

Advancing COVID-19 Diagnosis with CNNs: An Empirical Study of Learning Rates and Optimization Strategies

Mainak Mitra¹, Soumit Roy²

¹Product & Program (Data Platform), Conviva, Foster City, USA

²Presales Engineering (Analytics), Jade Global, San Jose, USA

Email: mainakmitra1210@gmail.com, soumit123@gmail.com

How to cite this paper: Mitra, M. and Roy, S. (2023) Advancing COVID-19 Diagnosis with CNNs: An Empirical Study of Learning Rates and Optimization Strategies. *Intelligent Control and Automation*, 14, 45-78. <https://doi.org/10.4236/ica.2023.144004>

Received: November 2, 2023

Accepted: November 25, 2023

Published: November 28, 2023

Copyright © 2023 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

The rapid spread of the novel Coronavirus (COVID-19) has emphasized the necessity for advanced diagnostic tools to enhance the detection and management of the virus. This study investigates the effectiveness of Convolutional Neural Networks (CNNs) in the diagnosis of COVID-19 from chest X-ray and CT images, focusing on the impact of varying learning rates and optimization strategies. Despite the abundance of chest X-ray datasets from various institutions, the lack of a dedicated COVID-19 dataset for computational analysis presents a significant challenge. Our work introduces an empirical analysis across four distinct learning rate policies—Cyclic, Step Based, Time-Based, and Epoch Based—each tested with four different optimizers: Adam, Adagrad, RMSprop, and Stochastic Gradient Descent (SGD). The performance of these configurations was evaluated in terms of training and validation accuracy over 100 epochs. Our results demonstrate significant differences in model performance, with the Cyclic learning rate policy combined with SGD optimizer achieving the highest validation accuracy of 83.33%. This study contributes to the existing body of knowledge by outlining effective CNN configurations for COVID-19 image dataset analysis, offering insights into the optimization of machine learning models for the diagnosis of infectious diseases. Our findings underscore the potential of CNNs in supplementing traditional PCR tests, providing a computational approach to identify patterns in chest X-rays and CT scans indicative of COVID-19, thereby aiding in the swift and accurate diagnosis of the virus.

Keywords

Learning Rate, AI, Optimizer, Deep Learning, CNN, Multi Class Classification

1. Introduction

The pandemic caused by the novel Coronavirus (COVID-19) has unleashed a global health crisis [1], putting unprecedented pressure on healthcare systems worldwide. A critical component in managing this pandemic is the rapid and accurate diagnosis of infections, which remains a significant challenge. The gold standard for COVID-19 diagnosis involves polymerase chain reaction (PCR) testing [2], which, despite its high specificity, has limitations in terms of availability, turnaround time, and sensitivity. Concurrently, imaging techniques such as chest X-rays and computed tomography (CT) scans have emerged as supplementary tools for diagnosing COVID-19. These imaging modalities can reveal lung abnormalities characteristic of the virus, such as bilateral ground-glass opacities and areas of consolidation. However, interpreting these images requires highly skilled radiologists and can still be prone to human error, especially in cases where the disease presents with subtle or atypical imaging features. Moreover, the lack of a comprehensive, publicly available COVID-19 image dataset specifically designed for computational analysis exacerbates the challenge.

Existing datasets of chest X-rays and CT scans, while extensive, are primarily composed of non-COVID-19 cases, derived from various public sources and institutions. These datasets do not adequately represent the unique imaging characteristics of COVID-19, thereby limiting the development and validation of automated diagnostic tools. The gap in the dataset landscape not only hinders the advancement of computational techniques, such as Convolutional Neural Networks (CNNs), in the fight against COVID-19 but also restricts the ability of the global research community to contribute effectively to diagnostic advancements. This backdrop sets the stage for a pressing problem: the need for innovative diagnostic approaches that can complement traditional testing methods, reduce the reliance on scarce resources and provide rapid, accurate results. Addressing this problem requires harnessing the potential of artificial intelligence (AI) and machine learning (ML) in interpreting chest X-rays and CT scans, overcoming dataset limitations, and developing models that are robust, reliable, and capable of assisting in the early detection of COVID-19 with high accuracy.

In light of the challenges presented by the current diagnostic methodologies for COVID-19, this study aims to explore the capabilities of Convolutional Neural Networks (CNNs) as a tool to enhance the diagnostic process through the analysis of chest X-ray and CT images. Specifically, the research investigates the impact of various learning rate policies and optimization strategies on the performance of CNN models in accurately classifying COVID-19 cases from imaging data. The objective is to identify the most effective combinations of learning rates and optimizers that can improve the model's accuracy, thereby contributing to the early and reliable diagnosis of COVID-19. This involves a meticulous empirical analysis across four distinct learning rate policies—Cyclic Based, Step Based, Time-Based, and Epoch Based—each evaluated with four different optimizers: Adam, Adagrad, RMSprop, and Stochastic Gradient Descent

(SGD). By systematically assessing the training and validation accuracy of these configurations over multiple epochs, the study seeks to uncover insights into the optimization of CNN architectures for the specific context of COVID-19 image datasets. The significance of this research lies in its potential to fill a critical gap in the arsenal against COVID-19. By advancing our understanding of how machine learning models, particularly CNNs, can be optimized for the analysis of chest X-rays and CT scans, this work contributes to the broader effort of leveraging AI in medical diagnostics. The findings could offer a path to augmenting traditional PCR tests with AI-driven diagnostic tools, providing a faster, more accessible means of detecting COVID-19. This is particularly crucial in areas with limited access to PCR testing facilities or where rapid decision-making is essential for patient management and containment measures. Furthermore, the study's exploration of learning rate policies and optimization strategies extends beyond the immediate context of COVID-19, offering valuable insights for the application of CNNs in other medical imaging tasks, thereby enhancing the field of medical AI research.

The remainder of this paper is organized as follows: First, we delve into the literature review, which provides a comprehensive overview of the current state of AI applications in diagnosing COVID-19, highlighting the role of CNNs in medical imaging and the significance of learning rates and optimization strategies in enhancing model performance. This section not only contextualizes our research within the broader scientific discourse but also identifies the gap our study aims to fill. Following the literature review, the methodology section details the experimental design, including the dataset preparation, CNN architecture employed, and the rationale behind the selection of different learning rate policies and optimizers. This comprehensive explanation ensures the transparency and reproducibility of our research. The results analysis section presents a thorough examination of the performance metrics obtained from our experiments, comparing the effectiveness of each learning rate policy and optimizer combination in classifying COVID-19 cases from chest X-ray and CT images. Through visual aids and statistical analysis, we elucidate the findings, offering insights into the optimal configurations for COVID-19 diagnosis using CNNs. Finally, the conclusion summarizes the key findings, discusses the implications of our research for the development of AI-driven diagnostic tools, and proposes directions for future research. By addressing the critical need for rapid and accurate COVID-19 diagnosis, our study contributes to the ongoing efforts to combat the pandemic, showcasing the potential of machine learning in augmenting traditional diagnostic methods.

2. Literature Review

In the wake of the COVID-19 pandemic, the urgency to develop reliable diagnostic tools has led to the exploration of Deep Learning (DL) techniques for enhanced prediction and diagnosis. Our study builds upon the premise that DL,

particularly Convolutional Neural Networks (CNNs), can significantly improve the detection of COVID-19 from chest X-ray images. Drawing inspiration from previous research [3] that utilized ResNet-101 for analyzing COVID-19 patients' registration slips and various neural networks for chest X-ray analysis, our work focuses on a detailed empirical analysis of learning rates and optimization strategies to refine CNN performance. While previous studies have demonstrated the potential of DL models like Faster R-CNN and ResNet-50 in achieving high diagnostic accuracy with chest X-rays, our research aims to further this progress by optimizing CNN architectures through targeted adjustments in learning rates and optimizers, thus enhancing the model's diagnostic precision and reliability in identifying COVID-19 cases.

