, model checkpoints were being created at every epoch, slowing down the training process and using up too much storage space. To address that, a user-defined validation accuracy was set to start saving checkpoints and another to stop the program.

As the project grew, it became more obvious that a better system was needed, one that is more automated. Preferably, it would have features like early-stopping, to save the best model and retain GPU time, and online logging, to enable remote monitoring. Luckily, a presentation by one of the students at the university introduced us to a potential solution. Hydra-Lightning-Template [12] became the core of this program. It enables easy experimentation, hyperparameter search with Optuna [13], and online logging through Weights & Biases (Wandb) [14]. It is built on Pytorch Lightning, which relies on ‘.yaml’ configuration files to retrieve experiment settings such as hyperparameters and model. This immensely simplified experimentation by minimizing the amount of boilerplate code, allowing us to produce different models without the need to manually change the code.

Optuna is an automatic hyperparameter optimization framework that is integrated into Pytorch Lightning. It has a simple and intuitive interface that is platform agnostic and lightweight with minimal infrastructure dependency, making it extremely versatile. Features include pruning to terminate unpromising trials and speed up optimization, easy parallelization to enable scalability across multiple machines, quick visualization through a built-in dashboard, as well as an abundance of optimization samplers [15]. This makes Optuna more suitable for our experimentation rather than other methods like GPyOpt, SageMaker, or GoogleVizier, that don’t have all of these advantages and features together [16].

2 Dataset

The Society for Imaging Informatics in Medicine (SIIM) partnered with the Foundation for the Promotion of Health and Biomedical Research of Valencia Region (FISABIO), Medical Imaging Databank of Valencia Region (BIMCV) and the Radiological Society of North America (RSNA) to create a Kaggle competition with the aim to advance the identification of COVID-19 Pneumonia from chest x-rays.

The dataset provided is composed of over 6000 scans in DICOM format belonging to four different classes: 1,676 ‘Negative for Pneumonia’, 2,855 ‘Typical Appearance’, 1,049 ‘Indeterminate Appearance’, and 474 ‘Atypical Appearance’.

ance’ [17]. To ensure data balance for binary experiments, 50% of the data was composed of ‘Negative for Pneumonia’, while the other half was a random mix of the other three classes.

To avoid confusing the model with lateral x-rays, 230 data folders that had more than one DICOM file were excluded. Furthermore, the dataset was stratified according to the assigned labels to ensure a balanced distribution of classes. In our experiments, we used a 70:20:10 ratio to split the data into training, validation, and testing datasets.

3 Methods

Optuna offers flexibility through a wide range of samplers to perform the optimization. These include Tree-structured Parzen Estimator (TPA), Gaussian Processes (GP), Covariance Matrix Adaptation Evolution Strategy (CMA-ES), and Random and Grid Search [18]. TPA is the default and recommended sampler by Optuna especially for experiments with less than 1000 trials and hence was used.

Within the Hydra-Lightning-Template, a configuration file can be created to define the optimization settings. For instance, number of trials, optimization metric, sampler or optimization technique, hyperparameters, and range of search are specified inside the ‘.yaml’ file. Optuna then uses the sampler to maximize (or minimize) the optimization metric, validation accuracy in our case, by trying different hyperparameter values within the search space. Based on the learning curve, Optuna may choose to prune or early-stop the experiment to save time and try another set of hyperparameters.

4 Setup

4.1 Image Pre-processing

Pydicom [19] library was used to read the x-ray images provided in DICOM format. The function ‘pixel_array’ is used to convert the images to a 2-dimensional NumPy array. However, Vanilla DenseNet-121 expects 3-channel input, while TorchXRayVision models expect a grayscale image of size 224 x 224. Therefore, we had to design our Data Module to handle files differently according to the model specified for the experiment.

According to the configuration of each experiment, geometrical transformations such as scaling, shear, rotation, translation, as well as horizontal and vertical flipping were applied to investigate the effects of augmentations on the accuracy of models. Scaling, shear, rotation, and translation were chosen because they naturally exist in x-ray imaging as a result of positioning error, while horizontal and vertical flipping can happen as a result of input or scanning error. These transformations were chosen as they make a good representation of what is seen in the real world. Color transformations however would not help since x-rays are always produced in grayscale.

