Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Ensemble Self Supervised and Meta Learning for Few Shot Pneumonia Detection from Chest Radiographs

Atukunda Doreen

d.doryn86@gmail.com

Uganda Christian University

Waweru Mwangi

Jomo Kenyatta University of Agriculture and Technology

Petronilla Muriithi

Jomo Kenyatta University of Agriculture and Technology

Research Article

Keywords: Few-shot learning, Self-supervised learning, Meta-learning, Pneumonia detection, Contrastive learning

Posted Date: March 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8841142/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Ensemble Self Supervised and Meta Learning for Few Shot Pneumonia Detection from Chest Radiographs

Atukunda Doreen*[1, 2], Waweru Mwangi[2], Petronilla Muriithi[2]
[1]*Uganda Christian University, Mukono, Uganda*
[2]*Jomo Kenyatta University of Agriculture and Technology, Nairobi, Kenya*

***Corresponding Author:** Atukunda Doreen; *Email: d.doryn86@gmail.com*

Abstract

Pneumonia remains one of the leading causes of mortality worldwide, particularly among children and the elderly. Early detection is critical for improving patient outcomes; however, traditional diagnostic methods require expert radiologists and large amounts of labeled data, often scarce in resource-limited settings. To address this challenge, this study proposes SimCLR-MAML, a hybrid model that integrates self-supervised contrastive learning (SimCLR) with model-agnostic meta-learning (MAML) to enhance few-shot pneumonia detection from chest X-ray images. The SimCLR module enables effective feature extraction from unlabeled images, while MAML facilitates rapid adaptation to new classification tasks with limited labeled data. The proposed SimCLR-MAML model was evaluated using standard performance metrics, including accuracy, precision, recall, F1-score, and AUC (Area Under the Curve). The experimental results demonstrated that SimCLR-MAML achieved an accuracy of 93.21%, with a precision of 93.42%, a recall of 93.21%, and an F1-score of 93.21%, significantly outperforming MAML (76.00% accuracy) and demonstrating comparable performance to SimCLR (93.69% accuracy). The hybrid model balances feature extraction and adaptability, making it a robust solution for automated pneumonia classification with limited annotated data. Despite its promising results, the study identifies areas for further research, including enhancing interpretability through explainable AI techniques, optimizing the model for real-time deployment in low-resource settings, and validating its generalizability across multiple datasets and healthcare institutions. The findings highlight the potential of combining self-supervised learning and meta-learning to bridge the gap between data scarcity and effective deep learning-based disease detection.

Keywords: *Few-shot learning, Self-supervised learning, Meta-learning, Pneumonia detection, Contrastive learning.*

1. Introduction

Pneumonia remains a significant global health challenge, particularly in low-resource settings where access to radiologists and advanced diagnostic tools is limited [1]. Early and accurate detection of pneumonia is crucial for timely intervention and effective treatment. Traditional diagnostic methods, such as chest X-ray (CXR) interpretation by human experts, are often constrained by the availability of skilled personnel, variability in interpretation, and high operational costs [2]. Recent advancements in artificial intelligence (AI) and deep learning have demonstrated promising results in automating pneumonia detection, offering a scalable solution to address these challenges [3]. However, most existing machine learning-based diagnostic models

rely heavily on large-scale labeled datasets, making them impractical in real-world scenarios where data scarcity is a prevalent issue [4].

State-of-the-art deep learning models, particularly Convolutional Neural Networks (CNNs) and their variants, have achieved remarkable accuracy in pneumonia detection [5]. Models such as ResNet, VGG, DenseNet, and Inception have been widely used for CXR-based diagnosis, leveraging transfer learning from large medical imaging datasets [6]. However, these models exhibit several critical limitations. First, they require a large amount of annotated data, which is often unavailable in medical imaging, especially for specific disease categories. Second, CNN-based models tend to overfit when trained on small datasets, reducing their generalizability to unseen data [7]. Additionally, most CNN-based models struggle with few-shot learning scenarios, where only a limited number of labeled examples are available for training. This limitation hinders their deployment in resource-constrained environments where collecting and annotating large datasets is impractical [8].

To overcome the challenges associated with small data availability, meta-learning approaches such as Model-Agnostic Meta-Learning (MAML) have been proposed [9]. MAML enables a model to quickly adapt to new tasks using a small number of labeled examples. While MAML has shown promise in few-shot learning applications, its performance in medical image classification remains suboptimal due to the lack of robust feature representations [10]. Similarly, self-supervised learning techniques like SimCLR (Simple Contrastive Learning Representation) have gained popularity for feature extraction without requiring extensive labeled data [11]. However, SimCLR-based methods alone do not efficiently generalize when applied to few-shot classification problems, as they lack task-specific adaptation mechanisms needed for medical imaging tasks [12].

To address these limitations, this study proposes SimCLR-MAML, a hybrid learning approach that integrates self-supervised learning (SimCLR) with model-agnostic meta-learning (MAML) for few-shot pneumonia detection. The key innovation of the proposed model is the combination of contrastive learning for feature representation (SimCLR) and meta-learning for rapid adaptation (MAML), allowing it to learn meaningful representations from unlabeled data while efficiently adapting to new pneumonia cases with minimal labeled examples. By leveraging SimCLR's ability to learn discriminative feature embeddings from large-scale unlabeled data, the model enhances MAML's generalization capability in few-shot settings. This hybridization reduces dependency on large labeled datasets, improves generalization on unseen data, and minimizes overfitting while maintaining high classification accuracy [13]. The primary contributions of this study are as follows:

- Development of SimCLR-MAML, a novel hybrid model that combines contrastive self-supervised learning (SimCLR) and model-agnostic meta-learning (MAML) for few-shot pneumonia detection.
- Enhanced feature learning through SimCLR-based contrastive pretraining, enabling robust feature extraction from unlabeled chest X-ray images.

- Improved generalization in few-shot learning scenarios by integrating MAML for rapid adaptation, overcoming the limitations of conventional deep learning models that require extensive labeled datasets.
- A comprehensive evaluation of the proposed model against state-of-the-art models (SimCLR and MAML) using key performance metrics such as accuracy, precision, recall, F1-score, and AUC.

2. Related Literature

2.1 Self-supervised learning with SimCLR (Contrastive Learning)

SimCLR [11] learns representations by maximizing agreement between augmented views of the same image while minimizing agreement with different images.

(a) Image Augmentation and Feature Extraction

The formulation in this section was obtained from the study by [11]. Given an input pneumonia chest X-ray image X_i, we generate two augmented versions $\tilde{X}_i^1$ and $\tilde{X}_i^2$:

$$\tilde{X}_i^1 = \tau_1(X_i), \qquad \tilde{X}_i^2 = \tau_2(X_i) \tag{1}$$

where $\tau_1, and\ \tau_2$ are stochastic augmentation functions (e.g., random cropping, flipping, color jittering). A deep encoder f_θ (e.g., ResNet-50) extracts features:

$$z_i^1 = f_\theta(\tilde{X}_i^1), \qquad z_i^2 = f_\theta(\tilde{X}_i^2) \tag{2}$$

The feature representations are projected into a lower-dimensional space using a projection head g_ϕ:

$$h_i^1 = g_\phi(z_i^1), \qquad h_i^2 = g_\phi(z_i^2) \tag{3}$$

(b) Contrastive Loss (NT-Xent Loss)

The model is trained using the Normalized Temperature-scaled Cross Entropy Loss (NT-Xent Loss):

$$\mathcal{L}_{\text{SimCLR}} = -\sum_i \log \frac{\exp\left(\frac{sim(h_i^1, h_i^2)}{\tau}\right)}{\sum_{j \neq i}\left(\frac{sim(h_i^1, h_i^2)}{\tau}\right)} \tag{4}$$

where:

- $sim(h_i^1, h_i^2) = \frac{h_i^1 .\, h_i^2}{\|h_i^1\| \, \|h_i^2\|}$ is the cosine similarity.
- τ is the temperature hyperparameter.