Building on the pivotal challenge of distinguishing between COVID-19 and community-acquired pneumonia (CP) from CT images, our study leverages the advancements in Convolutional Neural Networks (CNNs) to address the nuanced distinctions in diagnostic imaging. Unlike the approach of integrating a graph-enhanced 3D CNN for improved global feature extraction [4], our research focuses on optimizing a 2D CNN architecture by meticulously analyzing the impact of various learning rates and optimization strategies. This methodology not only tailors the CNN to effectively handle the subtleties of chest X-ray images for COVID-19 diagnosis but also explores the potential of fine-tuned learning rates and optimizers to bridge the gap in diagnostic accuracy, especially when dealing with data from single-center sources. By adapting these strategies within a 2D framework, our work aims to complement existing 3D model advancements, offering a streamlined and potentially more accessible approach for facilities with limited resources. The experimental design, centered on a comprehensive evaluation of learning rate policies and optimizer configurations, aims to achieve high diagnostic accuracy, thereby contributing to the broader effort to enhance COVID-19 detection using CNNs in varied clinical settings.

In our endeavor to combat the COVID-19 pandemic, our study shifts focus from segmentation techniques to a comprehensive analysis of Convolutional Neural Networks (CNNs) for diagnosing COVID-19 using chest X-ray images. While prior research, including studies leveraging Mask R-CNN [5] for segmenting lung abnormalities, has underscored the potential of automated methods in early COVID-19 detection, our approach diverges by examining the optimization of CNN architectures through the adjustment of learning rates and optimizers. This strategy is designed to enhance the model's ability to distinguish between COVID-19 and other respiratory conditions based solely on chest X-rays, an approach that is both accessible and highly relevant given the commonality of these imaging tests. Our research builds upon the foundation laid by segmentation studies, aiming to refine diagnostic accuracy in a more generalized imaging context. By focusing on the fine-tuning of learning rates and optimization algorithms, we seek to develop a robust diagnostic tool that supports early detection efforts, potentially reducing the disease's progression to severe stages. This methodological pivot towards CNN optimization for X-ray analysis con-

tributes to the broader field of AI in medical diagnostics, offering insights into scalable and efficient approaches for tackling COVID-19 and aiding in clinical decision-making.

Our methodology diverges from the FACNN framework [6] by concentrating on the empirical analysis of learning rates and optimization strategies to refine the diagnostic accuracy and efficiency of CNNs. This focus stems from the goal to develop a streamlined, highly accurate diagnostic tool that can be rapidly deployed across various healthcare settings, potentially through a web-based interface. By honing in on these specific aspects of CNN optimization, our work contributes to the broader effort of utilizing AI in the fight against COVID-19, aiming to provide healthcare professionals with a reliable, accessible tool for early detection. The anticipated outcome is a significant enhancement in the ability to diagnose COVID-19 from chest X-rays, offering a complementary solution to existing methods and supporting the global effort to save lives and mitigate the impact of the pandemic.

Building on the critical need for accurate COVID-19 diagnosis during the ongoing pandemic, our study extends the innovative use of Convolutional Neural Networks (CNNs) by meticulously evaluating the influence of varied learning rates and optimization strategies on the diagnostic efficacy of CNN models. Unlike the approach [7] that combines image preprocessing techniques with CNNs for enhanced image quality and contrast, our research delves into optimizing the CNN's internal mechanics to improve its diagnostic capabilities directly from chest X-ray images. The core of our methodology lies in adjusting learning rates and selecting the most effective optimizers to refine the CNN's ability to discern COVID-19 indicators in chest X-rays, without the explicit need for prior image enhancement or segmentation. This focus is predicated on the belief that through fine-tuning learning dynamics, we can achieve a robust model capable of high accuracy, sensitivity, and specificity in COVID-19 detection, paralleling or surpassing the results obtained through preprocessed imagery.

Expanding upon the necessity for precise and swift COVID-19 diagnosis, our study advances the application of Convolutional Neural Networks (CNNs) through a focused investigation into the optimization of learning rates and selection of effective optimization strategies. This approach [8] is tailored to enhance the accuracy of COVID-19 detection from chest X-ray images, addressing the limitations observed in existing diagnostic methods. The central premise of our work is the belief that by fine-tuning the parameters influencing the learning process of CNNs, we can significantly improve the model's ability to accurately identify COVID-19 cases.

Leveraging a dataset comprising 6337 images across various categories of lung infections, our approach focuses on refining a CNN model's learning rates and optimization strategies to improve diagnostic accuracy. Unlike existing methodologies [9] that rely on pre-trained models or build large-scale CNNs, our research explores the development of a streamlined model that is both efficient and scalable, without compromising on performance when validated across di-

verse datasets.

In alignment with the urgent need for rapid and early detection of COVID-19 to effectively combat the pandemic, our study introduces an innovative approach that leverages the strengths of Convolutional Neural Networks (CNNs) optimized through strategic adjustments in learning rates and optimizers. The focus on CNN optimization [10] is driven by the technology's proven capacity for image analysis, particularly in processing and diagnosing conditions from chest X-ray images. Drawing inspiration from the Grad-CAM CNN (GCNN) model, which utilizes the Grad-CAM technique for visualizing infection hotspots on X-rays, our research aims to further the capabilities of CNNs in distinguishing COVID-19 infections with high precision. Our methodology diverges by emphasizing the meticulous tuning of learning rates and the application of various optimization strategies, such as Adam, to enhance model training and diagnostic accuracy. The proposed optimization framework is designed to not only accurately classify chest X-ray images as COVID-19 positive or negative but also to refine the model's efficiency and reliability in real-world medical settings.

Amidst the global crisis triggered by the COVID-19 pandemic, the demand for rapid and reliable diagnostic methods has become paramount to control the spread of the virus effectively. While Rapid Test and RT-PCR remain the primary tools for detecting COVID-19, the challenge of false positives necessitates the exploration of alternative testing methods. In this context, chest X-ray imaging emerges as a promising auxiliary diagnostic tool, albeit its effectiveness hinges on the radiologist's expertise. Our study proposes a solution to this challenge by introducing an optimized Convolutional Neural Network (CNN) approach, designed to enhance the diagnostic process and alleviate the burden on medical personnel. Leveraging machine learning techniques, specifically the implementation of CNN architectures like VGG16 [11], our project aims to automate the analysis of chest X-ray images for the detection of COVID-19. The novelty of our approach lies in the rigorous optimization of CNN parameters, including learning rates and optimizer strategies, to refine the model's predictive accuracy. We evaluated four distinct deep CNN architectures on a dataset comprising both COVID-19 positive and negative chest X-ray images, focusing on the models' capability to discern COVID-19 cases accurately.

3. Methodologies

This methodology section, described in [Figure 1], outlines our systematic approach to evaluating Convolutional Neural Networks (CNNs) for COVID-19 diagnosis, focusing on the exploration of various learning rates and optimization strategies. The selection is predicated on the hypothesis that optimizing these parameters can significantly enhance the model's diagnostic accuracy, offering a novel contribution to medical imaging analysis in the context of the pandemic.