4.2 Models

DenseNet-121 model was chosen for this task due to its reported success with chest x-rays in literature [20]. In addition, there are nine DenseNet-121 TorchXRayVision models that are pretrained on different x-ray datasets such as CheXpert, NIH, and PadChest [11]. We encountered an issue with downloading the weights for the ‘JF Healthcare’ model so it was discarded.

Each of its pretrained models has 18 outputs corresponding to the following pathologies:

- | | | |
|-----------------|----------------------|------------------------------|
| – Atelectasis | – Effusion | – Lung Lesion |
| – Consolidation | – Pneumonia | – Fracture |
| – Infiltration | – Pleural Thickening | – Lung Opacity |
| – Pneumothorax | – Cardiomegaly | – Enlarged Cardiomediastinum |
| – Edema | – Nodule | |
| – Emphysema | – Mass | |
| – Fibrosis | – Hernia | |

When calling TorchXRayVision models, the number of classes is specified to 2 or 4 according to the experiment. However, vanilla DenseNet-121 doesn’t have such a feature, therefore the classifier was changed to ‘nn.Linear(num_features, num_classes)’ where num_features is in_features from the original classifier and num_classes is set to 2 or 4 depending on whether the experiment is binary or multi-class.

4.3 Search Space

To keep the architecture consistent across all experiments, the ResNet model from TorchXRayVision was not used. Furthermore, by using Optuna for hyperparameter optimization, augmentations, batch size, dropout rate, models, and learning rate were experimented with. The search limit was set to 0.0001-0.001 for learning rate, 0%-20% for dropout rate, -15%-15% rotation, -30%-30% scale, -30%-30% shear, and 0%-100% translation. Horizontal and vertical flipping and batch sizes of 8, 16, 32, 64 and 128 were also studied.

5 Results and Discussion

Adam optimizer and cross-entropy loss are used for all of our experiments.

5.1 Binary Classification

To begin, the models were trained using a learning rate of 0.0003, and batch size of 32. TorchXRayVision’s “ALL” model has been trained on all different datasets and achieved the highest baseline accuracy of 81.61% (see Table. 1). Therefore, “ALL” model was used for our binary experiments.

Table 1. Baseline validation accuracy - binary.

Model	Validation Accuracy (%)
DenseNet-121	72.14
ALL	81.61
CheXpert	72.14
MIMIC_CH	75.63
MIMIC_NB	77.80
NIH	78.45
PadChest	79.43
RSNA	50.38

We observed that the maximum validation accuracy peaked using 64 batches, and started decreasing on either side (see Table. 2). This is due to the fact that the ability of a model to generalize degrades beyond a certain batch size [21]. With validation accuracy of 82.81%, the decision was made to use batch size of 64 from that point for our binary experiments.

Table 2. Batch size vs validation accuracy - binary.

Batch Size	Validation Accuracy (%)
8	79.43
16	80.52
32	80.63
64	82.81
128	81.61

Instead of assuming 0.0003, we started searching for the optimal learning rate using “ALL” model and batch size of 64. Validation accuracy peaked at 1.737 (see Table. 3) and dropped above or below that. With too low of a learning rate, the model will take much longer to reach the optimal solution and might even underfit the data. On the other hand, having a learning rate that is very high means the model will be unable to converge as the steps are too large [22]. The highest validation accuracy of 82.48% was achieved with a learning rate of 0.0001737, and we estimate that the optimal learning rate is between 0.00016 and 0.00018.

Dropout is a technique used to boost accuracy and prevent models from overfitting by randomly dropping units from layers during training [23]. This way, a neural network is forced to learn with incomplete information, enabling it to generalize better. Using “ALL” model, batch size of 64, and learning rate of 0.0001737, the best model achieved a validation accuracy of 80.3% using 13.06%

Table 3. Learning rate vs validation accuracy - binary.

Learning Rate (e-4)	Validation Accuracy (%)
0.484	77.15
0.894	80.47
1.273	80.41
1.737	82.84
1.921	82.26
2.916	82.15
5.913	80.41
8.644	78.78

dropout rate (see Table. 4). Contrary to previous intuition, using dropout did not provide an improvement in the validation accuracy. It is possible that the network required more time to learn with dropout, and that a better accuracy could’ve been achieved with longer training.