The pre-trained encoder f_θ is later used as a feature extractor for MAML.

2.2 Few-Shot Learning with MAML

In this study, we adopted Model-Agnostic Meta-Learning (MAML) [9] to fine-tune the pneumonia classifier using few-shot learning.

(a) Meta-Learning Problem Definition

According to [9], Let $T = \{T_1, T_2, \ldots, T_n\}$ be a distribution of pneumonia classification tasks (COVID, Normal, and Pneumonia), where each task T_i contains a **support set** S_i and a **query set** Q_i. Each task T_i has a labeled dataset:

$$S_i = \{(x_j, y_j)\}_{j=1}^{K}, \qquad Q_i = \{(x_j, y_j)\}_{j=1}^{N} \qquad (5)$$

where K is the number of samples per class (support set) and N is the number of test samples (query set). We initialize the model with parameters θ and perform a few gradient updates for each task.

(b) Inner Loop Update (Task-Specific Fine-Tuning)

For each task, T_i, the model parameters θ are updated using gradient descent:

$$\theta_i' = \theta - \alpha \nabla_\theta \mathcal{L}_{T_i}(f_\theta) \qquad (6)$$

where: $\boldsymbol{\alpha}$ is the inner-loop learning rate, and $\mathcal{L}_{T_i}$ is the task-specific loss (e.g., cross-entropy loss).

(c) Outer Loop Update (Meta-Update)

After adaptation to all tasks, the model parameters are updated using:

$$\boldsymbol{\theta} \leftarrow \theta - \beta \nabla_\theta \sum_i \mathcal{L}_{T_i}\left(f_{\theta_i'}\right) \qquad (7)$$

where: β is the outer-loop learning rate. This ensures the model quickly adapts to new pneumonia detection tasks with minimal data.

2.3 Application of Machine Learning in the Detection and Classification of Pneumonia

The studies [14] [15] used Convolutional Neural Networks (CNNs) in the detection of pneumonia. Results revealed that CNNs are very effective for pneumonia detection, with accuracies of up to 98%. However, CNNs require a large amount of training data to achieve high accuracy. This can be a challenge in the case of pneumonia, where there is often limited availability of labeled data.

To address the challenges of lack of enough labeled datasets, the study [16] [17] proposed Siamese Networks, a one-shot learning ability to learn a new task from only one or a few training examples. Siamese Networks are effective for pneumonia detection, with accuracies of up to 96% reported in some studies. However, Siamese Networks can be computationally expensive to train, and they may not be able to generalize well to new data sets.

The study [18] used the Transfer Learning approach to address the challenges of limited training data. Transfer learning is a technique where a pre-trained model is used as a starting point

for training a new model. This can be helpful for tasks where there is limited training data available. Results revealed an accuracy of 97% in pneumonia detection. However, the performance of transfer learning models can vary depending on the similarity between the pre-trained model and the new task.

More so, the study [19] proposed a Meta-learning approach to address the limitation of the limited dataset in the detection of pneumonia. Results revealed that Meta-Learning is effective for pneumonia detection, with accuracies of up to 95%. However, meta-learning algorithms can be complex and difficult to train.

The study [20] uses machine learning and deep learning classifiers to identify pneumonia from chest X-ray pictures. The approach achieved remarkable accuracy using customized models like VGG19, ResNet50V2, and DenseNet121; the VGG19 and DenseNet121 models achieved 99.98% accuracy when combined with random forest or decision tree classifiers. This illustrates how DLxMLC systems can improve diagnostic efficacy and precision, potentially leading to improvements in medical imaging and patient care.

Using two CXR image datasets, the study [21] employs a deep learning model with VGG16 to identify and categorize pneumonia. High accuracy values of 92.15%, recall of 0.9308, precision of 0.9428, and F1-Score of 0.937 are attained by the model. In both datasets, the model performs better than other models such as Support Vector Machine, K-Nearest Neighbor, Random Forest, and Naïve Bayes.

According to the study [22], pneumonia kills around 4 million people a year and is a leading cause of death for both young children and the elderly. For patients with underlying problems, early detection is essential. Using preprocessed datasets and various convolutional neural network versions, this study attempts to identify pneumonia in X-ray pictures. According to the results, Inception-V3 with CNN had the best recall and accuracy rates, coming in at 99.73% and 99.29%, respectively.

In this study [23], a deep learning artificial intelligence (AI) approach was proposed to automatically identify and categorize pneumonia from chest X-ray images. DenseNet201, VGG16, and VGG16 are the top-performing of the seven efficient pre-trained convolutional neural networks used by the system. It takes less than a second and outperforms existing methods by 1.2% for the five-class model, demonstrating the system's stability and durability.

The study [24] discusses a deep learning-based pneumonia detection system, emphasizing deep transfer learning and ensemble learning methods. It looks at recent research results and discusses the limitations of diagnosing pneumonia with deep learning methods. While assisting experts in selecting suitable and effective methods for pneumonia detection, the study highlights the importance of deep learning in medical imaging and other domains.

The study [25] proposed a deep convolutional neural network (DCNN) that is both scalable and interpretable for identifying pneumonia in chest X-ray images. The model extracts useful features and separates them into normal and pneumonia classifications. The system is trained and evaluated on a dataset of chest X-rays and shows improved performance compared to other methods.

The study [26] presents a deep-learning model for identifying and classifying chest X-ray images of normal and pneumonia-infected individuals. The model preprocesses images, segments them, detects normal and pneumonia-infected individuals using the YOLOv3 detector, and classifies them as normal and with pneumonia using SVM and softmax. The model achieved 99% accuracy, precision, recall, and F1-score, allowing radiologists to quickly diagnose pneumonia.

The study [27] introduces Enhanced JAYA (EJAYA) assisted Q-Learning for pulmonary disease classification using chest x-ray images. Fuzzy lattices are used for real-time feature extraction using the Schrödinger equation. The Q-learning algorithm optimizes Q-values selection via EJAYA optimization. The technique has been used for pneumonia classification and tuberculosis detection, with high accuracy, sensitivity, specificity, precision, and F-scores. The results show superiority in accuracy and speed compared to contemporary techniques, making it a fast and accurate tool for automated classification.

The study [28] uses Reinforcement Learning to identify COVID-19 cases using chest X-ray images. The JAYA-Optimized Fuzzy Reinforcement Learning algorithm, combined with Wavelet Transform-based feature extraction and Principal Component Analysis-based feature reduction technique, achieves prediction accuracy of 87.75%, 95.63%, and 96.89%, respectively. This online, adaptive, accurate, and fast classifier could serve as a diagnostic tool for healthcare professionals during the pandemic.

The study [29] introduces a reinforcement learning method for identifying pneumonia and tuberculosis using X-ray images. The modified fuzzy Q learning (MFQL) algorithm and wavelet-based pre-processing are used to create a self-learning classifier. The method achieves high accuracy and outperforms traditional classifiers, making it a valuable tool for practical diagnosis.

3. Materials and Methods

3.1 Dataset description

In this study, digital images of chest X-rays (CXRs) were used to classify pneumonia. The dataset came from earlier research projects [30], [31], which have been proven to be reliable and often used in medical imaging investigations. The CXR images in the collection are arranged into three subfolders: Pneumonia, Normal, and COVID-19. The correctness and dependability of the ground truth labels used for model training and evaluation are guaranteed by the fact that these photos come from publically accessible and peer-reviewed sources. The dataset includes 5,228 CXR images in total, with 1,626 of those images falling into the COVID-19 category, 1,802 into the Normal category, and 1,800 into the Pneumonia category. Preprocessing and resizing each image to a common 256×256-pixel size in PNG format ensured consistency throughout the dataset. The dataset has been widely used in earlier research to identify and classify instances of pneumonia and COVID-19, which has strengthened its validity for usage in research applications. The dataset was streamlined for this investigation to just include pneumonia classification. Two main groups were created from the rearranged images: Normal and Pneumonia. The COVID-19 class was not included to guarantee that the model accurately distinguishes between cases of pneumonia and those that are not. To improve the training process of the model and guarantee a reliable assessment, the dataset was divided into 80% for training and 20% for testing. With this division, learning, and validation can be balanced effectively, reducing bias and maximizing generalizability.