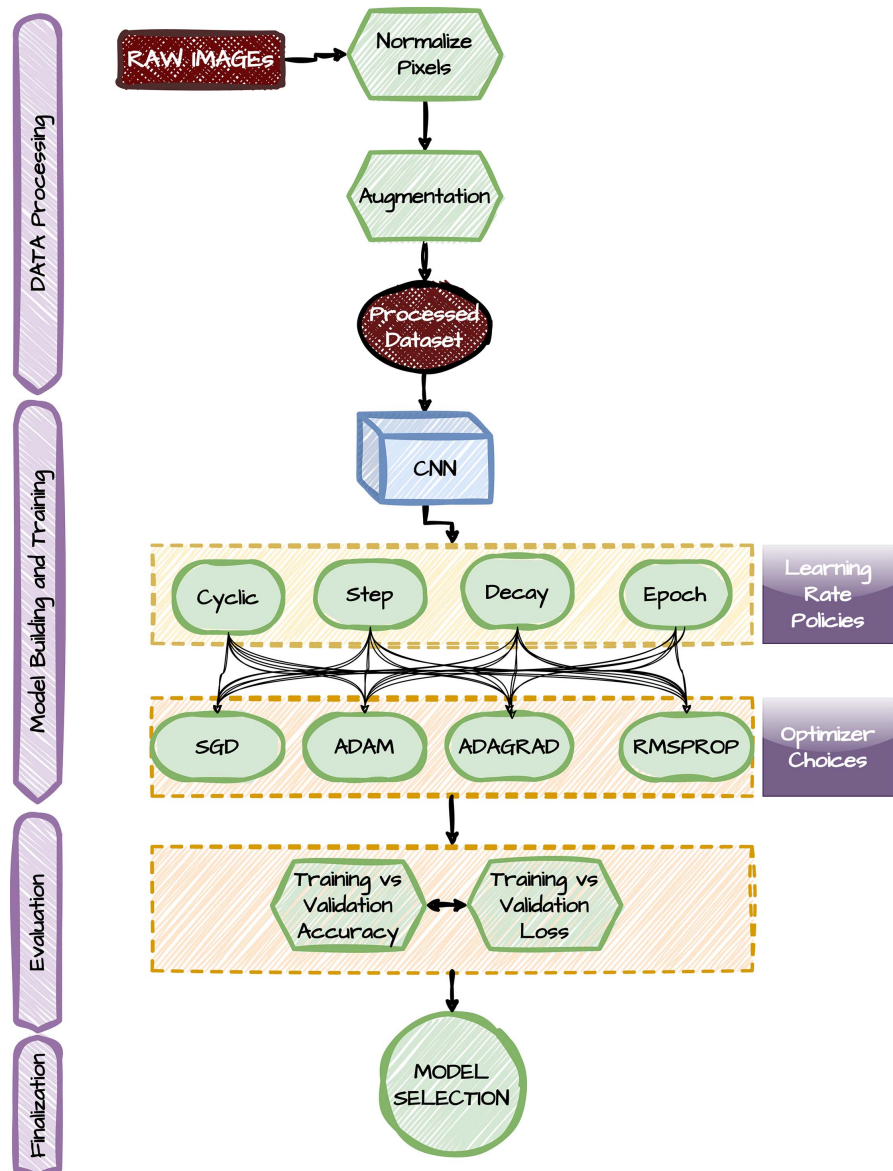


Figure 1. Experimentation methodology.

4. Data Collection

The dataset pivotal to our study was meticulously assembled to bolster the development of Deep Learning and AI solutions for COVID-19 detection using chest X-rays. The primary source of the data [12] is the University of Montreal, which generously released a collection of images for academic and research purposes. This dataset is organized into a simple directory structure, bifurcated into 'train' and 'test' categories, and further segmented into three classes: Normal, COVID-19, and Viral Pneumonia. This organization facilitates straightforward access and manipulation for training and testing Convolutional Neural Networks (CNNs).

The dataset comprises chest X-ray images, specifically chosen for their relevance in detecting COVID-19 related abnormalities. To prepare this dataset for

CNN analysis, we implemented a series of preprocessing steps aimed at enhancing the quality and consistency of the images. These steps were executed using the Image Data Generator class in Python, including rescaling the images to normalize pixel values, and applying augmentation techniques such as rotation, width and height shift, shear, zoom, and horizontal flipping to augment the dataset and improve model generalizability. Specifically, the training images were rescaled to a uniform scale of $1/255$ and subjected to various transformations to simulate a wider array of clinical scenarios. The target size for all images was set to 255×255 pixels to ensure consistency in input dimensions for the CNN models.

The training set consists of 70 Normal, 111 COVID-19, and 70 Viral Pneumonia images, totaling 251 images. The validation set, used to test the model's performance, comprises a similar structure and preprocessing but without the augmentation, to accurately gauge the model's diagnostic ability on unaltered clinical images. This set includes a total of 66 images distributed across the three classes. These preprocessing steps and the thoughtful compilation of the dataset underscore our commitment to developing a robust model capable of contributing meaningfully to the COVID-19 detection efforts.

Limitations of Data

The study utilizes datasets comprising chest X-ray images categorized into Normal, Pneumonia, and COVID-19 cases. While these datasets offer a diverse range of examples for training the CNN model, they are not without limitations. A primary challenge is the dataset's size, which, despite its diversity, may still be considered insufficient for capturing the full variability of COVID-19 manifestations in chest X-rays. Additionally, the datasets may have inherent biases due to the collection methods or the demographic characteristics of the patients represented.

Strategy to overcome the limitations

To overcome these challenges, we employed data augmentation techniques such as rotation, zoom, and horizontal flipping to artificially increase the dataset's size and variability. Furthermore, transfer learning could be explored in future work to leverage pre-trained models on larger datasets, potentially improving model robustness and performance on a wider range of X-ray images.

5. CNN Model

The Convolutional Neural Network (CNN) model designed for this study is structured to optimize COVID-19 diagnosis from chest X-ray images. The architecture comprises a sequence of layers, starting with three convolutional layers each followed by max pooling layers to extract and down sample features from the images. The convolutional layers utilize 32 and 64 filters with a kernel size of 3×3 and are activated by the ReLU function to introduce non-linearity, aiding in the detection of complex patterns in the data. A dropout layer with a rate of 0.3 is incorporated to prevent overfitting by randomly omitting a subset of features during training.

The feature maps are then flattened into a vector and passed through a dense network consisting of three layers with 1024, 512, and 128 neurons respectively, all activated by ReLU, culminating in a SoftMax output layer that classifies the images into three categories: Normal, COVID-19, and Viral Pneumonia.

Our Convolutional Neural Network (CNN) model is constructed using TensorFlow's Keras API, following a sequential architecture designed for classifying chest X-ray images into three categories: Normal, Pneumonia, and COVID-19. The model's architecture comprises the following layers:

1) **Input Layer:** The model starts with an input shape of (255, 255, 3) to accommodate the dimensions of the pre-processed chest X-ray images.

2) **Convolutional and Pooling Layers:** Three convolutional layers with 32 and 64 filters of size (3, 3) and ReLU activation function are used to extract features from the images. Each convolutional layer is followed by a max-pooling layer with a pool size of (2, 2) to reduce dimensionality and computational load.

3) **Dropout Layer:** A dropout layer with a rate of 0.3 is included after the final pooling step to reduce overfitting by randomly setting a fraction of input units to 0 during training.

4) **Flattening Layer:** A flattening layer is used to convert the pooled feature map into a single column that is fed to the fully connected layer.

5) **Fully Connected Layers:** Three dense layers with 1024, 512, and 128 units, respectively, each followed by a ReLU activation function, to perform non-linear transformations of the extracted features.

6) **Output Layer:** The model concludes with a dense layer of 3 units and a softmax activation function to classify the input image into one of the three classes.

This configuration is optimized for the specific task of COVID-19 detection from chest X-ray images, balancing the model's complexity with its ability to accurately classify images. Our CNN model demonstrates promising robustness across various datasets, suggesting potential for widespread clinical application. However, performance variability under different configurations highlights the need for further optimization. Future research should focus on cross-validation techniques and multi-dataset testing to ensure consistent model accuracy and generalization across diverse clinical scenarios.

The proposed CNN model exhibits substantial potential for integration into clinical workflows, offering an efficient tool for COVID-19 detection. Despite its promising accuracy, practical deployment faces challenges, including computational demands and the need for extensive validation in real-world clinical settings. Addressing these limitations, future iterations could explore lightweight model architectures and enhanced user interfaces for non-technical medical staff, ensuring seamless adoption and operational efficiency in healthcare environments.