Table 4. Dropout rate vs validation accuracy - binary.

Dropout Rate (%)	Validation Accuracy (%)
3.678	79.98
6.328	79.98
11.91	80.20
13.06	80.30
18.59	79.76

Due to the decrease in accuracy using dropout, we decided to proceed without any to experiment with augmentations, using “ALL” model, batch size of 64, and learning rate of 0.0001737. The model achieved its highest validation accuracy of 80.85% using scale of +/- 24.12%, shear of +/-18.09%, and translation of 64.25% without any horizontal or vertical flipping (see Table. 5).

5.2 Multi-Class Classification

After binary classification, we experimented with the classification of all four labels separately. Using a learning rate of 0.0003, baseline performance of all eight models was obtained. The highest validation accuracy of 62.15% was achieved using the “ALL” model (see Table. 6). Therefore, we decided to use the “ALL” model for our multi-class experiments as well.

Batch size of 32 achieved the highest validation accuracy of 62.39% (see Table. 7). Hence, we decided to use a batch size of 32 for any following experiments.

Table 5. Augmentations vs validation accuracy - binary.

Horizontal Flip	Off	On	Off	Off	On	Off	On
Vertical Flip	Off	Off	Off	On	Off	Off	Off
Rotation (+/- °)	14	13	0	0	14	14	13
Scale (+/- %)	13.3	9.97	24.12	17.46	5.94	26.60	19.94
Shear (+/- %)	23.91	20.18	18.09	20.55	9.48	28.69	24.22
Translation (+/- %)	46.19	62.71	64.25	89.95	69.92	46.19	62.71
Validation Accuracy (%)	80.30	80.30	80.85	79.65	80.09	80.30	80.41

Table 6. Baseline validation accuracy - multiclass.

Model	Validation Accuracy (%)
DenseNet-121	56.74
ALL	62.15
CheXpert	56.08
MIMIC_CH	57.88
MIMIC_NB	57.42
NIH	60.17
PadChest	61.49
RSNA	27.86

Table 7. Batch size vs validation accuracy - multiclass.

Batch Size	Validation Accuracy (%)
8	62.33
16	62.09
32	62.39
64	62.33
128	62.15

Instead of assuming 0.0003 again, we used Optuna to search for the optimal learning rate with “ALL” model and batch size of 32. The highest validation accuracy of 63.18% was reached with a learning rate of 0.0005913 (see Table. 8). We estimate that the optimal learning rate is between 0.0005 and 0.0006.

Table 8. Learning rate vs validation accuracy - multiclass.

Learning Rate (e-4)	Validation Accuracy (%)
1.737	62.39
1.921	61.79
2.916	62.82
5.153	62.15
5.913	63.18
6.765	62.39
8.684	60.71

Using “ALL” model, batch size of 32, and learning rate of 0.0005913, we noticed a slight improvement using dropout. The best model achieved a validation accuracy of 63.6% using 13.94% dropout rate (see Table. 9). We estimate that the best learning rate is between 13% and 14%.

Table 9. Dropout rate vs validation accuracy - multiclass.

Dropout Rate (%)	Validation Accuracy (%)
3.068	63.30
4.746	8.516
8.516	63.24
13.94	63.30
15.82	63.18
18.59	63.12

With “ALL” model, batch size of 32, learning rate of 0.0005913, and dropout rate of 13.94, we started experimenting with augmentations. The best performing model achieved a validation accuracy of 64.38% using +/- 1.437% scale, +/- 5.252% shear, and 27.61% translation (see Table. 10).

As shown above, the accuracy of our multi-class models hardly exceeded the 60% mark before overfitting the data. After inspecting our confusion matrices (see Fig. 1) and F1 scores (see Fig. 2), we discovered that the models are simply unable to recognize the third (‘Indeterminate Appearance’) and forth (‘Atypical Appearance’) classes, especially the latter. This might be a direct cause of having a smaller amount of data for them compared to the other two.

Table 10. Augmentations vs validation accuracy - multiclass.