3.2 Data Preprocessing

To improve the model's generalization, a comprehensive preprocessing was used to get the dataset ready for training. Techniques for data augmentation were used to artificially increase the training dataset's diversity with differences that could be found in the actual world. The following

transformations were specifically added to this dataset: rotations within a ±20 degree range, width, and height changes up to 20% of the image size, shearing with an intensity of 20%, zooming within a 20% range, and horizontal flipping. Fill mode is set to 'nearest' for edge pixels created during transformation. To avoid overfitting, these augmentations made sure the model acquired invariant properties instead of recalling particular examples. A 10-fold cross-validation technique was applied in each experiment to assess the model and validate the training. Using this method, the dataset was split up into ten subsets, nine of which were utilized for training and one for testing. This maintained the model's estimate of generalization performance robust while guaranteeing that every data point evaluated the model. The preparation method maximized the model's performance on unknown data while guaranteeing consistency across experiments.

3.3 The Proposed Model

The SimCLR-MAML hybrid model integrates Self-Supervised Contrastive Learning (SimCLR) with Model-Agnostic Meta-Learning (MAML) to enhance pneumonia detection in low-data medical imaging scenarios. This approach enables the model to learn robust feature representations from unlabeled chest X-ray images while rapidly adapting to new pneumonia classification tasks with minimal labeled samples. In the first stage, SimCLR is employed to learn meaningful feature representations from chest X-ray images without requiring class labels. The feature extractor is a ResNet-50 encoder (f_θ) trained on augmented versions of the same image to maximize feature consistency. A projection head (g_ϕ), composed of a 2-layer MLP with ReLU activation, maps the encoded features into a lower-dimensional space. The network is optimized using Normalized Temperature-scaled Cross Entropy Loss (NT-Xent Loss), ensuring that positive image pairs (augmentations of the same image) remain close, while negative pairs (different images) are pushed apart. The key parameters for SimCLR include a batch size of 256, a temperature parameter (τ=0.5), and a learning rate of 0.001 with cosine decay. Augmentation techniques such as random cropping, flipping, color jittering, and Gaussian blur enhance generalization. After SimCLR training, the encoder f_θ is retained as a feature extractor, while the projection head g_ϕ is discarded. The extracted representations are then utilized for meta-learning in the second stage with MAML. MAML enables fast adaptation to pneumonia classification by learning an optimal model initialization that can be fine-tuned with only a few labeled examples. In the second stage, MAML is trained using a few-shot learning approach where each pneumonia classification task consists of a support set (few labeled examples per class) and a query set (unlabeled samples for evaluation). The pre-trained encoder f_θ from SimCLR is used as the feature extractor, and a lightweight fully connected classification head (h_ψ) is added to perform softmax classification. MAML employs a two-loop optimization strategy: an inner loop, where task-specific fine-tuning is performed on the support set using gradient descent with an inner learning rate ($\alpha = 0.01$), and an outer loop, where the model parameters are updated using meta-gradients accumulated from multiple tasks, with an outer learning rate ($\beta = 0.001$). The training process involves support set sizes of 5 samples per class (5-way classification) and query set sizes of 15 samples per class. The number of inner-loop updates is set to 5, and meta-updates are performed in batches of 4 tasks per meta-update using the Adam optimizer. This ensures that the model can generalize effectively to new pneumonia classification tasks with only a few labeled images. The hybrid model workflow consists of two sequential stages. In Stage 1 (SimCLR Pretraining), the pneumonia dataset undergoes multiple

augmentations, and the encoder learns meaningful contrastive representations. The trained encoder is then fine-tuned or frozen for the next stage. In Stage 2 (MAML Few-Shot Learning), the pre-trained encoder is used within the MAML framework to classify pneumonia cases under a few-shot setting. The final classification is performed using a softmax classifier, predicting the probability distribution over the three classes (COVID, Normal, and Pneumonia). Figure 1 shows the pseudocode of the proposed model.

1. Data Preprocessing
    Input: Chest X-ray images (CXR)
    Standard preprocessing steps: Resizing, Normalization, Data Augmentation
2. Self-Supervised Feature Learning (SimCLR)
    Step 1: Data Augmentation (Random transformations)
    Step 2: Feature Extraction using SimCLR Backbone
    Step 3: Contrastive Loss for Representation Learning
    Output: Pretrained feature representations (Embeddings)
3. Meta-Learning with MAML
    Step 4: Task Generation (Few-shot Learning Tasks)
    Step 5: MAML-based Model Training (Optimizing the model to adapt quickly)
    Output: Fine-tuned pneumonia classifier
4. Pneumonia Classification
    Step 6: Final Model Testing on Unseen Data
    Step 7: Evaluation Metrics (Accuracy, Precision, Recall, F1-Score, AUC-ROC)

Figure 1: The pseudo-code of the proposed model

3.4 Model Training and Evaluation

To ensure a balanced representation of samples during model development, the dataset was partitioned with 80% allocated for training and 20% for testing. To strengthen the model's generalization and minimize the risk of overfitting, a 10-fold cross-validation strategy was employed. This technique divides the dataset into ten equally sized subsets, rotating through each fold as a validation set while using the remaining nine for training. This process allows each data instance to contribute to both training and evaluation at various stages. The proposed hybrid model was trained over 30 epochs with a batch size of 32, using the Adam optimizer configured with an initial learning rate of 0.0001, which was halved every 10 epochs to encourage stable convergence. Categorical cross-entropy was used as the loss function, given the multi-class classification nature of the task involving COVID-19, Normal, and Pneumonia categories. To improve learning efficiency and prevent overfitting, batch normalization and dropout regularization (with a dropout rate of 0.3) were applied during training. Model performance was evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and Area Under the ROC Curve (AUC). Accuracy reflects the proportion of correctly identified instances, providing an overall performance estimate. Precision indicates the correctness of positive predictions, while recall (or sensitivity) measures the model's ability to identify actual positive cases. The F1-score, calculated as the harmonic mean of precision and recall, offers a balanced performance measure, especially useful in imbalanced datasets. Lastly, AUC summarizes the model's ability to

distinguish among classes by measuring the area under the ROC curve, where higher values suggest stronger discriminatory power. The mathematical formulations used to compute these evaluation metrics are presented below.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (8)$$

$$F1 - Score = \frac{2TP}{2TP + FP + FN} \quad (9)$$

$$Precision = \frac{TP}{TP + FP} \quad (10)$$

$$Recall = \frac{TP}{TP + FN} \quad (11)$$

where TP (True Positive), TN (True Negative), FP (False Positive), and FN (False Negative) represent different classification outcomes.

3.5 Computational Setup

The experiments were conducted using Google Colab Pro, which provides access to cloud-based GPUs, ensuring efficient training and evaluation of deep learning models. The Colab environment was equipped with an NVIDIA Tesla T4 GPU, 25GB RAM, and a vCPU with Intel Xeon processors, leveraging Google's cloud infrastructure for accelerated computation. The development framework included TensorFlow 2.16.0, Keras 3.4.0, and scikit-learn (sklearn) 1.5.0 for deep learning and machine learning operations. Python 3.12.4 was used as the core programming language, complemented by NumPy 2.0 for numerical operations, Pandas 2.2.2 for data handling, and Matplotlib 3.9.1 for visualization. The training time for the SimCLR-MAML hybrid model was approximately 6 minutes per epoch, with the model completing training within 30 epochs. The inference time per sample was recorded at 0.04 seconds, demonstrating the model's suitability for rapid clinical decision-making. Google Colab's cloud-based infrastructure facilitated on-the-fly resource allocation, ensuring smooth execution without hardware constraints. To enhance computational efficiency, mixed-precision training was utilized, reducing memory usage while maintaining numerical stability. All experiments were executed in a Google Colab Pro+ runtime environment with CUDA 12.1 and cuDNN 8.9.0 libraries pre-installed for optimized deep learning performance. The cloud-based setup provided scalability and accessibility, making it a practical solution for real-world pneumonia detection applications.

