COVID-19 INFECTION DETECTION FROM RADIOGRAPHS

Using Neural Networks and Vector Machines for Rapid and Accurate Medical Testing

Authors
Benjamin Kleyner, Haochen Yang, Elijah Rosen, Kieran Carroll
Carnegie Mellon University - 02-251 Professor Philip Compeau
May 2, 2024

1 Abstract

Building upon the foundation laid by Erdaw and Tachbele [1], who developed a machine learning model for the automatic detection of COVID-19 in chest X-ray images, this study extends the application to differentiate COVID-19 pneumonia from other types of pneumonia. Employing a combination of Convolutional Neural Networks (CNN) and Support Vector Machines (SVM), our models replicate and advance previous findings, achieving high classification accuracy. In particular these models achieve 98.5% accuracy in a binary classification task (COVID-19 Pneumonia vs. Other) and 89.13% accuracy in a ternary classification task (COVID-19 Pneumonia vs. Normal vs. Viral Pneumonia). The SVM model, originally utilized by Erdaw and Tachbele, has been complemented with a CNN approach to enhance feature extraction and classification tasks, thereby broadening the scope to identify pneumonia caused by pathogens other than SARS-CoV-2. Our results confirm the viability of using advanced machine learning techniques to improve diagnostic accuracy in medical imaging, contributing valuable insights into the automated differentiation of pulmonary diseases. This extension not only underscores the adaptability of AI in medical diagnostics but also paves the way for more targeted and swift healthcare responses in distinguishing between COVID-19 and other respiratory illnesses.

2 Introduction

Medical imaging, such as X-rays and CT scans, plays a critical role in the non-invasive diagnosis of diseases, including respiratory conditions like pneumonia, where imaging is the current standard diagnostic practice [2]. The onset of the COVID-19 pandemic has heightened the need for rapid and accurate diagnostic tools capable of distinguishing between similar illnesses, a task which medical professionals struggle with [3]. Currently, radiologists assess chest X-ray images for signs of COVID-19 by looking for several key features indicative of the disease. Consolidations, presenting as denser, whiter areas, indicate fluid or inflammatory infiltration, often located at the lung bases and possibly affecting both lungs. COVID-19 often shows bilateral involvement and affects the outer sections of the lungs, distinguishing it from some other types of pneumonia. Linear opacities, representing thickened interlobular septa or filled alveolar walls, might also be visible, along with vascular enlargement due to inflammation. Radiologists also look for the absence of features typical of other respiratory conditions, such as nodules and pleural effusions. These visual inspections are complemented by considering the patient's symptoms and exposure history, though it's important to note that the disease's radiographic presentation can vary with the stage of infection[4]. While traditional imaging requires expert interpretation, which, as mentioned earlier, is often inaccurate, advances in artificial intelligence (AI), specifically machine learning (ML) and deep learning (DL), offer significant enhancements. AI technologies automate the detection of patterns in imaging data, increasing diagnostic accuracy and consistency. This study leverages Convolutional Neural Networks (CNN) and Support Vector Machines (SVM) to extend the capabilities of chest radiography, aiming to improve the speed and reliability of diagnosing COVID-19 relative to other respiratory diseases. By advancing AI-driven diagnostics, we aim to enhance treatment efficacy and healthcare responsiveness in managing respiratory health crises.

3 Background

Building upon the foundation laid by Erdaw and Tachbele [1], who developed a machine learning model for the automatic detection of COVID-19 in chest X-ray images, this study extends the application to differentiate COVID-19 pneumonia from other types of pneumonia. Employing a combination of Convolutional Neural Networks (CNN) and Support Vector Machines (SVM), our models replicate and advance previous findings, achieving high classification accuracy. The SVM model, originally utilized by Erdaw and Tachbele, has been complemented with a CNN approach to enhance feature extraction and classification tasks, thereby broadening the scope to identify pneumonia caused by pathogens other than SARS-CoV-2. Our results confirm the viability of using advanced machine learning techniques to improve diagnostic accuracy in medical imaging, contributing valuable insights into the automated differentiation of pulmonary diseases. This extension not only underscores the adaptability of AI in medical diagnostics but also paves the way for more targeted and swift healthcare responses in distinguishing between COVID-19 and other respiratory illnesses.

This study aims to fill this gap by employing both CNN and SVM models to not only replicate existing diagnostic accuracies but also extend them to more nuanced distinctions between multiple types of pneumonia. This approach seeks to harness the strengths of both

model types to improve diagnostic granularity and reliability, contributing to more targeted therapeutic interventions.