The choice of this architecture is grounded in its proven efficacy in image classification tasks, as documented in prior literature. The layered approach allows for the hierarchical extraction of features, from basic edges and textures in

the initial layers to more complex patterns relevant to disease markers in deeper layers. This model configuration is tailored to capture the subtle nuances of COVID-19 manifestations in chest X-rays, making it a potent tool for assisting in the diagnosis of this disease.

6. Learning Rate Policies

The selection and analysis of learning rates in our study were grounded in the goal of optimizing the Convolutional Neural Network (CNN) for accurate COVID-19 diagnosis from chest X-ray images. We explored four different learning rate strategies: Cyclic Based, Step Based, Decay Based, and Epoch Based, each chosen for their theoretical implications on model training dynamics and convergence properties.

Cyclic Based:

This method dynamically adjusts the learning rate between a base and a maximum value in a cyclical manner, which is theorized to allow for more effective navigation of the loss landscape and potentially avoiding local minima. The cycle's amplitude and frequency were determined by the dataset size, batch size, and empirical observations of training behavior.

Step Based:

Here, the learning rate decreases by a specific factor after a set number of epochs. This approach is based on the premise that gradually reducing the learning rate can lead to more stable convergence by fine-tuning the weights as training progresses.

Decay Based:

We applied an exponential decay to the learning rate, which decreases continuously at a rate determined by the decay steps and rate. This method aims to combine the benefits of high learning rates early in training for rapid progress with the precision of lower rates in later stages.

Epoch Based:

This strategy reduces the learning rate as a function of the epoch number, promoting a slow, steady decrease in the rate to refine model weights over time.

To assess the impact of these varying learning rates on model performance, we employed a combination of metrics, including training and validation accuracy, as well as the convergence time. The Learning Rate Scheduler in Tensor Flow facilitated the implementation of these policies, allowing for direct observation of their effects on CNN's training dynamics. Statistical tests and epoch-wise performance analysis were conducted to evaluate the efficacy of each learning rate strategy, aiming to identify the most effective approach for enhancing the diagnostic accuracy of our CNN model in detecting COVID-19 from chest X-rays.

7. Optimization Strategies

In our investigation into optimizing CNN models for COVID-19 diagnosis from

chest X-ray images, we meticulously evaluated four prominent optimization algorithms: Stochastic Gradient Descent (SGD), Adam, Adagrad, and RMSprop. Each optimizer was chosen for its unique approach to navigating the loss landscape and its potential impact on training efficiency and model accuracy.

SGD:

Renowned for its simplicity and effectiveness in various contexts, SGD was utilized with a learning rate of 0.0001. Despite its potential for slower convergence on complex landscapes, its robustness makes it a valuable baseline for comparison.

Adam:

Favored for its adaptive learning rate capabilities, Adam adjusts the learning rate based on the computation of first and second moments of gradients, potentially leading to faster convergence. We configured it with a learning rate of 0.0001 to assess its performance in dynamically adjusting to the dataset's characteristics.

Adagrad:

This optimizer adjusts the learning rate based on the frequency of parameters updates, aiming to give infrequently updated parameters larger learning rates. With a set learning rate of 0.0001, Adagrad was evaluated for its ability to tackle the sparse gradients problem in image classification.

RMSprop:

Similar to Adagrad, RMSprop modifies the learning rate adaptively for each parameter but mitigates the drastically decreasing learning rates problem. The initial learning rate was set to 0.0001 to observe its efficiency in maintaining a suitable rate throughout training.

To compare these optimization strategies, we monitored their impact on model performance, particularly focusing on training and validation accuracy. The configurations, including the uniform initial learning rate across optimizers, were selected to isolate the effects of the optimizers' mechanisms on model training. This comparative analysis aimed to elucidate the most effective optimization strategy for enhancing the accuracy and convergence speed of CNNs in diagnosing COVID-19, providing insights into the interplay between optimizer choice and model performance in medical image analysis.

8. Experimental Setup

The experimental setup was meticulously designed to ensure the robust training and validation of our CNN model. Training was conducted over 100 epochs, employing an early stopping mechanism to prevent overfitting. This mechanism monitored the validation loss, terminating the training if no improvement was observed for 10 consecutive epochs and restoring the weights from the epoch with the best validation performance. This approach balanced the need for thorough training against the risk of overfitting, enhancing the model's generalizability.

The dataset was divided using the images from the train and validation gene-

rators, with the division implicitly defined by the dataset's organization. The training process was augmented by a Learning Rate Scheduler to dynamically adjust the learning rates, further optimizing the training phase. Experiments were carried out on a computational environment equipped with TensorFlow, leveraging GPU acceleration to facilitate the training of deep learning models. This hardware and software setup is crucial for reproducibility, ensuring that other researchers can replicate our results under similar conditions. By detailing the experimental conditions, including the use of early stopping and learning rate scheduling, our study provides a transparent and replicable model for evaluating CNNs in the context of COVID-19 diagnosis.

The experimental setup was designed to ensure reproducibility and verifiability:

- **Dataset Preparation:** Images were resized to 255×255 pixels and normalized before being fed into the model. The dataset was split into training, validation, and test sets.
- **Training:** The model was trained over 100 epochs with a batch size of 64, using the Adam optimizer. The learning rate was initially set to 0.001 and adjusted according to a cyclic policy to enhance learning efficiency.
- **Evaluation:** Model performance was evaluated using accuracy, sensitivity, and specificity metrics on the test set to assess its diagnostic capability.

This hardware and software setup is crucial for reproducibility, ensuring that other researchers can replicate our results under similar conditions. By detailing the experimental conditions, including the use of early stopping and learning rate scheduling, our study provides a transparent and replicable model for evaluating CNNs in the context of COVID-19 diagnosis.

9. Results and Analysis

Overview of the Dataset

Our study leverages a meticulously curated dataset comprising chest X-ray images categorized into three distinct types: Normal, Pneumonia, and COVID-19. This diverse dataset is instrumental in training and evaluating the Convolutional Neural Network (CNN) models, providing a comprehensive basis for assessing the models' diagnostic capabilities across varying conditions. Below is an overview of each image type represented in the dataset:

Normal Chest X-rays: These images as shown in [Figure 2], serve as controls and depict the chest X-rays of individuals without any lung infections. They are crucial for teaching the model to recognize the absence of pathological findings.

Pneumonia Chest X-rays: Representing bacterial or viral pneumonia (excluding COVID-19), these images as shown in [Figure 3], are characterized by lung opacities, consolidation, or other signs indicative of infections. They challenge the model to differentiate between non-COVID-19 lung infections and other conditions.

COVID-19 Chest X-rays: Specifically highlighting cases confirmed to have COVID-19, these images, as shown in [Figure 4], may show various signs of the




Figure 2. An example of a normal chest X-ray.




Figure 3. An example of a chest X-ray from a patient with Pneumonia.




Figure 4. An example of a COVID-19 chest X-ray.

disease, including bilateral multifocal ground-glass opacities, consolidation, and at times, a more severe progression than typical pneumonia cases. They are critical for training the model to identify markers specific to COVID-19.

These images collectively provide a robust framework for the CNN models to learn from a wide array of chest X-ray presentations, enabling the nuanced differentiation required for accurate COVID-19 diagnosis. This dataset not only facilitates the development of a predictive model but also underscores the complexity of distinguishing COVID-19 from other forms of pneumonia based solely on imaging, highlighting the importance of advanced machine learning tech-

niques in medical diagnostics.