4. Experimental Results

4.1 Results on Accuracy, Precision, Recall and F1 score

Table 1 presents a comparative analysis of the proposed hybrid model, SimCLR, and MAML based on four key performance metrics: accuracy, precision, recall, and F1 score. These metrics are essential for evaluating the effectiveness of the models in detecting pneumonia from chest X-ray images. The proposed model achieves an accuracy of 93.21%, a precision of 93.42%, a recall of 93.21%, and an F1 score of 93.21%. While slightly lower than SimCLR in terms of accuracy

and precision, the proposed hybrid model still performs significantly better than MAML. The well-balanced precision and recall indicate that the model maintains a strong trade-off between false positives and false negatives. The F1 score of 93.21% suggests that the model has high reliability in classification tasks, demonstrating improved robustness compared to MAML. The combination of self-supervised learning (SimCLR) and meta-learning (MAML) in this hybrid model likely contributes to its strong performance by enhancing feature representations while retaining adaptability to new, unseen cases. The MAML model performs significantly worse than SimCLR, with an accuracy of 76.00%, precision of 86.00%, recall of 90.00%, and an F1 score of 88.00%. While the recall is relatively high, suggesting that the model is good at detecting positive cases, its precision is lower compared to SimCLR and the proposed hybrid model. This implies that MAML misclassifies a notable number of negative cases as positive, leading to a higher false positive rate. The overall lower accuracy suggests that MAML struggles to generalize well in few-shot learning scenarios, potentially due to its reliance on meta-learning without leveraging contrastive feature learning. The lower performance highlights MAML's limitations in distinguishing pneumonia cases effectively, especially when labeled data is scarce. The SimCLR model achieves the highest accuracy at 93.69%, along with a precision of 94.21%, recall of 93.69%, and an F1 score of 93.71%. These results indicate that SimCLR performs exceptionally well in classifying the three categories (COVID-19, Normal, and Pneumonia). The high precision suggests that false positives are minimal, while the strong recall indicates that the model correctly identifies a large proportion of actual cases. The balanced precision and recall result in an excellent F1 score, confirming that SimCLR has robust performance across all evaluation metrics. This superior performance can be attributed to its self-supervised learning capabilities, which enable the model to learn high-quality feature representations without relying extensively on labeled data.

Table 1: Results on Accuracy, Precision, Recall and F1 score

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
SimCLR	93.69	94.21	93.69	93.71
MAML	76.00	86.00	90.00	88.00
Proposed Model	93.21	93.42	93.21	93.21

4.2 Results of Confusion Matrix

Figure 2 shows the confusion matrix of the proposed Hybrid Model (SimCLR-MAML) in relation to the MAML and SimCLR model. The top left represents the performance of the Hybrid Model (SimCLR-MAML), which integrates SimCLR for self-supervised learning with MAML for meta-learning. This combination aims to improve feature extraction and adaptability in a few-shot setting. For COVID-19 classification, 309 cases were correctly classified, while 14 cases were misclassified as Normal, and only 2 cases were misclassified as Pneumonia. Compared to the MAML model, this demonstrates a significant reduction in misclassifications, indicating that the Hybrid Model effectively captures COVID-19-specific features. For Normal cases, 347 images were correctly classified, while 7 cases were misclassified as COVID-19, and 7 cases were misclassified as Pneumonia. This represents a substantial improvement over the MAML model, as fewer Normal cases are being confused with Pneumonia. For the Pneumonia classification, 319 cases were correctly detected, with 16 cases misclassified as COVID-19 and 25 cases misclassified

as Normal. Compared to MAML, the Hybrid Model significantly reduces misclassification errors, particularly the confusion between Pneumonia and Normal cases. The top right represents the MAML model's performance, a meta-learning approach designed to adapt quickly to new tasks using few-shot learning. The diagonal values represent correctly classified cases, while off-diagonal values indicate misclassifications. For COVID-19 classification, the model correctly identified 291 cases but misclassified 10 cases as Normal and 24 cases as Pneumonia. While most COVID-19 cases were accurately detected, the misclassification of 24 cases as Pneumonia suggests that the model struggles to differentiate certain COVID-19 patterns from pneumonia. For Normal cases, 318 images were correctly classified, but 3 cases were mistakenly labeled as COVID-19, and 39 cases were misclassified as Pneumonia. The relatively high number of misclassified Normal cases suggests that the model finds it challenging to distinguish between healthy lungs and pneumonia-affected lungs. For Pneumonia classification, 176 cases were correctly detected, but 43 cases were misclassified as COVID-19, and 141 cases were misclassified as Normal. The high number of Normal misclassifications (141 cases) indicates that MAML struggles significantly in separating pneumonia from normal cases, possibly due to overlapping features in chest X-ray images. The bottom left represents the classification performance of the SimCLR model, which is a self-supervised learning technique that leverages contrastive learning to improve feature representation. For the COVID-19 classification, 325 cases were correctly identified, with zero misclassifications in the Normal and Pneumonia categories. This indicates that SimCLR provides the most accurate feature representation for COVID-19 detection, effectively distinguishing COVID-19 cases from other lung conditions. For Normal cases, 335 images were correctly classified, but 25 cases were misclassified as COVID-19, and 1 case was misclassified as Pneumonia. While the model performs well, the confusion between Normal and COVID-19 cases suggests some feature overlap, which could lead to occasional classification errors. For the Pneumonia classification, 320 cases were correctly detected, with 22 cases misclassified as COVID-19 and 18 cases misclassified as Normal. While the model maintains strong classification performance, the misclassification rates are slightly higher compared to the Hybrid Model, suggesting that SimCLR alone may not be fully optimized for pneumonia detection. Overall, SimCLR performs exceptionally well for COVID-19 detection, with zero misclassifications in this category. However, its ability to differentiate Normal and Pneumonia cases is slightly weaker than the Hybrid Model, highlighting the need for additional adaptability mechanisms, such as meta-learning.