Horizontal Flip	On	Off	On	Off	Off	Off	Off
Vertical Flip	Off	Off	Off	Off	Off	Off	Off
Rotation (+/- °)	14	14	13	0	14	0	2
Scale (+/- %)	13.94	7.98	5.98	14.47	5.46	1.44	2.87
Shear (+/- %)	3.16	9.56	8.07	7.24	6.99	5.25	1.46
Translation (+/- %)	7.36	18.48	25.08	25.70	5.17	27.61	28.53
Validation Accuracy (%)	63.42	61.91	63.24	63.24	64.32	65.38	62.52




Fig. 1. Confusion matrix – multiclass.




Fig. 2. F1 heatmap - multiclass.

6 Conclusion

According to a study [24], up to 29% of negative PCR tests can become positive when repeated. With careful tuning of augmentations, hyperparameters, and model, we are able to produce a model with 83% accuracy that can be used as a quick way to detect COVID-19 early and take precautionary measure even without a positive PCR result. However, for multi-class classification, even with further hyperparameter optimization, we don't believe it is possible to exceed an accuracy of 65-66%. In other words, we can classify the presence of COVID-19 pneumonia from chest x-rays but the severity of the infection.

Augmentations are widely accepted as a method to improve accuracy, especially when there is lack or imbalance of data [26]. However, in the case of our binary experiments, this concept did not hold true. Adding dropout was also detrimental to the performance of our binary models. In fact, it caused the training accuracy to peak then decrease, preventing the model from overfitting.

After examining the final results of the competition, we discovered that the highest score achieved was 63.5% [25]. DICE score was used because the challenge asked for bounding boxes and not just classification. With that score in mind, we believe that there is a limitation within the dataset itself regardless of model, learning rate, dropout, batch size, and/or augmentation selection. For future work, we would still like to experiment with other models, optimizers, losses, as well as GAN-based data augmentations [27] to improve our multi-class model.

References

1. WHO Coronavirus (COVID-19) Dashboard, <https://covid19.who.int/>. Last accessed 9 January 2022
2. Reyes, R.: 'Omicron' could cause 140 MILLION new infections from January to March, <https://www.dailymail.co.uk/news/article-10337845/Enormous-spread-Omicron-cause-140-MILLION-new-infections-January-1-March-1.html>. Last accessed 9 January 2022
3. PCR Test for COVID-19: What It Is, How Its Done, What The Results Mean, <https://my.clevelandclinic.org/health/diagnostics/21462-covid-19-and-pcr-testing>. Last accessed 9 January 2022
4. The varying costs of a PCR test at airports around the world, <https://skytraxratings.com/the-varying-costs-of-a-pcr-test-at-airports-around-the-world>. Last accessed 9 January 2022
5. Clarke, K.: Delayed PCR results and travel rule changes stall residents' return to the UAE, <https://www.thenationalnews.com/coronavirus/2022/01/03/delayed-pcr-results-and-travel-rule-changes-stall-residents-return-to-the-uae/>. Last accessed 9 January 2022
6. Guerrero, M.: Delays in COVID-19 Test Results Leave North Texans Frustrated, <https://www.nbcdfw.com/news/local/delays-in-covid-19-test-results-leave-north-texans-frustrated/2851765/>. Last accessed 9 January 2022