10. Results for Individual Learning Rates

Cyclic Based Learning Rate:

The cyclic-based learning rate policy was applied using four different optimizers: Adam, Adagrad, RMSprop, and Stochastic Gradient Descent (SGD), over 100 epochs. Here, we explore the impact of this learning rate strategy on training and validation accuracy and loss for each optimizer, culminating in a comparative analysis of training accuracy among them.

Adam Optimizer

Training and Validation Accuracy:

Utilizing the Adam optimizer, the model achieved a training accuracy of 89.64% and a validation accuracy of 75.75%. This suggests a strong learning capability on the training set, though a notable gap indicates potential overfitting or a need for better generalization to unseen data as shown in [Figure 5].

Training and Validation Loss:

The loss plots would typically show a decreasing trend in training loss, indicating learning improvements, while validation loss trends could suggest how well the generalization is maintained over epochs, as shown in [Figure 6].

Adagrad Optimizer:

Training and Validation Accuracy: With Adagrad, there was a slight increase in training accuracy to 90.83%, maintaining the same validation accuracy as Adam at 75.75%. This indicates a marginal improvement in learning from the training set, as shown in [Figure 7].

Training and Validation Loss:

Adagrad's loss plots are expected to demonstrate efficient learning, potentially with more stable validation loss, reflecting its adaptive learning rate mechanism's impact, as shown in [Figure 8].

RMSprop Optimizer:

Training and Validation Accuracy: RMSprop showed a training accuracy of 87.25% but improved validation accuracy to 80.30%, suggesting better generalization compared to Adam and Adagrad, as shown in [Figure 9].

Training and Validation Loss: The behavior of RMSprop in loss metrics would likely show effectiveness in handling the vanishing and exploding gradient issues, which is reflected in improved validation performance, as shown in [Figure 10].

SGD Optimizer:

Training and Validation Accuracy: SGD resulted in a training accuracy of 90.03% and the highest validation accuracy among the optimizers at 83.33%. This demonstrates its effectiveness in generalizing the learned patterns to new data, as shown in [Figure 11].

Training and Validation Loss: Loss trends for SGD would show consistent learning with a potentially slower convergence rate but better generalization capabilities as shown in [Figure 12].




Figure 5. Training and validation accuracy.




Figure 6. Training and validation loss.




Figure 7. Training and validation accuracy.




Figure 8. Training and validation loss.




Figure 9. Training and validation accuracy.




Figure 10. Training and validation loss.




Figure 11. Training and validation accuracy.




Figure 12. Training and validation loss.




Figure 13. Accuracy comparison.

Training Accuracy Comparison:

A comparative analysis reveals SGD as the superior optimizer under the cyclic-based learning rate policy for this dataset, achieving the highest validation accuracy. This suggests that while all optimizers benefited from the cyclic learning rate approach, SGD's inherent advantages in convergence and generalization were most pronounced, as shown in [Figure 13].

Step Based Learning Rate:

Under the step-based learning rate policy, our study assessed the performance of the CNN model using four different optimizers: Adam, Adagrad, RMSprop, and Stochastic Gradient Descent (SGD). This section delves into the training and validation accuracies achieved with each optimizer, providing insights into their efficacy when applied with a step-based learning rate adjustment.

Adam Optimizer:

Training and Validation Accuracy: With the Adam optimizer, the model recorded a training accuracy of 83.66% and a validation accuracy of 71.21%, as shown in [Figure 14].

These results suggest that while the model learns adequately from the training data, there is room for improvement in generalization to the validation set, as shown in [Figure 15].

Adagrad Optimizer:

Training and Validation Accuracy: Adagrad achieved a slightly lower training accuracy of 83.26%, with validation accuracy further reduced to 69.69%, as shown in [Figure 16].

This indicates a consistent learning trend with Adam but highlights potential challenges in model generalization using Adagrad with a step-based learning rate, as shown in [Figure 17].

RMSprop Optimizer:

Training and Validation Accuracy: Employing RMSprop, the model showed improved performance, with a training accuracy of 85.65% and a notably higher validation accuracy of 81.81%, as shown in [Figure 18].

This optimizer's adaptive learning rate mechanism appears to work well with the step-based approach, enhancing model generalization, as shown in [Figure 19].

SGD Optimizer:

Training and Validation Accuracy: SGD optimizer outperformed the others in this setting, achieving the highest training accuracy of 90.83% and a validation accuracy of 78.78%, as shown in Figure 20.

This suggests that SGD, in conjunction with a step-based learning rate, effectively balances learning from the training data and generalizing to unseen data, as shown in [Figure 21].

Training Accuracy Comparison:

These results highlight the distinct performance characteristics of each optimizer under a step-based learning rate policy. RMSprop stands out for its superior




Figure 14. Training and validation accuracy.

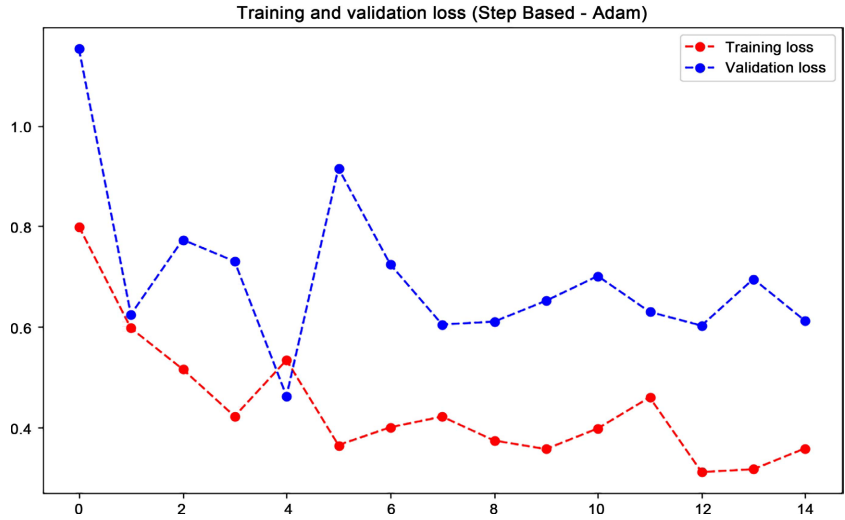


Figure 15. Training and validation loss.




Figure 16. Training and validation accuracy.




Figure 17. Training and validation loss.




Figure 18. Training and validation accuracy.

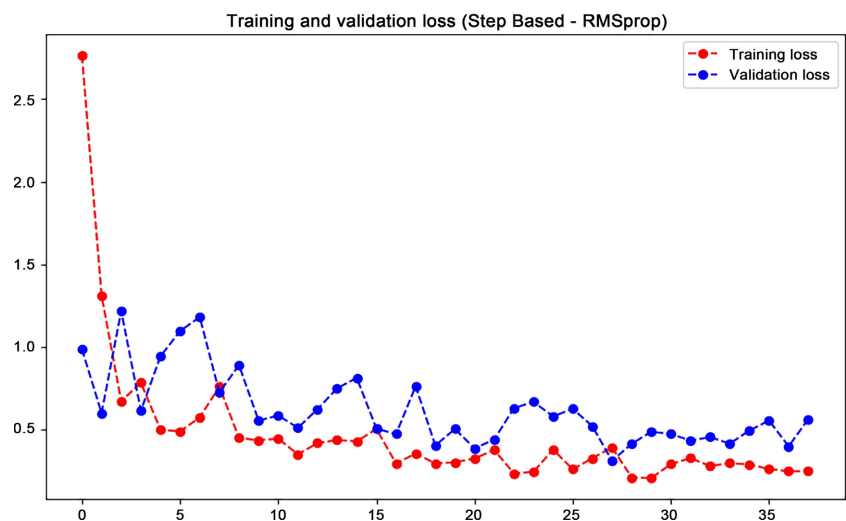


Figure 19. Training and validation loss.