4 Defining a Computational Problem

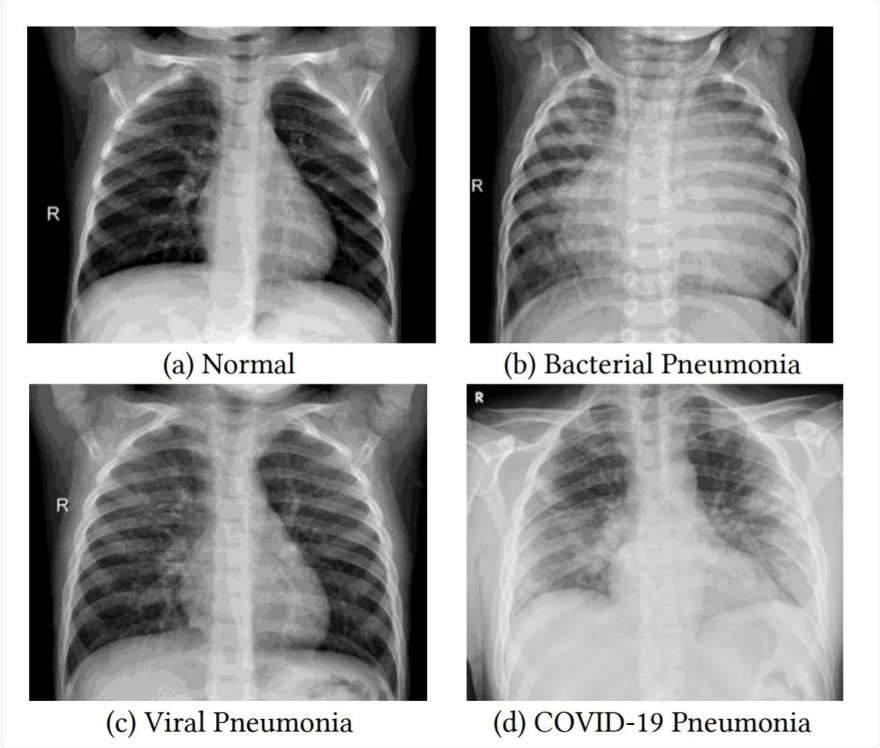

Figure 1: An idealized comparison of chest x-rays caused by different conditions [5]

As is visible above, differentiating between normal and patients with a respiratory illness from a chest x-ray appears to a be a trivial task for a medical professional. However, when asked to differentiate between pneumonia caused by COVID-19 or another viral/bacterial infection doctors can only do so with 62.8% sensitivity[3]. As a result, we define the following classification task to be solved with a machine learning approach.

Input

- A dataset of N chest X-ray images, denoted as $X = \{x_1, x_2, \ldots, x_N\}$, where each image x_i is represented by a high-dimensional pixel array.
- Each image x_i is pre-labeled with one of C classes $\{c_1, c_2, \ldots, c_C\}$ where C includes categories such as normal, various types of pneumonia (e.g., viral, bacterial, fungal, COVID-19), and other thoracic diseases.

Output

- A classification label $\hat{y}_i$ for each image x_i, that minimizes the cross-entropy loss function defined by $H(P^*|P) = -\sum_i P^*(i) \log P(i)$. This is the function used to optimize the models used in this project.

5 Algorithms and Computational Approaches

5.1 Convolutional Neural Network

A **Convolutional Neural Network** (CNN) is a machine learning model that is especially good at recognizing features of image and classifying them according to a sequence of layered computations. A general outline of the computation is as follows:

5.1.1 Input Layer

The CNN takes an image as input, represented as an array of its pixels. In the case of colored images, there are 3 separate arrays for each of the color channels, but the case of x-ray image data (grayscale), we only have 1 matrix representing the intensity of each pixel.

5.1.2 Convolutional Layers

The convolutional layers apply a convolution operation to the images by sliding a set of learnable filters/kernels which are applied via computing dot products to produce feature maps. As the filter slides across the width and height of the input volume, a 2-dimensional activation map is produced giving the the response of the filter at every position. The network learns filters that activate when they see specific features at some position in the input.

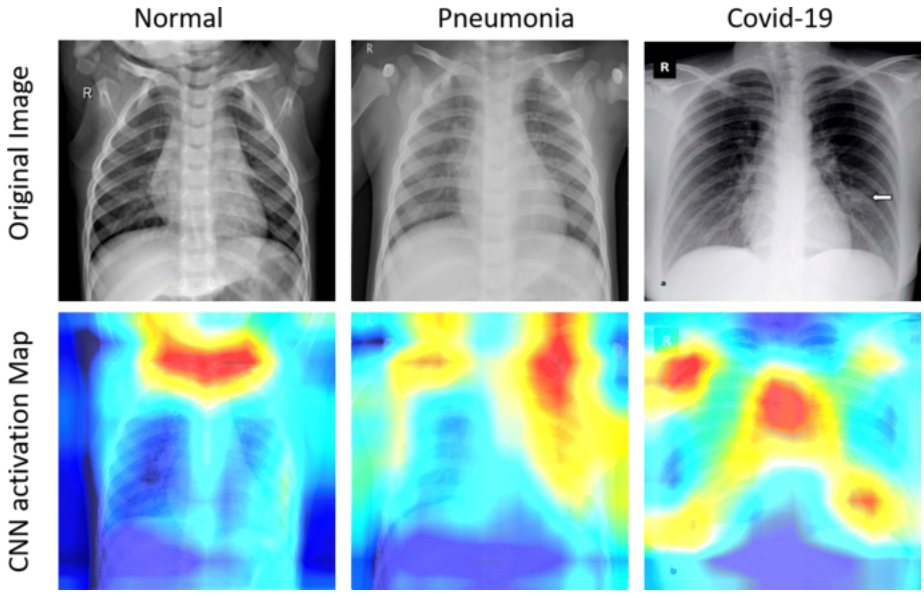

Figure 2: Feature maps created by convolutional layers [6]

5.1.3 Activation Function

After the convolution operations, an activation function such as Recitfied Linear Unit (ReLU) is applied. The ReLU function introduces non-linearity into the model, allowing it to learn more complex patterns. The ReLU activation function is defined as follows:

$$\mathsf{ReLU}(x) = \max(0, x)$$

The ReLU activation function promotes sparsity in the neural network. Sparseness is