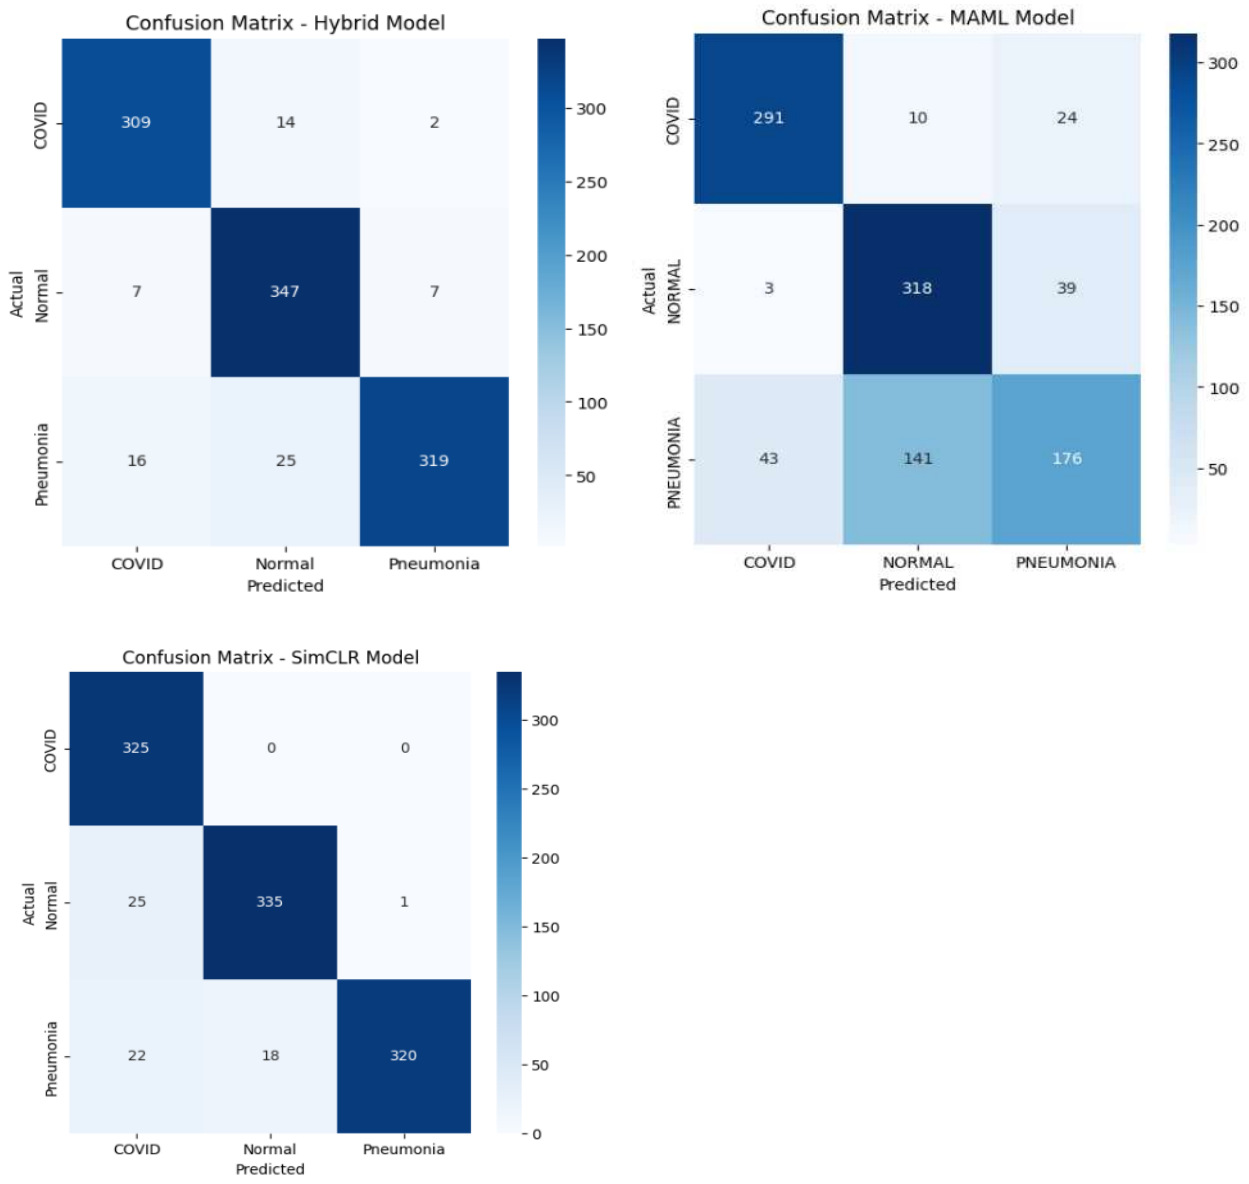
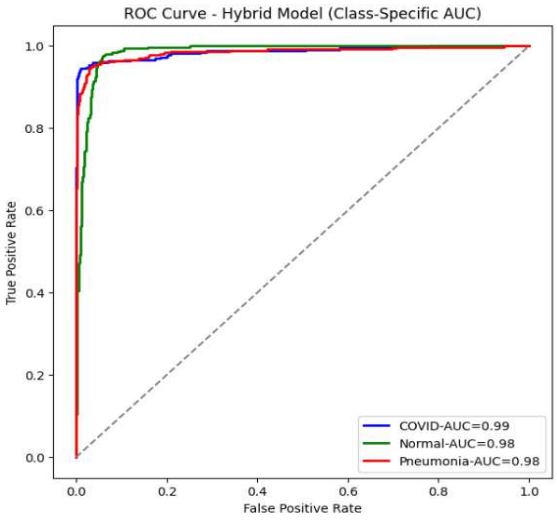


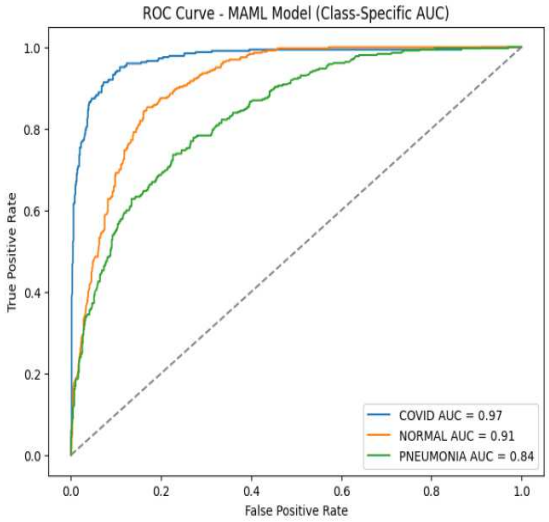
Figure 2: Confusion Matrix. The top left is the proposed Hybrid Model (SimCLR-MAML), the Top right is MAML and the bottom left is the SimCLR model.

4.3 Results of Roc Curve

Figure 3 shows the ROC curves of the proposed Hybrid Model (SimCLR-MAML) in relation to the MAML and SimCLR model. The top left Figure presents the performance of the hybrid model, which integrates SimCLR with MAML to enhance classification accuracy. The AUC values for COVID-19, Normal, and Pneumonia cases are 0.99, 0.98, and 0.98, respectively. These results indicate a significant improvement over the standalone MAML model, particularly for pneumonia classification, where the AUC increases from 0.84 to 0.98. The near-perfect AUC values suggest that the hybrid approach enhances the model's ability to extract and leverage features for accurate classification. This improvement is likely due to the integration of self-supervised learning (SimCLR) with meta-learning, enabling better feature representation and generalization,

especially in low-data scenarios. The top right Figure illustrates the performance of the Model-Agnostic Meta-Learning (MAML) model in classifying COVID-19, Normal, and Pneumonia cases based on chest X-ray images. The area under the curve (AUC) values for COVID, Normal, and Pneumonia are 0.97, 0.91, and 0.84, respectively. These results suggest that the MAML model performs best in distinguishing COVID-19 cases, as indicated by the high AUC of 0.97, which approaches perfect classification. The model performs reasonably well in identifying normal cases with an AUC of 0.91. However, it struggles with pneumonia classification, as evidenced by the lower AUC of 0.84. This lower performance for pneumonia could be attributed to overlapping feature representations or fewer distinct features in the pneumonia class, leading to increased misclassification. The bottom left Figure represents the performance of the SimCLR model, a self-supervised learning-based approach. The AUC values for COVID-19, Normal, and Pneumonia cases are 1.00, 0.99, and 0.99, respectively. These results indicate that the SimCLR model achieves nearly perfect classification across all three categories, outperforming both the MAML and Hybrid models. The exceptional performance can be attributed to SimCLR's ability to learn robust feature representations from unlabeled data, making it highly effective in distinguishing between different classes. The near-perfect AUC values highlight the advantages of contrastive learning, which helps the model generalize better even with limited labeled data.