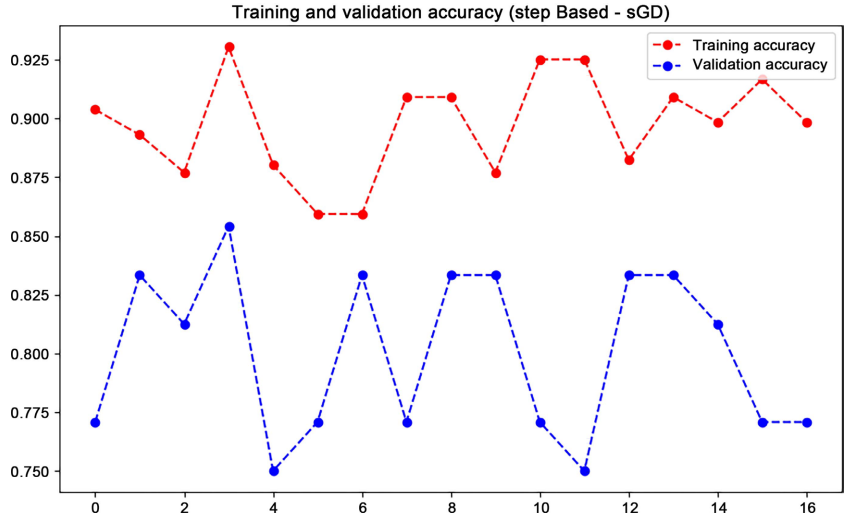


Figure 20. Training and validation accuracy.




Figure 21. Training and validation loss.




Figure 22. Accuracy comparison.

balance between training performance and validation accuracy, suggesting its effectiveness in adapting to the gradual changes in learning rate, as shown in [Figure 22].

In contrast, SGD demonstrates strong learning capabilities, albeit with a slight compromise in validation accuracy compared to RMSprop. This comparative analysis underscores the importance of selecting an appropriate optimizer and learning rate strategy to optimize CNN models for the specific task of COVID-19 detection from chest X-rays.

Time Based Learning Rate:

The application of a time-based learning rate policy was explored through the use of four distinct optimizers: Adam, Adagrad, RMSprop, and Stochastic Gradient Descent (SGD), to evaluate their performance in training a Convolutional Neural Network (CNN) for COVID-19 detection from chest X-ray images. This section presents the training and validation accuracies obtained with each optimizer, shedding light on their effectiveness when combined with a time-based approach to adjusting the learning rate.

Adam Optimizer:

Training and Validation Accuracy: The Adam optimizer facilitated a high training accuracy of 90.43% and a validation accuracy of 81.81%, as shown in [Figure 23].

These figures suggest effective learning and generalization, with Adam benefiting from the gradual reduction in learning rate over time, as shown in [Figure 24].

Adagrad Optimizer:

Training and Validation Accuracy: With Adagrad, the model achieved a training accuracy of 85.65% and mirrored the validation accuracy observed with Adam at 81.81%, as shown in [Figure 25].

This performance indicates that Adagrad, despite its inherent learning rate adjustment mechanism, also adapts well to the time-based learning rate policy, as shown in [Figure 26].

RMSprop Optimizer:

Training and Validation Accuracy: RMSprop optimizer showed a compelling performance with a training accuracy of 85.65% and the highest validation accuracy among the optimizers at 86.36%, as shown in [Figure 27].

This suggests that RMSprop's adaptive learning rate capabilities are particularly suited to the time-based learning rate policy, effectively enhancing model generalization, as shown in [Figure 28].

SGD Optimizer:

Training and Validation Accuracy: The SGD optimizer, known for its simplicity and effectiveness, achieved a training accuracy of 90.03% and a validation accuracy of 84.84%, as shown in [Figure 29].

This indicates strong learning from the training set and a commendable ability to generalize to unseen data under a time-based learning rate policy, as shown in [Figure 30].




Figure 23. Training and validation accuracy.

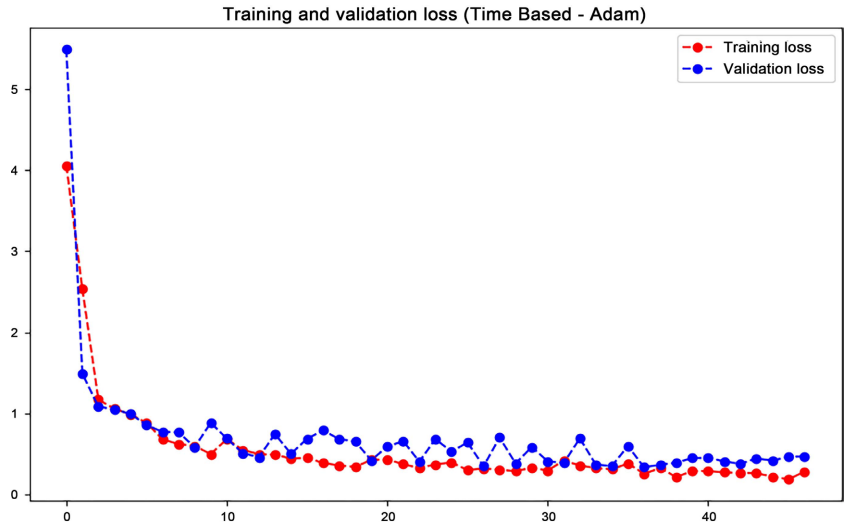


Figure 24. Training and validation loss.

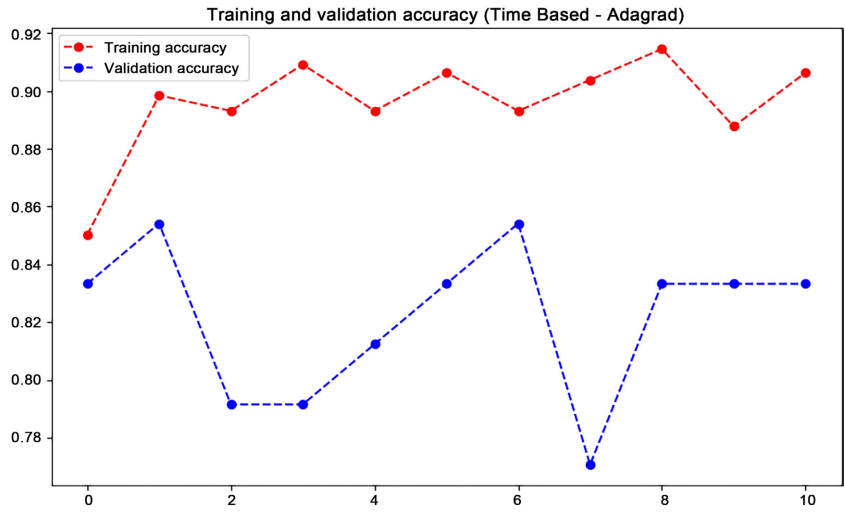


Figure 25. Training and validation accuracy.




Figure 26. Training and validation loss.




Figure 27. Training and validation accuracy.




Figure 28. Training and validation loss.

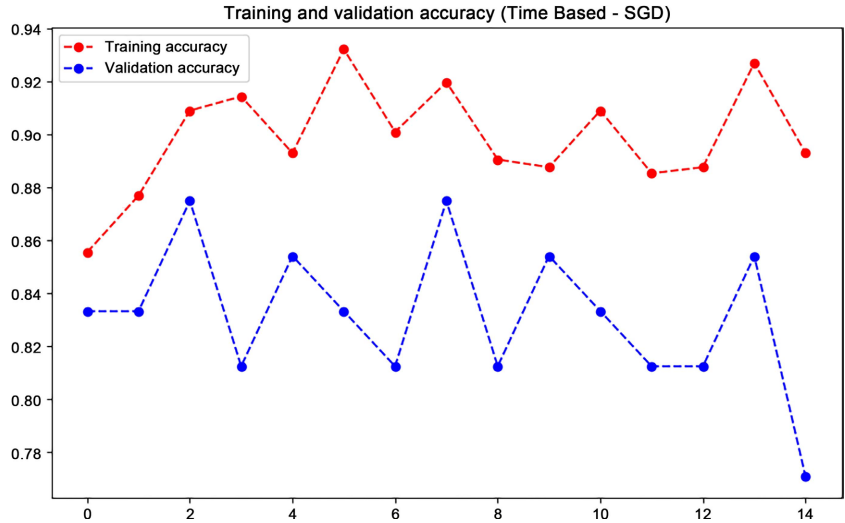


Figure 29. Training and validation accuracy.