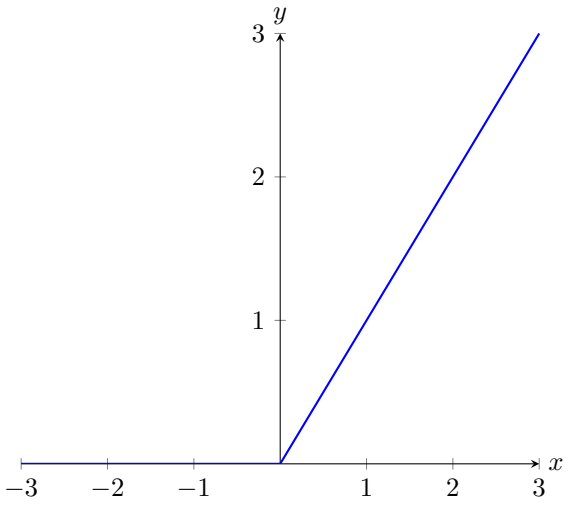

Figure 3: Graph of the ReLU Function

important for the success of our model not only because it improves computational efficiency by reducing the amount active neurons, but also by reducing the risk of overfitting. This allows for better generalization to novel data and allow the network to quickly converge to a set of salient features in an image.

5.1.4 Pooling Layers

These layers reduce the width and height of the input volume for the next convolutional layer by a technique called max pooling where the max value from a set of pixels is retained. This downsampling reduces computational complexity and reduces *overfitting* by abstracting some features. Note that *overfitting* is meant to describe the idea that a machine learning model may learn the dataset that it is trained on too well and recognize features that are not relevant in classifying novel datasets.

5.1.5 Fully Connected Layers

After several layers of pooling and convolution operations, **fully-connected layers** handle the high-level reasoning of the neural network. Each neuron in a fully-connected layer computes a weighted sum of inputs adds a bias term, passes the result through an activation function such as ReLU, and then makes the output available to the next layer.

5.1.6 Output Layer

The last of the **Fully Connected Layers** usually uses a softmax or sigmoid activation functions to give a final classificatioin. The softmax function gives a probability distribution over the classes in the form of a vector whose entries sum to 1 and the sigmoid function returns a value $0 \leq p \leq 1$ and is used for binary classfication tasks.

5.1.7 The Sigmoid and Softmax Activation Functions

The sigmoid activation function is defined as:

$$S(x) = \frac{1}{1 + e^{-x}}$$

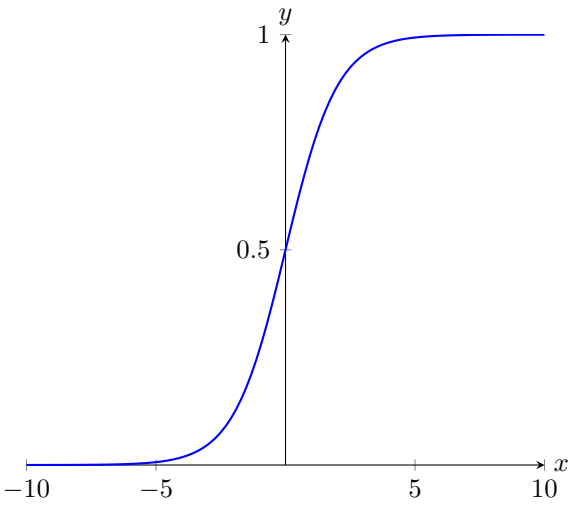

Figure 4: Graph of the Sigmoid Function

This function allows us, in the last layer, to classify each image as being representative of a COVID-19 infection or otherwise. We utilize a threshold value of $p = 0.5$. So all all inputs that sigmoid maps to $x \geq 0.5$ are considered to be COVID-19 and all others are considered to be representative of a COVID-19 infection.

The Softmax function is defined as follows:

$$\sigma(\vec{z})_i = \frac{e^{z_i}}{\sum_{j=1}^{K}} e^{z_j}$$

where K is the number of classifications there are for the given task.

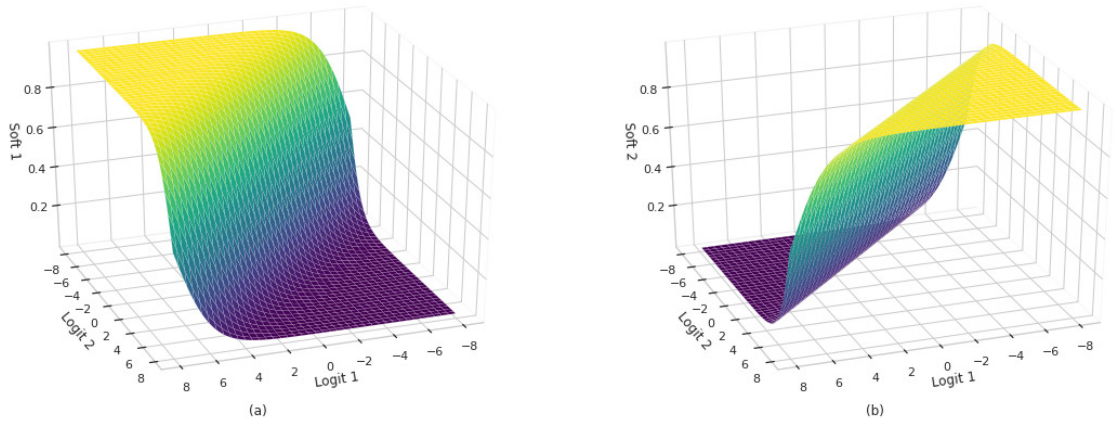

Figure 5: Graph of the Softmax Function[7]

Essentially, this function can be described as a generalization of the Sigmoid function. It transforms the raw outputs of the neural network into a vector of probabilities (that sum to 1) which resembles probability distribution over the input classes. The class with the highest probability value, will be the predicted label of the image.