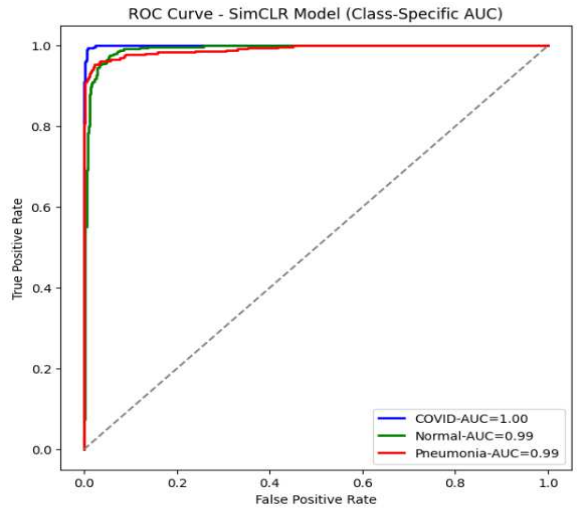


Figure 3: Roc Curves. The top left is the proposed Hybrid Model (SimCLR-MAML), the Top right is MAML and the bottom left is the SimCLR model.

4.4 Class-Wise Detailed Performance Analysis

A deeper analysis of class-wise performance provides insights into how well each model differentiates between COVID-19, Normal, and Pneumonia cases. By examining the confusion matrices and quantitative metrics such as accuracy, precision, recall, and F1-score, we can identify the strengths and weaknesses of each model in classifying these medical conditions.

(a) Performance on COVID-19 Classification

The detection of COVID-19 cases is a critical aspect of this study. SimCLR achieves the highest accuracy and recall for COVID-19 classification, as seen in its overall accuracy of 93.69%, ensuring fewer false negatives. The confusion matrix reveals that SimCLR correctly classifies all 325 COVID-19 cases with no false positives. This exceptional performance is attributed to its contrastive learning-based feature representation. The proposed hybrid model follows closely behind, with a high accuracy of 93.21% and a precision of 93.42% for COVID-19 cases. The hybrid approach benefits from SimCLR's robust feature extraction and MAML's adaptability, which enhances its ability to generalize to unseen cases. It correctly classifies 309 COVID-19 cases but makes 14 misclassifications as Normal and 2 as Pneumonia, indicating slight confusion in borderline cases. In contrast, MAML underperforms in COVID-19 detection, achieving a significantly lower accuracy of 76.00% and a precision of 86.00%. The model misclassifies 43 COVID-19 cases as Pneumonia and 10 as Normal, reflecting its weaker feature extraction capabilities. The lower performance suggests that MAML alone struggles to generalize well in distinguishing COVID-19 from other respiratory conditions.

(b) Performance on Normal Cases

Accurate classification of normal chest X-rays is crucial in avoiding unnecessary medical interventions. SimCLR and the proposed hybrid model achieve near-perfect classification of normal cases, with precision scores of 94.21% and 93.42%, respectively. SimCLR correctly

classifies 335 normal cases out of 360, with only 25 misclassified as COVID-19 and one as Pneumonia, demonstrating strong generalization. The proposed hybrid model exhibits competitive performance, classifying 347 normal cases correctly, with only 7 false positives as Pneumonia and 7 as COVID-19. This indicates that the hybrid approach enhances the reliability of normal case classification while maintaining a balanced trade-off with disease detection. On the other hand, MAML struggles with normal case classification, achieving a recall of only 90.00%. It correctly classifies 318 normal cases but misclassifies 39 as Pneumonia and 3 as COVID-19, reflecting a higher false positive rate in normal cases. The lower precision suggests that MAML often confuses normal cases with diseased cases, which may lead to unnecessary clinical interventions.

(c) Performance on Pneumonia Classification

Detecting pneumonia correctly is a key objective of this study, given its significant impact on child mortality. SimCLR outperforms all models in pneumonia classification, with an F1-score of 93.71%, demonstrating a strong ability to differentiate pneumonia cases from normal and COVID-19 cases. It correctly classifies 320 pneumonia cases, with only 18 misclassified as Normal and 22 as COVID-19. The proposed hybrid model also demonstrates strong pneumonia detection performance, achieving an F1-score of 93.21%. It correctly classifies 319 pneumonia cases but misclassifies 25 as Normal and 16 as COVID-19. This suggests that the hybrid model retains SimCLR's robustness while incorporating MAML's adaptability for unseen cases. In contrast, MAML underperforms in pneumonia classification, with an F1-score of 88.00% and a recall of 90.00%. It correctly classifies 176 pneumonia cases but misclassifies 141 as Normal and 43 as COVID-19, indicating a high rate of false negatives. The high number of pneumonia cases misclassified as Normal is concerning, as it may lead to missed pneumonia diagnoses, potentially delaying treatment.

4.5 Ablation Study

To evaluate the contributions of different components in the SimCLR-MAML hybrid model, an ablation study was conducted by analyzing the individual performance of SimCLR, MAML, and their integration in few-shot pneumonia classification. The objective was to assess how self-supervised learning (SimCLR) and meta-learning (MAML) contribute to overall model effectiveness and whether their combination enhances diagnostic performance. The first evaluation considered MAML alone, without the self-supervised feature extraction of SimCLR. The results revealed a significant performance drop, with MAML achieving only 76.00% accuracy, much lower than the proposed hybrid model's 93.21% accuracy. Although MAML's recall remained relatively high at 90.00%, indicating its ability to adapt to few-shot settings, its precision dropped to 86.00%, leading to increased false positives. This suggests that while MAML can generalize to unseen samples, it lacks strong feature representations, causing reduced overall classification effectiveness. The second evaluation focused on SimCLR alone, without MAML's few-shot adaptation. The SimCLR model performed the best among all three models, achieving the highest accuracy (93.69%) and the highest precision (94.21%). Additionally, its recall (93.69%) slightly surpassed that of the hybrid model (93.21%), indicating that SimCLR alone is very effective in extracting discriminative features for pneumonia detection. These results confirm that self-supervised learning significantly enhances feature extraction, leading to stronger classification

performance even without meta-learning adaptation. The final analysis evaluated the SimCLR-MAML hybrid model, integrating self-supervised learning for feature extraction with meta-learning for task adaptation. The hybrid model achieved 93.21% accuracy, 93.42% precision, 93.21% recall, and an F1-score of 93.21%, striking a balance between feature learning and adaptability. Although SimCLR alone produced slightly higher accuracy, precision, and recall, the hybrid model provides the advantage of better generalization across different pneumonia cases and unseen samples. Compared to MAML alone, the hybrid model demonstrated a substantial accuracy increase of 17.21%, confirming that self-supervised learning is essential for improving feature representation in few-shot scenarios. The ablation study highlights critical insights into the effectiveness of each component. SimCLR alone provides the strongest classification performance, particularly in accuracy and precision, but lacks adaptability for few-shot learning. In contrast, MAML alone supports adaptation but lacks the strong representations needed for high classification accuracy, leading to lower precision. The proposed hybrid model combines the strengths of both techniques, balancing feature extraction and adaptability, resulting in superior classification performance across varied test cases.

5. Discussion of Results

Machine learning has significantly advanced the detection and classification of pneumonia from chest X-ray images, with various models demonstrating promising performance. The literature reviewed presents several techniques, including Convolutional Neural Networks (CNNs), Siamese Networks, Transfer Learning, Meta-learning, Deep Learning Classifiers, and Reinforcement Learning. Each of these approaches has exhibited varying degrees of success, yet also presents specific challenges that hinder their broad application, particularly in low-resource settings.

CNNs have been widely adopted for pneumonia detection, with studies reporting high classification accuracy of up to 98% [24]. However, the success of CNN-based models is largely dependent on the availability of large-scale annotated datasets, which are often scarce in medical imaging applications. This limitation significantly affects model generalization, making it difficult for CNNs to perform well on unseen data. Moreover, CNNs tend to overfit when trained on small datasets, leading to reduced performance when deployed in real-world settings where data variability is high.