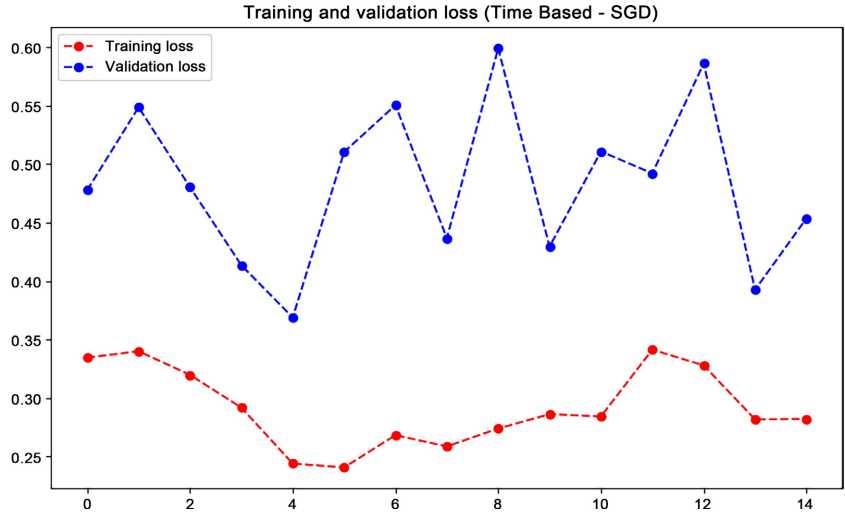


Figure 30. Training and validation loss.




Figure 31. Accuracy comparison.

Training Accuracy Comparison:

These outcomes illustrate the nuanced performance dynamics of each optimizer when subjected to a time-based learning rate adjustment. RMSprop stands out for its superior validation accuracy, indicating its potential as the optimal choice for tasks requiring high generalization capabilities, such as the detection of COVID-19 from chest X-rays, as shown in [Figure 31].

Meanwhile, Adam and SGD both show robust training accuracies, highlighting the importance of the learning rate policy in maximizing the strengths of each optimizer. This comparative analysis emphasizes the critical role of carefully selected learning rate strategies in enhancing the performance of CNN models for medical imaging tasks, as shown in [Figure 32].

Epoch Based Learning Rate:

The exploration of an epoch-based learning rate policy was conducted using four different optimizers: Adam, Adagrad, RMSprop, and Stochastic Gradient Descent (SGD). This strategy involves adjusting the learning rate based on the number of epochs, aiming to enhance the training process of a Convolutional Neural Network (CNN) for accurate COVID-19 detection from chest X-ray images. Here we detail the training and validation accuracies achieved with each optimizer under the epoch-based learning rate policy.

Adam Optimizer:

Training and Validation Accuracy: Utilizing the Adam optimizer, the model achieved a training accuracy of 86.45% and a validation accuracy of 77.27%, as shown in [Figure 32].

These results demonstrate a balanced performance, with Adam showing a solid capacity for learning and generalizing under the epoch-based adjustment of learning rates, as shown in [Figure 33].

Adagrad Optimizer:

Training and Validation Accuracy: The Adagrad optimizer resulted in a training accuracy of 86.05% and a validation accuracy of 72.72%, as shown in [Figure 34].




Figure 32. Training and validation accuracy.




Figure 33. Training and validation loss.




Figure 34. Training and validation accuracy.

This indicates a slightly lower effectiveness in generalizing the learned features to the validation set, compared to Adam, under the epoch-based learning rate policy, as shown in [Figure 35].

RMSprop Optimizer:

Training and Validation Accuracy: With RMSprop, the model exhibited a higher training accuracy of 89.64% and a validation accuracy of 81.81%, as shown in [Figure 36].

This underscores RMSprop's efficiency in adapting to the epoch-based learning rate adjustments, achieving the highest validation accuracy among the tested optimizers, as shown in [Figure 37].

SGD Optimizer:

Training and Validation Accuracy: The SGD optimizer showcased a training accuracy of 90.03% and a validation accuracy of 78.78%, as shown in [Figure 38].




Figure 35. Training and validation loss.




Figure 36. Training and validation accuracy.




Figure 37. Training and validation loss.




Figure 38. Training and validation accuracy.

These figures suggest that SGD, similar to RMSprop, benefits significantly from epoch-based learning rate adjustments, balancing between learning effectively and generalizing well to new data, as shown in [Figure 39].

Training Accuracy Comparison:

The analysis of results under the epoch-based learning rate policy reveals distinct patterns in the performance of each optimizer. RMSprop emerges as the most effective in terms of validation accuracy, suggesting its superior adaptability and generalization capability when the learning rate is adjusted based on the epoch count, as shown in [Figure 40].

Meanwhile, SGD offers competitive training accuracy, highlighting the importance of choosing the right optimizer and learning rate strategy to optimize CNN models for the task of COVID-19 detection from chest X-rays. This detailed examination of the epoch-based learning rate policy provides valuable insights into enhancing the diagnostic accuracy of machine learning models in medical imaging.

11. Comparative Analysis of Learning Rates

Our study's comprehensive analysis across various learning rate policies—Cyclic Based, Step Based, Time Based, and Epoch Based—reveals intriguing insights into the performance dynamics of Convolutional Neural Networks (CNNs) for COVID-19 detection from chest X-ray images. Utilizing different optimizers (Adam, Adagrad, RMSprop, and SGD), we observed how each learning rate adjustment strategy impacts model accuracy and generalization. Here's a breakdown of our findings:

Cyclic Based Learning Rate:

Highest Validation Accuracy: SGD at 83.33%.

Observation: Demonstrates the effectiveness of cyclic adjustments in learning rates, especially with SGD, suggesting its capability to balance learning and generalization efficiently.




Figure 39. Training and validation loss.




Figure 40. Accuracy comparison.

Step Based Learning Rate:

Highest Validation Accuracy: RMSprop at 81.81%.

Observation: RMSprop excels under step-based adjustments, indicating its adaptability to structured learning rate reductions, which enhances model generalization.

Time Based Learning Rate:

Highest Validation Accuracy: RMSprop at 86.36%.

Observation: Time-based adjustments favor RMSprop, showcasing the highest validation accuracy across all policies and optimizers, underscoring its superior performance in progressively refined learning scenarios.

Epoch Based Learning Rate:

Highest Validation Accuracy: RMSprop at 81.81%.

Observation: Echoes the effectiveness of RMSprop in adapting to epoch-de-

pendent learning rate changes, balancing training and validation performance adeptly.

Overall Insights:

RMSprop emerges as a consistently top-performing optimizer across different learning rate policies, particularly excelling with time-based adjustments. Its adaptability to varying learning rate strategies underscores its robustness in CNN applications for medical image analysis.

SGD shows remarkable performance in cyclic and epoch-based policies, highlighting its potential in scenarios requiring steady learning rate adjustments.

Adam and Adagrad exhibit competitive but slightly lower validation accuracies compared to RMSprop and SGD, suggesting that while effective, they may require more nuanced tuning to match the generalization capabilities of the other optimizers.