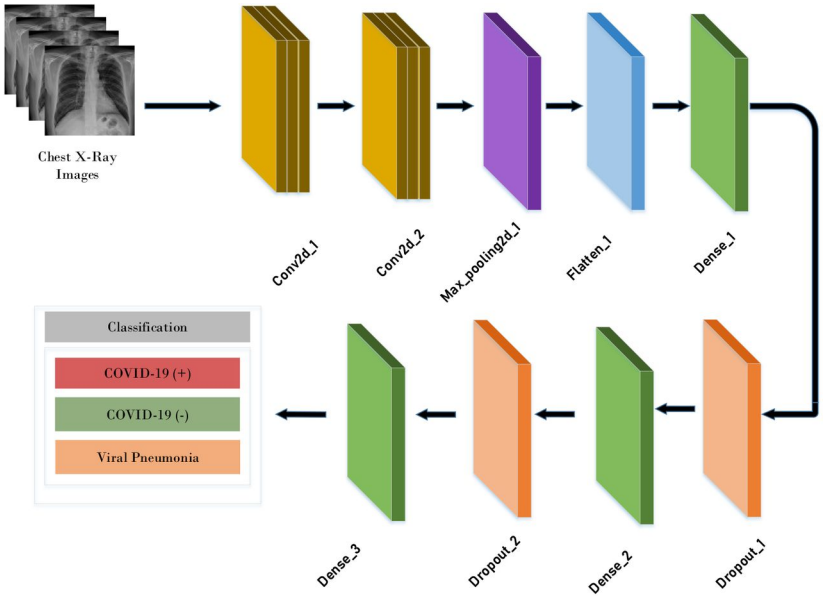

Figure 6: An overview of the layers of the CNN model [8]

5.2 Support Vector Machines

A Support Vector Machine (SVM) is a powerful supervised machine learning model used for classification tasks, including medical image analysis. In the context of our project, we employ an SVM to differentiate chest X-ray images into two categories: those that exhibit indicators of COVID-19 and those that do not.

5.2.1 Classification Mechanism

At its core, the SVM operates by identifying the optimal hyperplane that separates the two classes with the greatest margin. This hyperplane is the decision boundary: points on one side belong to one class, while points on the other side fall into the other class. What makes SVM powerful is the fact that only support vectors need to be stored. This minimizes space complexity and improves efficiency.

5.2.2 Preprocessing and feature extraction

For the classification of chest X-rays using an SVM, each image is first processed to extract relevant features, which typically include textures, edges, and specific patterns pertinent to pulmonary abnormalities. These features are transformed into a high-dimensional space (using techniques like Histogram of Oriented Gradients (HOG) or Principal Component Analysis (PCA)) where the SVM's kernel trick is applied to facilitate a separation even in non-linear distributions.

5.2.3 Kernel functions

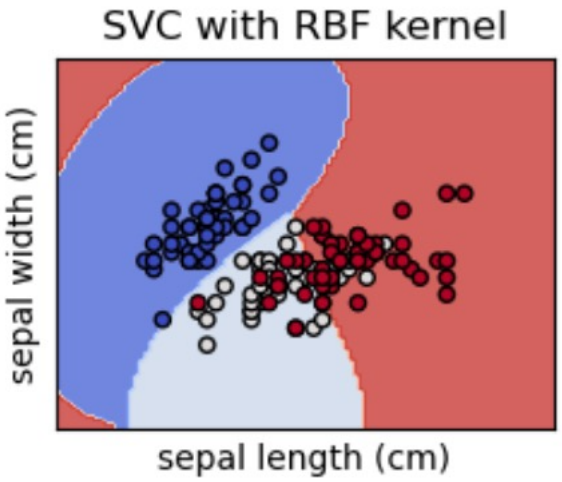

Figure 7: A visualization of a 2D projection of the hyperplane used to separate datapoints in the SVM [9]

The kernel choice, often a Gaussian Radial Basis Function (RBF), critically influences the SVM's performance by dictating how the separation line bends around data points in the feature space. The two sample RBF kernel equation $K(x', x) = \exp\left(-\gamma||x' - x||^2\right)$

corresponds to the above example, with tuning variable γ that dictates how jagged or smooth the hyperplane is. Though only a projection of the actual hyperplane used by the SVM, Figure 7 gives some intuition as to how γ values could effect the generated separation lines. A larger γ would allow for more curvature and inflection points (potentially adding another inflection point to cover more of the red dots), while a smaller γ would reduce the number of inflection points (potentially uncovering some of the white dots). SVMs are prized for their effectiveness in high-dimensional spaces and their ability to handle overfitting well, particularly in scenarios where the number of dimensions exceeds the number of samples.

5.2.4 Training the SVM

To train an SVM for image classification, each image is first converted into a $n \times 1$ feature vector whose values represent the intensities/colors of pixels. The SVM then learns from a training set of such vectors, each labeled with its category (COVID-19 vs no COVID-19). Once trained, the SVM can classify new images as a COVID-19 patient or not by computing the side of the decision boundary on which the new image's feature vector falls, offering a valuable diagnostic tool that adds to physicians' capabilities in fast and accurate COVID-19 detection.

6 Methodology

6.1 Data Acquisition

For the CNN model, a Kaggle data set [10] of approximately 15,000 images containing chest radiographs representing a variety of upper respiratory infections (normal, COVID-19 pneumonia, viral pneumonia) was split 80/20 for training and validation, respectively.