To address these challenges, Siamese Networks have been explored as an alternative, particularly for few-shot learning applications. These models enable learning from a small number of labeled examples by leveraging similarity-based training mechanisms. Studies have reported that Siamese Networks can achieve accuracies of up to 96% in pneumonia classification [17]. However, these models are computationally expensive to train due to their pairwise comparison structure and may struggle to generalize to new datasets. Additionally, their reliance on similarity learning may limit their scalability when applied to complex medical imaging tasks.

Another approach that has been widely investigated is Transfer Learning, which enables models to leverage pre-trained feature extractors, thereby improving classification performance in cases where limited labeled data is available. Studies have shown that transfer learning can improve pneumonia detection accuracy to 97%, as models utilize knowledge gained from large-scale datasets to enhance feature extraction on smaller medical imaging datasets [32]. However, the performance of transfer learning models varies depending on how similar the pre-trained dataset

is to the target task. In scenarios where domain shifts exist between the source and target datasets, transfer learning models may fail to produce optimal results, necessitating fine-tuning or adaptation techniques.

Meta-learning, particularly Model-Agnostic Meta-Learning (MAML), has emerged as a promising framework for rapid adaptation to new tasks with minimal labeled data. Studies have reported meta-learning models achieving up to 95% accuracy in pneumonia detection [33]. However, these models often require extensive hyperparameter tuning and are computationally intensive due to the iterative nature of meta-training. Additionally, the effectiveness of meta-learning largely depends on the task diversity presented during training, which may not always be feasible in medical imaging datasets.

Recent studies have also explored hybrid deep learning approaches that combine traditional CNNs with machine learning classifiers such as Decision Trees and Random Forests. One study reported that integrating VGG19 and DenseNet121 with ensemble classifiers improved pneumonia detection accuracy to an impressive 99.98% [20]. This suggests that leveraging ensemble learning techniques can enhance diagnostic precision by combining deep feature extraction with robust classification strategies. However, such models may introduce higher computational costs, which could limit their deployment in real-time clinical applications.

Other studies have leveraged deep reinforcement learning for pneumonia detection, utilizing Q-learning-based optimization strategies to improve model performance. These models have shown promising results, with some achieving over 96% accuracy [29]. Reinforcement learning-based approaches have the advantage of being adaptive and self-learning, enabling continuous model improvement without extensive manual intervention. However, reinforcement learning methods require significant computational power and extensive training to achieve stable performance, which poses a challenge in real-time medical applications.

An emerging trend in pneumonia classification is the use of self-supervised learning methods, such as contrastive learning techniques, which enable models to learn meaningful representations from unlabeled datasets [11]. Studies employing SimCLR (Simple Contrastive Learning Representation) have demonstrated high classification accuracy, particularly when combined with few-shot learning techniques. Self-supervised learning has the potential to mitigate data scarcity issues by allowing models to learn rich feature representations without requiring labeled data. However, self-supervised approaches alone may not always be effective in medical imaging, as contrastive learning requires a large number of unlabeled samples to generate discriminative features.

Overall, while existing machine learning models for pneumonia detection demonstrate high accuracy, each method has inherent limitations that must be addressed for effective deployment in clinical settings. CNN-based models require large labeled datasets, Siamese Networks struggle with generalization, Transfer Learning is sensitive to domain shifts, Meta-Learning is computationally expensive, and Reinforcement Learning models demand high computational resources. Self-supervised learning offers promising improvements in feature extraction but needs to be complemented with adaptive learning techniques to enhance generalization in few-shot classification settings.

The proposed SimCLR-MAML model aims to address these limitations by combining self-supervised learning (SimCLR) with model-agnostic meta-learning (MAML). This hybrid

approach enables the model to learn robust feature representations from unlabeled data while efficiently adapting to new pneumonia cases using a few labeled examples. By leveraging the strengths of both SimCLR and MAML, the proposed model seeks to reduce dependency on large labeled datasets, improve generalization, and enhance classification performance in real-world pneumonia detection tasks. The integration of contrastive feature learning with meta-learning optimization represents a significant advancement in few-shot medical image classification, paving the way for scalable and efficient AI-driven diagnostic tools in healthcare.

6. Conclusion and Future Work

This study proposed a hybrid model, SimCLR-MAML, integrating self-supervised contrastive learning (SimCLR) with model-agnostic meta-learning (MAML) to enhance few-shot pneumonia detection from chest X-ray images. The experimental results demonstrated that the proposed model effectively leverages SimCLR's feature extraction capabilities and MAML's adaptability to new tasks with limited labeled data. The SimCLR-MAML model achieved an accuracy of 93.21%, a precision of 93.42%, a recall of 93.21%, and an F1-score of 93.21%, showing competitive performance compared to standalone SimCLR (93.69% accuracy) and significantly outperforming MAML (76.00% accuracy). The results validate the model's robustness in classifying pneumonia cases while maintaining high precision and recall, crucial for reducing false positives and false negatives in medical diagnosis. Although SimCLR alone achieved slightly higher accuracy and precision, the SimCLR-MAML hybrid model demonstrated more balanced performance across all evaluation metrics, making it a more reliable option for real-world medical applications where both precision and recall are critical. The ability of the hybrid model to generalize effectively with limited data highlights its potential applicability in resource-constrained healthcare settings, where access to large labeled datasets is often limited.

Despite the promising results, there are several areas for future research. First, exploring alternative self-supervised learning techniques, such as BYOL (Bootstrap Your Own Latent) or MoCo (Momentum Contrast), could potentially enhance feature representations further. Additionally, different meta-learning algorithms, such as ProtoNets or Reptile, could be investigated to improve adaptability to new, unseen tasks. Another important research direction is improving model interpretability. Although the hybrid model achieves high accuracy, explainability remains a major challenge in deep learning-based medical diagnostics. Incorporating explainable AI (XAI) techniques such as Grad-CAM (Gradient-weighted Class Activation Mapping) or SHAP (SHapley Additive Explanations) would enhance model transparency, making it more trustworthy for clinicians. Furthermore, domain adaptation and cross-hospital validation should be performed to ensure model robustness across diverse datasets. The current study used a single dataset, and testing the model on external datasets will confirm its ability to generalize across different imaging protocols and patient populations. Federated learning approaches could also be explored to enable privacy-preserving collaborative training across multiple hospitals, addressing data-sharing concerns in medical applications. Lastly, the model should be optimized for real-time deployment in clinical workflows. Implementing edge AI or cloud-based solutions could enable rapid inference in low-resource settings, allowing for real-time pneumonia screening in remote healthcare facilities.

DECLARATIONS

Ethical approval - Not Applicable

Consent to participate -Not Applicable

Consent to publish -Not Applicable

Competing interests: The authors declare no competing interests.

Clinical trial number: Not applicable.

Acknowledgements
Not Applicable.

Author Contribution Declaration

Atukunda Doreen (AD): Writing the manuscript, Data Interpretation and Analysis, Methodology, Implementation, Manuscript Writing, and Review, read and approved the final version of the manuscript for publication. **Waweru Mwangi (WM):** Methodology, Manuscript Review, supervision, and approval of the final version of the manuscript for publication. **Petronilla Muriithi (PM):** Methodology, Manuscript Review, supervision, and approval of the final version of the manuscript for publication. All authors have read and approved the final version of the manuscript.

Funding
We acknowledge the financial support from DAAD.

Data availability

The data utilized in this study is sourced from the Mendeley Data which is publicly accessible at the following URL: https://data.mendeley.com/datasets/rscbjbr9sj/2. doi 10.17632/rscbjbr9sj.2. Accessed 15-October-2024.

Code availability: Not applicable.

References

[1] P. Radočaj, D. Radočaj, and G. Martinović, "Pediatric Pneumonia Recognition Using an Improved DenseNet201 Model with Multi-Scale Convolutions and Mish Activation Function," *Algorithms*, vol. 18, no. 2, pp. 1–17, 2025, doi: 10.3390/a18020098.