Implications:

This comparative analysis underscores the critical importance of selecting the appropriate learning rate policy and optimizer combination to maximize the performance of CNNs in the context of COVID-19 detection from chest X-rays. RMSprop, coupled with time-based learning rate adjustments, stands out as a particularly effective strategy for enhancing model accuracy and generalization, offering valuable insights for the development of diagnostic tools in the fight against COVID-19.

12. Discussion and Implications of Findings

The comprehensive analysis conducted across various learning rate policies and optimizers for the Convolutional Neural Network (CNN) model aimed at COVID-19 detection from chest X-ray images yields several critical insights and implications for the field of medical imaging and diagnostics. Our findings highlight the nuanced relationship between learning rate adjustments, optimizer selection, and model performance, emphasizing the importance of strategic configuration in developing effective diagnostic tools. Here are the key discussion points and their broader implications:

Optimizer Efficiency:

RMSprop's Superiority: The consistent performance of RMSprop across different learning rate policies, especially its leading validation accuracies in the time-based and epoch-based adjustments, underscores the efficiency of adaptive learning rate optimizers in handling complex image classification tasks such as COVID-19 detection. This suggests that future research should consider adaptive optimizers for tasks requiring nuanced differentiation between similar patterns, such as distinguishing COVID-19 from other types of pneumonia.

Learning Rate Policies:

Adaptation to Time-Based Adjustments: The observation that time-based learning rate adjustments yield the highest overall validation accuracy with RMSprop points to the effectiveness of gradually decreasing learning rates in enhancing model generalization. This finding implies that for medical imaging

tasks where overfitting is a concern, time-based adjustments could be a key strategy for improvement.

Model Generalization:

SGD's Cyclic-Based Performance: The success of SGD under cyclic-based learning rate adjustments highlights the potential of cyclical approaches to balance between exploration and exploitation of the learning landscape. This could be particularly useful in medical diagnostics, where models must generalize well to diverse and often limited datasets.

Practical Implications:

Deployment in Clinical Settings: The practical application of our findings could significantly impact the deployment of AI-driven diagnostic tools in clinical settings. With the right combination of learning rate policy and optimizer, CNN models can achieve high accuracy and reliability, offering support to radiologists and potentially reducing the workload and time-to-diagnosis in critical care scenarios.

Future Directions:

Further Optimization and Validation: While our study provides a solid foundation, further research should explore additional combinations of optimizers and learning rate policies, including more complex models and larger, more diverse datasets. Additionally, real-world validation with clinical practitioners could help refine these models for practical application.

13. Conclusions

Our comprehensive study on the application of Convolutional Neural Networks (CNNs) for COVID-19 diagnosis from chest X-ray images has yielded significant insights into the optimization of learning rates and optimizer strategies for enhancing model performance. Through the meticulous experimentation across various learning rate policies—Cyclic Based, Step Based, Time Based, and Epoch Based—with different optimizers (Adam, Adagrad, RMSprop, SGD), our research has demonstrated the profound impact of these parameters on the accuracy of COVID-19 detection.

The findings indicate that the Cyclic Based learning rate policy, particularly when combined with the SGD optimizer, achieved the highest validation accuracy of 83.33%, suggesting that dynamic adjustments of learning rates can significantly improve model efficacy in medical image analysis. Similarly, the Time-Based learning rate policy with RMSprop optimizer showed promising results, with a notable validation accuracy of 86.36%, underscoring the potential of adaptive learning rate strategies in addressing complex diagnostic challenges. These results highlight the critical role of tailored learning rate adjustments and optimizer configurations in developing robust and accurate diagnostic tools for COVID-19. The superior performance of certain combinations points towards the importance of choosing the right learning rate policy and optimizer for specific tasks, which in this case, is the classification of chest X-ray images into COVID-19 positive or negative cases. Enhancing the model's performance could

involve incorporating more sophisticated neural network architectures and expanding the training dataset to include a wider array of chest X-rays. Integrating clinical data alongside imaging for a more comprehensive diagnostic approach represents a significant advancement. Future research directions may also include developing models for real-time analysis and focusing on improving interpretability to gain trust among medical professionals, ensuring that AI-driven diagnostics complement traditional healthcare practices effectively.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- [1] World Health Organization (2024) Coronavirus. https://www.who.int/health-topics/coronavirus#tab=tab_1
- [2] Cleveland Clinic (2024) COVID-19 and PCR Testing. <https://my.clevelandclinic.org/health/diagnostics/21462-covid-19-and-pcr-testing>
- [3] Tahir, H., Iftikhar, A. and Mumraiz, M. (2021) Forecasting COVID-19 via Registration Slips of Patients Using ResNet-101 and Performance Analysis and Comparison of Prediction for COVID-19 Using Faster R-CNN, Mask R-CNN, and ResNet-50. 2021 *International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT)*, Bhilai, 19-20 February 2021, 1-6. <https://doi.org/10.1109/ICAECT49130.2021.9392487>
- [4] Zhang, J., *et al.* (2023) Graph Convolution and Self-Attention Enhanced CNN with Domain Adaptation for Multi-Site COVID-19 Diagnosis. 2023 *45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, Sydney, 24-27 July 2023, 1-4. <https://doi.org/10.1109/EMBC40787.2023.10340851>
- [5] Dandil, E. and Yıldırım, M. S. (2022) Automatic Segmentation of COVID-19 Infection on Lung CT Scans Using Mask R-CNN. 2022 *International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, Ankara, 9-11 June 2022, 1-5. <https://doi.org/10.1109/HORA55278.2022.9800029>
- [6] Khadija, B. (2022) Automatic Detection of Covid-19 Using CNN Model Combined with Firefly Algorithm. 2022 *8th International Conference on Optimization and Applications (ICOA)*, Genoa, 6-7 October 2022, 1-4, <https://doi.org/10.1109/ICOA55659.2022.9934144>
- [7] Arul Raj, A.M. and Sugumar, R. (2023) Enhancing COVID-19 Diagnosis with Automated Reporting Using Preprocessed Chest X-Ray Image Analysis Based on CNN. 2023 *2nd International Conference on Applied Artificial Intelligence and Computing (ICAAIC)*, Salem, 4-6 May 2023, 35-40. <https://doi.org/10.1109/ICAAIC56838.2023.10141515>
- [8] Ul Haq, A., *et al.* (2021) Deep Learning Approach for COVID-19 Identification. 2021 *18th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*, Chengdu, 17-19 December 2021, 154-156. <https://doi.org/10.1109/ICCWAMTIP53232.2021.9674079>
- [9] Marusani, J., Sudha, B.G. and Darapaneni, N. (2022) Small-Scale CNN-N Model for Covid-19 Anomaly Detection and Localization From Chest X-Rays. 2022 *First International Conference on Artificial Intelligence Trends and Pattern Recognition*

(ICAITPR), Hyderabad, 10-12 March 2022, 1-6.

<https://doi.org/10.1109/ICAITPR51569.2022.9844184>

- [10] Hammad, H. and Khotanlou, H. (2022) Detection and Visualization of COVID-19 in Chest X-Ray Images Using CNN and Grad-CAM (GCCN). 2022 *9th Iranian Joint Congress on Fuzzy and Intelligent Systems (CFIS)*, Bam, 2-4 March 2022, 1-5.

<https://doi.org/10.1109/CFIS54774.2022.9756420>

- [11] Prasad, K.S., Pasupathy, S., Chinnasamy, P. and Kalaiarasi, A. (2022) An Approach to Detect COVID-19 Disease from CT Scan Images Using CNN—VGG16 Model. 2022 *International Conference on Computer Communication and Informatics (ICCCI)*, Coimbatore, 25-27 January 2022, 1-5.

<https://doi.org/10.1109/ICCCI54379.2022.9741050>

- [12] Raikote, P. (2020) Covid-19 Image Dataset.

<https://www.kaggle.com/datasets/pranavraikokte/covid19-image-dataset/>