For the SVM model, a Kaggle data set was used [11]. 96 images of combined COVID-19 x-rays and non-COVID-19 x-rays were used to fit the SVM model. To validate the accuracy of the SVM, the remaining images were randomly selected and run through the SVM in batches of 20 due to the long image processing time.

6.2 Data Analysis

For this project **Amazon Web Services** and **Google Colab** were used to analyze the data sets and train/test the CNN and SVM models. Firstly, to perform data analysis and build our models, we installed numpy, pandas, tensorflow, matplotlib, and keras. We prepared the data sets and added layers to the CNN with ReLU and Sigmoid activation functions so that it can be trained to perform binary classification on images. Note that since the last activation function used in the last layer of the CNN is a sigmoid activation function, the final prediction is outputted as a value p, where $0 \leq p \leq 1$. The model then uses this predication value p to classify each image as not COVID-19 or COVID-19. Next, we prepared the training and validation datasets using an 80/20 split. Using line graphs, we visualized the test and training accuracy and loss.

Next, we prepared the model for testing on a novel dataset. First, the classes and labels of the test and training generators were loaded. For each image in the novel dataset, the CNN model made a classification prediction (COVID-19 or no COVID-19). We set the cut-off at 0.5. Thus, for all predictions $p < 0.5$, the image was classified as no COVID-19, and $p \geq 0.5$ was classified as a COVID-19 patient.

7 Results

7.1 CNN Results

After training the CNN using the methodology described in the previous section, we obtained a training accuracy of 98.25% and a validation accuracy of 98.5%.

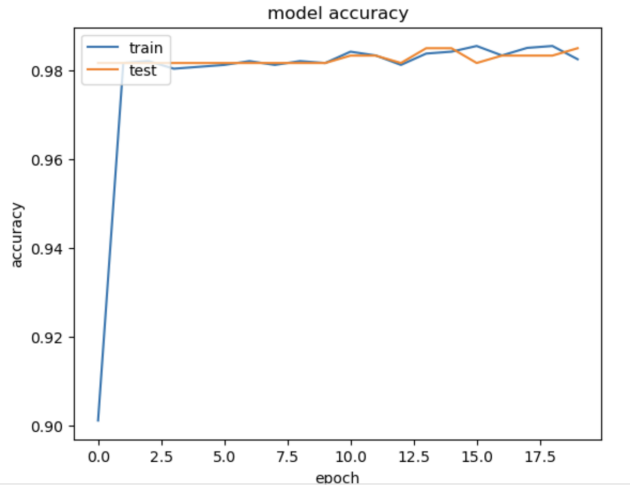

Figure 8: CNN training accuracy in binary classification task

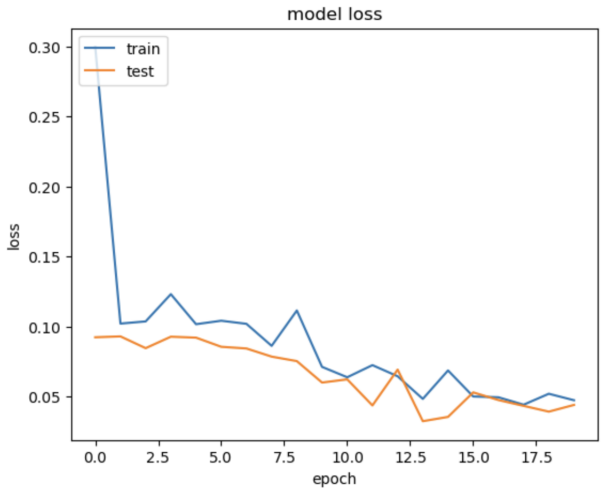

Figure 9: CNN model loss in binary classification task

Following training the model was run on a novel dataset (no shared images with the training set). On this data set, it classified 624 of 624 images correctly. More testing is necessary to determine whether this unprecedented accuracy is in fact the result of overfitting or some other oversight.

However, our ultimate goal was to go beyond binary classification and achieve a ternary classification model capable of separating chest radiographs from COVID-19, normal, and other chest radiographs of viral / bacterial pneumonia with better accuracy than physicians.

In this task, our model achieved 92.68% training accuracy and 90.97% validation accuracy, as shown in the figure below.

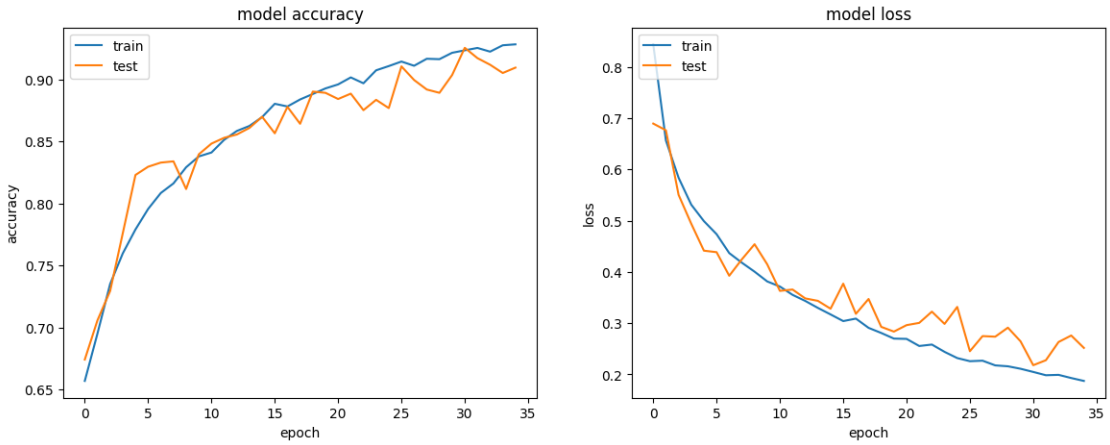

Figure 10: Ternary classification model accuracy statistics

This lower accuracy is explained by the following confusion matrix.