[2] R. K. Id, R. Das Id, Z. Woo, G. Id, G. H. Id, and R. S. Id, "Pneumonia detection in chest X-ray images using an ensemble of deep learning models," *PLoS One*, 2021, doi: 10.1371/journal.pone.0256630.

[3] P. Rajpurkar *et al.*, "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," pp. 3–9, 2017, [Online]. Available: http://arxiv.org/abs/1711.05225

[4] G. Litjens *et al.*, "A Survey on Deep Learning in Medical Image Analysis," *Med. Image Anal.*, 2017, doi: 10.1016/j.media.2017.07.005.

[5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2015, doi: 10.1109/CVPR.2016.90.

[6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," *Proc. - 30th IEEE Conf. Comput. Vis. Pattern Recognition, CVPR 2017*, vol. 2017-Janua, pp. 2261–2269, 2017, doi: 10.1109/CVPR.2017.243.

[7] C. Shorten and T. M. Khoshgoftaar, "A survey on Image Data Augmentation for Deep Learning," *J. Big Data*, vol. 6, no. 1, 2019, doi: 10.1186/s40537-019-0197-0.

[8] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a Few Examples: A Survey on Few-shot Learning," *ACM Comput. Surv.*, vol. 53, no. 3, pp. 1–34, 2020, doi: 10.1145/3386252.

[9] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," *34th Int. Conf. Mach. Learn. ICML 2017*, vol. 3, pp. 1856–1868, 2017.

[10] A. Antoniou, A. Storkey, and H. Edwards, "How to train your MAML," *7th Int. Conf. Learn. Represent. ICLR 2019*, pp. 1–11, 2019.

[11] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," *37th Int. Conf. Mach. Learn. ICML 2020*, vol. PartF168147-3, no. Figure 1, pp. 1575–1585, 2020.

[12] A. Kolesnikov *et al.*, "Big Transfer (BiT): General Visual Representation Learning," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 12350 LNCS, pp. 491–507, 2020, doi: 10.1007/978-3-030-58558-7_29.

[13] M. Raghu, C. Zhang, J. Kleinberg, and S. Bengio, "Transfusion: Understanding transfer learning for medical imaging," *Adv. Neural Inf. Process. Syst.*, vol. 32, no. NeurIPS, 2019.

[14] N. Krishnamoorthy, K. Nirmaladevi, T. Kumaravel, K. S. S. Nithish, S. Sarathkumar, and M. Sarveshwaran, "Diagnosis of Pneumonia Using Deep Learning Techniques," *2022 Second Int. Conf. Adv. Electr. Comput. Commun. Sustain. Technol.*, pp. 1–5, 2022, doi: 10.1109/ICAECT54875.2022.9807954.

[15] P. Szepesi, "Detection of pneumonia using convolutional neural networks and deep learning," vol. 42, pp. 1012–1022, 2022, doi: 10.1016/j.bbe.2022.08.001.

[16] T. R. Ornob, G. Roy, and E. Hassan, "CovidExpert: A Triplet Siamese Neural Network framework for the detection of COVID-19," *Informatics Med. Unlocked*, vol. 37, 2023.

[17] M. Shorfuzzaman and M. S. Hossain, "MetaCOVID: A Siamese neural network framework with contrastive loss for n-shot diagnosis of COVID-19 patients," *Pattern Recognit.*, vol. 113, p. 107700, 2021, doi: 10.1016/j.patcog.2020.107700.

[18] A. Kareem, H. Liu, and P. Sant, "Review on Pneumonia Image Detection : A Machine

Learning Approach," *Human-Centric Intell. Syst.*, vol. 2, no. 1, pp. 31–43, 2022, doi: 10.1007/s44230-022-00002-2.
[19] R. K. Singh, R. Pandey, and R. N. Babu, "COVIDScreen: explainable deep learning framework for differential diagnosis of COVID-19 using chest X-rays," *Neural Comput. Appl.*, vol. 33, no. 14, pp. 8871–8892, 2021, doi: 10.1007/s00521-020-05636-6.
[20] R. Vyas and D. R. Khadatkar, "Ensemble of Deep Learning Architectures with Machine Learning for Pneumonia Classification Using Chest X-rays," *J. Imaging Informatics Med.*, 2024, doi: 10.1007/s10278-024-01201-y.
[21] S. Sharma and K. Guleria, "A Deep Learning based model for the Detection of Pneumonia from Chest X-Ray Images using VGG-16 and Neural Networks," *Procedia Comput. Sci.*, pp. 357–366, 2023, doi: 10.1016/j.procs.2023.01.018.
[22] M. Mujahid, F. Rustam, R. Álvarez, J. Luis Vidal Mazón, I. de la T. Díez, and I. Ashraf, "Pneumonia Classification from X-ray Images with Inception-V3 and Convolutional Neural Network," *Diagnostics*, vol. 12, no. 5, pp. 1–16, 2022, doi: 10.3390/diagnostics12051280.
[23] Nillmani *et al.*, "Four Types of Multiclass Frameworks for Pneumonia Classification and Its Validation in X-ray Scans Using Seven Types of Deep Learning Artificial Intelligence Models," *Diagnostics*, vol. 12, no. 3, 2022, doi: 10.3390/diagnostics12030652.
[24] M. R. Hasan, S. M. Azmat Ullah, and S. M. Rabiul Islam, "Recent advancement of deep learning techniques for pneumonia prediction from chest X-ray image," *Med. Reports*, vol. 7, no. June, p. 100106, 2024, doi: 10.1016/j.hmedic.2024.100106.
[25] R. Yi, L. Tang, Y. Tian, J. Liu, and Z. Wu, "Identification and classification of pneumonia disease using a deep learning-based intelligent computational framework," *Neural Comput. Appl.*, vol. 35, no. 20, pp. 14473–14486, 2023, doi: 10.1007/s00521-021-06102-7.
[26] N. W. Asnake, A. O. Salau, and A. M. Ayalew, "X-ray image-based pneumonia detection and classification using deep learning," *Multimed. Tools Appl.*, 2024, doi: 10.1007/s11042-023-17965-4.
[27] A. Kukker, R. Sharma, G. Pandey, and M. Faseehuddin, "Fuzzy lattices assisted EJAYA Q-learning for automated pulmonary diseases classification," *Biomed. Phys. Eng. Express*, 2024, doi: 10.1088/2057-1976/ad72f8.
[28] A. Kukker and R. Sharma, "JAYA-Optimized Fuzzy Reinforcement Learning Classifier for COVID-19," *IETE J. Res.*, 2022, doi: 10.1080/03772063.2022.2027287.
[29] A. Kukker and R. Sharma, "Modified Fuzzy Q Learning Based Classifier for Pneumonia and Tuberculosis," *IRBM*, 2021, doi: 10.1016/j.irbm.2020.10.005.
[30] S. Shastri, I. Kansal, and S. et al. Kumar, "CheXImageNet: a novel architecture for accurate classification of Covid-19 with chest x-ray digital images using deep convolutional neural networks," *Heal. Technol.*, 2022, doi: 10.1007/s12553-021-00630-x.
[31] S. Kumar *et al.*, "LiteCovidNet: A lightweight deep neural network model for detection of COVID-19 using X-ray images," *Int. J. Imaging Syst. Technol.*, 2022, doi: 10.1002/ima.22770.
[32] A. Hosna, E. Merry, J. Gyalmo, Z. Alom, Z. Aung, and M. A. Azim, "Transfer learning: a friendly introduction," *J. Big Data*, vol. 9, no. 1, 2022, doi: 10.1186/s40537-022-00652-w.
[33] A. Rafiei, R. Moore, S. Jahromi, F. Hajati, and R. Kamaleswaran, "Meta-learning in Healthcare: A Survey," *SN Comput. Sci.*, vol. 5, no. 6, pp. 1–27, 2024, doi: 10.1007/s42979-024-03166-9.