7. Farooqui, S.: Canadians struggle with flight delays, testing backlogs as they return from U.S., <https://www.theglobeandmail.com/canada/article-canadians-struggle-with-flight-delays-testing-backlogs-as-they-return/>. Last accessed 9 January 2022
8. History of Medicine: Dr. Roentgen's Accidental X-Rays, <https://columbiasurgery.org/news/2015/09/17/history-medicine-dr-roentgen-s-accidental-x-rays>. Last accessed 10 January 2022
9. X-ray (Radiography) - Chest, <https://www.radiologyinfo.org/en/info/chestrad>. Last accessed 10 January 2022
10. Huang, G., Liu, Z., van der Maaten, Laurens., Weinberger, K.Q.: Densely Connected Convolutional Networks. arXiv:1608.06993 (2016)
11. GitHub - mlmed/torchxrayvision: TorchXRayVision: A library of chest X-ray datasets and models., <https://github.com/mlmed/torchxrayvision>. Last accessed 10 January 2022
12. GitHub - ashleve/lightning-hydra-template: PyTorch Lightning + Hydra. A feature-rich template for rapid, scalable and reproducible ML experimentation with best practices. url<https://github.com/ashleve/lightning-hydra-template>. Last accessed 10 January 2022
13. Optuna - A hyperparameter optimization framework, <https://optuna.org/>. Last accessed 10 January 2022
14. Weights & Biases – Developer tools for ML, <https://wandb.ai/>. Last accessed 10 January 2022
15. GitHub - optuna/optuna: A hyperparameter optimization framework, <https://github.com/optuna/optuna>. Last accessed 14 January 2022
16. Optuna: A Define by Run Hyperparameter Optimization Framework - SciPy 2019, https://www.youtube.com/watch?v=J_aymk4YXhg. Last accessed 14 January 2022
17. SIIM-FISABIO-RSNA COVID-19 Detection - Kaggle, <https://www.kaggle.com/c/siim-covid19-detection>. Last accessed 10 January 2022
18. optuna.samplers — Optuna 2.10.0 documentation, <https://optuna.readthedocs.io/en/stable/reference/samplers.html>. Last accessed 14 January 2022
19. Pydicom, <https://pydicom.github.io/>. Last accessed 10 January 2022
20. Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M.P., Ng, A.Y.: CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning. arXiv:1711.05225 (2017)
21. Keskarm, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. arXiv:1803.09820 (2018)
22. Smith, L.N.: A disciplined approach to neural network hyper-parameters: Part 1 – learning rate, batch size, momentum, and weight decay. arXiv:1609.04836 (2016)
23. Srivastava, N., et al.: Dropout: A Simple Way to Prevent Neural Networks from Overfitting. JMLR. 15(56), 1929-1958 (2014)
24. Watson, J., Whiting, P.F.: Interpreting a covid-19 test result. BMJ 2020;369:m1808
25. SIM-FISABIO-RSNA Covid-19 Detection Challenge - Society for Imaging Informatics in Medicine, https://siim.org/page/covi-19_classification_challenge?_zs=GhQGc&_zl=6Ej12/. Last accessed 13 January 2022
26. Rebuffi, S-A., Gowal, S., Calian, D.A., Stimberg, F., Wiles, O., Mann, T.: Data Augmentation Can Improve Robustness. arXiv:2111.05328 (2021)
27. Sundaram, S., Hulkund, N.: GAN-based Data Augmentation for Chest X-ray Classification. arXiv:2107.02970 (2021)

A Appendix




Fig. 3. Model vs validation accuracy – binary.




Fig. 4. Batch size vs validation accuracy - binary.




Fig. 5. Learning rate vs validation accuracy – binary.




Fig. 6. Dropout rate vs validation accuracy - binary.




Fig. 7. Top 8 augmentations vs validation accuracy - binary.




Fig. 8. Model vs validation accuracy - multiclass.




Fig. 9. Batch size vs validation accuracy - multiclass.




Fig. 10. Learning rate vs validation accuracy - multiclass.

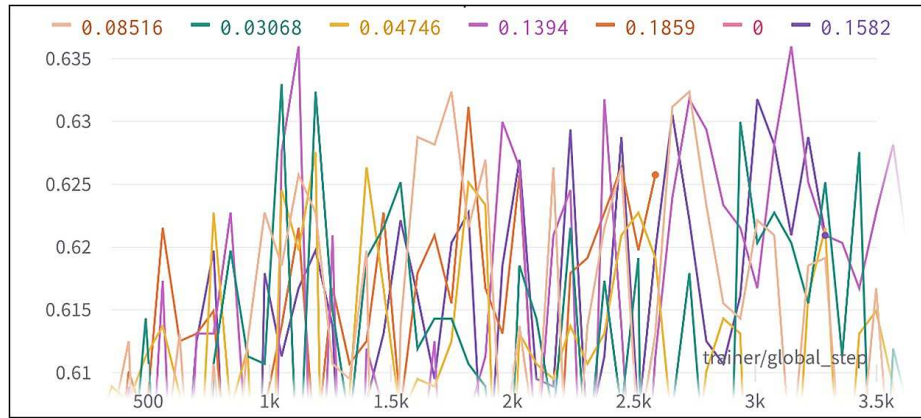


Fig. 11. Dropout rate vs validation accuracy – multiclass.




Fig. 12. Top 8 augmentations vs validation accuracy – multiclass.




Fig. 13. Top 8 augmentations vs f1 macro score - multiclass.