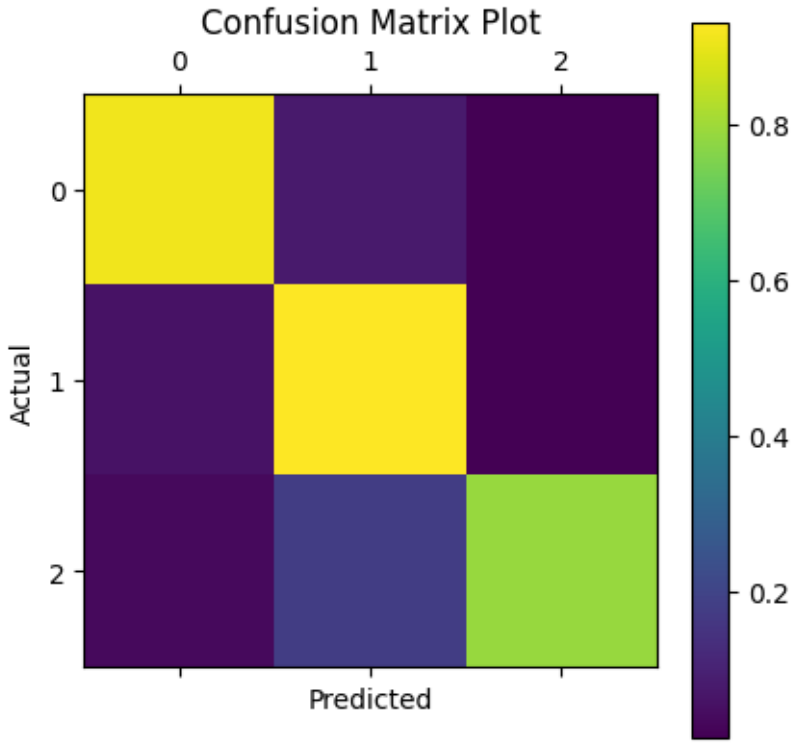

Figure 11: Confusion matrix for ternary classification task

As the plot demonstrates, the CNN identifies COVID-19 and normal chest x-rays with high accuracy, but struggles with differentiating between normal x-rays and x-rays attributable to other viruses or bacteria. However, the precision over 80% for this class far

exceeds the ability of physicians to differentiate between these conditions by radiography. This specific classification is likely the most difficult in the context of the images presented earlier. As such, a potential resolution to this issue may include adding more layers to the network or varying the activation functions used.

7.2 SVM Results

Since the SVM utilizes an N-dimensional hyperplane for binary classification, it is not possible to generate a comprehensive visual. However, an intuitive 2D visualization still exists; an example is Figure 4. The SVM model, when fed non-COVID-19 images, produced an accuracy of 98% and a false positive of 2%. When fed COVID-19 images, the model produced an accuracy of 98.4% and a false negative of 1.6%. For an overall accuracy of 98.2% (981 out of 999 was correct).

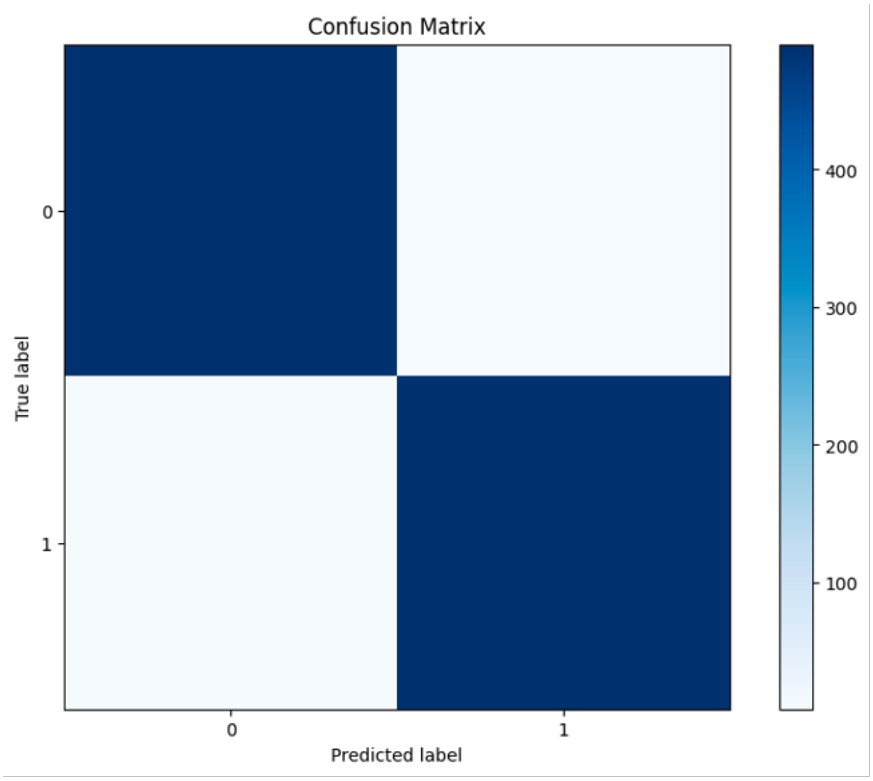

Figure 12: Confusion Matrix of the SVM

8 Discussion

Our study effectively builds on the foundation laid by Erdaw and Tachbele [1], using two models, one using a Convolutional Neural Network (CNN) and the other using a Support Vector Machine (SVM) to classify chest radiograph images for the detection of COVID-19 pneumonia versus other types of pneumonia. The integration of CNNs was crucial for enhanced feature extraction, enabling the identification of subtle but critical variations in pulmonary patterns that are characteristic of different pathogens, such as COVID-19. This advanced extraction of characteristics and the separation of data classes allowed for high classification accuracy, a significant increase in performance compared to 62.8% sensitivity of healthcare professionals. These results underscore the ability of advanced machine learning techniques to significantly improve diagnostic accuracy in medical imaging.

We visually inspected the images that were wrongly classified to analyze any potential trends in their visual features and reasons for wrong classification. Two possible trends were observed. Firstly, it seemed that most of the wrongly classified images were images that had visual features in between different classes. As an example, Figure 13 shows an image that was predicted to be normal (no COVID-19), but its true label was COVID-19. This image shows mild signs of inflammation and is more similar to normal patients. Additionally, the watermarks in the x-ray images could have influenced the classification. The images have different watermarks (possibly suggesting that they were taken with different machines). If for example, one hospital saw a bunch of COVID-19 cases, and thus the model learns the features of these images as being COVID-19, it could think that the watermark is a sign of COVID-19. If a different hospital uses the same machine with the same watermarks, then on normal patients (no COVID-19) it could detect the watermark as COVID-19, wrongly classifying the image. This could be the reason for the wrong classification of Figure 14, which was classified as COVID-19 but actually has true label normal, highly possible due to many COVID-19 images having the same "R" watermark.

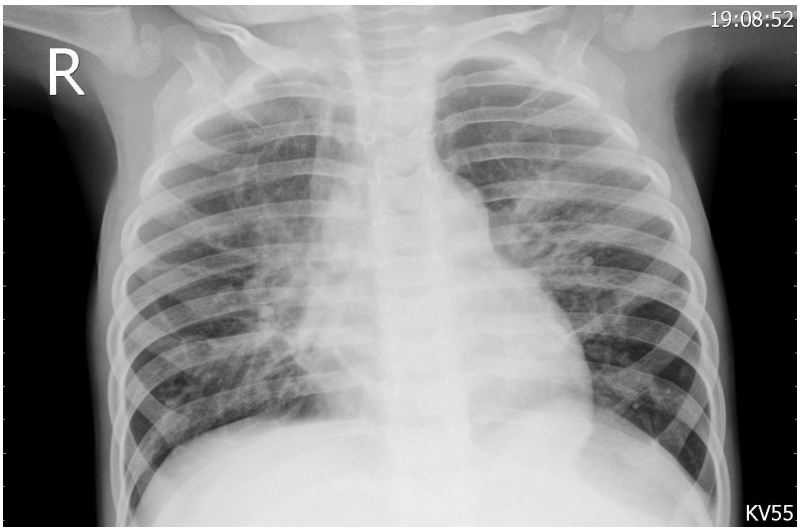

Figure 13: COVID-19 image wrongly classified as normal

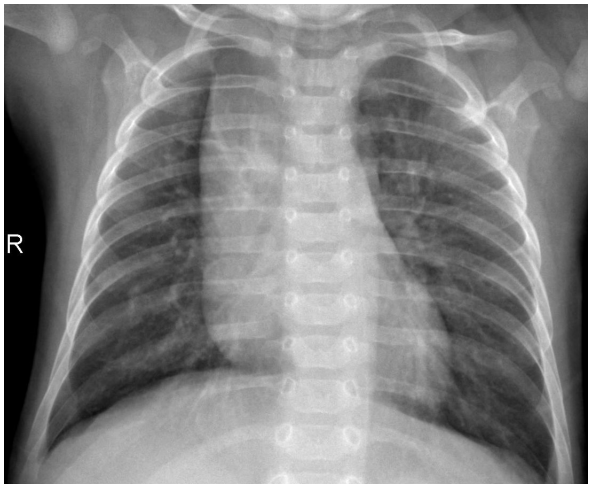

Figure 14: Normal image wrongly classified as COVID-19

Moreover, overfitting remains a critical concern for models like CNNs. To combat this, we used k-fold cross-validation in the testing & training phase to ensure the effectiveness of the model in various data subsets and to prevent its performance from being distorted by any particular partition of the data. As a result, our analysis of the model's loss curves during the epochs suggests that our model is generalizing well. Both training and validation losses demonstrate a downward trend, indicating an effective learning which is minimizing the loss function during back-propagation and therefore increasing its prediction capabilities. Furthermore, there is no upward trend in validation loss after a certain point. This upward trend in the validation loss is a common sign of overfitting because in this case the model learns the training data and its noise too well, not generalizing to unseen data. Therefore, the absence of this upward trend in the validation loss is typically indicative of the model having learned to capture the underlying patterns in the data rather than memorizing the training set. Despite the high validation accuracy achieved, these characteristics of the loss curves lend confidence to our model's ability to generalize to novel data and not be overfit.

Next, future research could extend these results in several ways. First, expanding the data set to include a greater variety of X-ray images and balance the representation of different classes more would likely help to develop a more robust and generalized model. Furthermore, integrating clinical data such as patient demographics, symptoms, and time between presentation with illness and imaging could provide additional information, improving the predictive capability and complexity of these models. Advanced image preprocessing techniques could also be explored to reduce noise and image artifacts more effectively, thus refining the input quality for model training.

The successful implementation of both a Convolutional Neural Network (CNN) and a Support Vector Machine (SVM) to differentiate COVID-19 from people without COVID-19 or other forms of pneumonia through chest X-ray images marks significant advancement and achievement in medical diagnostics. The high accuracy achieved by this model underscores its potential to transform healthcare practices by providing rapid and precise diagnoses, thereby facilitating early and more effective treatment interventions. This capability is

particularly critical for the management of infectious diseases such as COVID-19, where swift diagnosis is essential to contain outbreaks and efficiently manage public health responses.

In addition, the application of such advanced diagnostics can significantly reduce healthcare costs by minimizing the reliance on traditional diagnostic methods that require more labor and human input. This efficiency makes high-level medical care more accessible, especially in remote and underserved areas, potentially equalizing the disparities in healthcare availability. The predictive accuracy of the model also paves the way for personalized medicine, where treatments are tailored to individual genetic profiles and physiological conditions, improving treatment efficacy and reducing side effects. Furthermore, such AI-driven tools can enhance public health surveillance by providing authorities with real-time data on disease trends and epidemiological patterns, thus improving decision making in public health management and resource allocation.

However, the integration of AI into medical diagnostics is not without challenges. Ethical, legal, and privacy concerns, particularly related to personal data handling, algorithm bias, and the lack of human influence in diagnostics, must be seriously considered to ensure these technologies are used responsibly and equitably.

As we continue to refine these technologies, the prospect of smoothly integrating AI into clinical settings offers a promising pathway to a smarter and more responsive healthcare system that is both effective and equitable. This study not only demonstrates the capability of the model in enhancing diagnostic precision but also highlights the broader implications for healthcare systems worldwide, suggesting a substantial impact on how medical care is delivered in the future.

References

[1] Y. Erdaw and E. Tachbele, "Machine Learning Model Applied on Chest X-Ray Images Enables Automatic Detection of COVID-19 Cases with High Accuracy," *International Journal of General Medicine*, vol. 14, pp. 4923–4931, Aug. 2021. Publisher: Dove Press.

[2] "Pneumonia - Diagnosis | NHLBI, NIH," Mar. 2022.

[3] E. Saez de Gordoa, A. Portella, J. Escudero-Fernández, and J. Andreu Soriano, "Usefulness of chest X-rays for detecting COVID 19 pneumonia during the SARS-CoV-2 pandemic," *Radiologia*, vol. 64, no. 4, pp. 310–316, 2022.

[4] A. Tafti and D. W. Byerly, "X-ray Image Acquisition," in *StatPearls*, Treasure Island (FL): StatPearls Publishing, 2024.

[5] D. L. T. MD, "Transfer learning exploits chest-Xray to diagnose COVID-19 pneumonia," Dec. 2020. Section: Medical News.

[6] T. Zebin and S. Rezvy, "COVID-19 detection and disease progression visualization: Deep learning on chest X-rays for classification and coarse localization," *Applied Intelligence*, vol. 51, pp. 1010–1021, Feb. 2021.

[7] L. David, "Activation, Cross-Entropy and Logits."

[8] M. S. Satu, K. Ahammed, M. Z. Abedin, M. A. Rahman, S. M. S. Islam, A. K. M. Azad, S. A. Alyami, and M. A. Moni, "Convolutional Neural Network Model to Detect COVID-19 Patients Utilizing Chest X-ray Images," Mar. 2021. Pages: 2020.06.07.20124594.

[9] "Scikit-learn SVM Tutorial with Python (Support Vector Machines)."

[10] "CoronaHack -Chest X-Ray-Dataset."

[11] "COVID-19 Patients Lungs X Ray Images 10000."

Group Member	Contribution
Benjamin Kleyner	● Identified classification task computational problem and provided formal formulation of ternary vs. binary classification task ● Found datasets for CNN model training ● Managed training of CNN model ● Designed slides and presentation regarding the technical aspects of this process ● Curated statistics regarding efficacy and accuracy of the model ● Authored section of paper detailing technical specifics of the model ● Authored portion of discussion section comparing efficacy of radiologist diagnosis to diagnosis by the CNN ● Worked with Kieran to integrate aspects of research paper selected as inspiration into our project ● Also worked with Kieran to provide necessary data for his analysis of the validity of the model (overfitting/identifying trends in mistaken classifications)
Kieran Carroll	● Found datasets for CNN model novel testing ● Managed set up and foundational code for CNN model ● Designed slides and presentation regarding the technical aspects of CNN model ● Authored portion of introduction section on the current diagnosis of COVID-19 on x-rays by doctors. ● Authored background of the paper ● Authored entire discussion section of the paper ● Authored technical description of sigmoid activation function ● Worked with Ben to integrate aspects of research paper selected as inspiration into our project
Elijah Rosen	● Worked with Haochen to run the SVM model on our dataset

	● Edited and expanded on SVM section in writeup ● Designed presentation style and edited slides for consistency ● Authored section of paper describing potential future medical impacts of this technology
Haochen Yang	● Worked with Elijah to design slides explaining technical aspects of SVM model and its efficacy ● Conducted initial SVM model training ● Worked with Elijah to tune model based on suggestions made in paper used as inspiration and trends seen by visual observation ● Created visuals for the presentation of the technical aspects of the SVM model ● Analyzed accuracy of model ● Generating visuals to represent model's classification accuracy ● Authored section of paper analyzing results of experimentation of SVM model