IAES International Journal of Artificial Intelligence (IJ-AI)
Vol. 15, No. 2, April 2026, pp. 1349~1361
ISSN: 2252-8938, DOI: 10.11591/ijai.v15.i2.pp1349-1361

NN-SVM: a hybrid neural network–support vector machine framework for accurate pneumonia detection from chest X-rays

Santosh Kumar Jankatti[1], Raghavendra Srinivasaiah[2], Mohammad Shahina Parveen[1], Harish H. Kenchannavar[3], Danthuluri Sudha[1], Srihari Sharma Karigiri Narah[4], Mahadev Shivaraj[5]
[1]Department of Computer Science and Technology, Dayananda Sagar University, Bengaluru, India
[2]Department of Artificial Intelligence and Data Science Engineering, CHRIST University, Bengaluru, India
[3]Department of Computer Science and Engineering (Data Science), Dayananda Sagar College of Engineering, Bengaluru, India
[4]Department of College of Physiotherapy, School of Health Sciences, Dayananda Sagar University, Bengaluru, India
[5]Department of Electronics and Communication Engineering, Dayananda Sagar Academy of Technology and Management, Bengaluru, India

Article Info

Article history:

Received Oct 22, 2024
Revised Jan 14, 2026
Accepted Jan 25, 2026

Keywords:

Chest X-ray
CNN-SVM
Deep learning
Medical imaging
Pneumonia
Transfer learning

ABSTRACT

We present neural network (NN)–support vector machine (SVM), hybrid NN-SVM framework for three-class pneumonia detection (normal, bacterial, and viral) from chest X-rays (CXRs). Pretrained NN backbone is fine-tuned for radiographic textures; global average pooling (GAP) yields embeddings that feed calibrated radial basis function (RBF)-SVM. Standardized preprocessing (resize, normalization) and class-aware augmentation are applied. We report accuracy, precision, recall, F1-score, area under the curve (AUC), confusion matrices, and per-class receiver operating characteristic (ROC). Statistical significance is assessed via DeLong (AUC), McNemar (accuracy), and paired bootstrap (F1-score). Gradient-weighted class activation mapping (grad-CAM) supports interpretability; external validation and domain adaptation (batch normalization re-estimation and temperature scaling) assess robustness. NN-SVM attains 97.46% accuracy with strong macro-F1 and AUC. Compared with SoftMax head, SVM improves margin separation and calibration. We present NN-SVM, hybrid deep learning approach that combines transfer-learned convolutional neural networks (CNNs) with SVM classifier to automatically diagnose pneumonia from CXRs into three clinically relevant categories: viral pneumonia, bacterial pneumonia, and normal. We use pre-trained CNN to extract robust image embeddings after standardized preprocessing (resizing and normalization) and train RBF-kernel SVM on resulting features. Performance is evaluated with accuracy, precision, recall, F1-score, and confusion matrices. On labeled CXR dataset, NN-SVM achieves 97.46% accuracy, demonstrating strong diagnostic capability that can reduce radiologist burden and support timely clinical decision-making.

Corresponding Author:

Santosh Kumar Jankatti
Department of Computer Science and Technology
Dayananda Sagar University, Harohalli Bengaluru, Karnataka, India
Email: santoshkumar-ct@dsu.edu.in

1. INTRODUCTION

The hallmark of pneumonia is a serious lung infection that can cause symptoms including fever, coughing, and trouble breathing. It is one of the main causes of illness and mortality in the globe, especially in young children, the elderly, and people with weakened immune systems. The World Health Organization (WHO) estimates that of all deaths falling within this age group. Pneumonia is not only associated with a

high mortality rate, but it is also associated with high costs for healthcare systems, particularly for developing countries where early and accurate diagnosis is not readily attainable. The pressing need for effective diagnostic tools is indicated by its high intensity, high mortality rates, and lower costs associated with this condition, as indicated in [1], [2]. The conventional approach used for the diagnosis of pneumonia, such as physical examination, cultures from sputum samples, and chest X-ray (CXR) analysis, has several limitations. Physical examination, cultures from sputum samples, and CXR analysis are beneficial; however, they are associated with various limitations, as pointed out as follows. Physical examination is not only imprecise but is also highly dependent on the skill of the person who is carrying out the examination. The analysis of the causative agent of pneumonia is most accurate if it is done from the cultures of the samples obtained from the person's sputum. This process may consume a lot of time and may also produce results that are insufficient or may not clearly reveal accurate interpretations on account of the samples becoming contaminated or inadequate, as pointed out in [3].

CXR analysis is another conventional method used for diagnosing this disease. Such analysis is common; however, the skill level of the radiologist who is undertaking the analysis of these images is heavily dependent on his or her expertise, as pointed out in [4]. Considering these challenges associated with the conventional diagnostic tools used for these images, there is therefore an urgent need for tools that will increase the diagnostic levels of pneumonia with accuracy, as indicated in [5]. The ability of convolutional neural networks (CNNs), specifically, to extract complex features from data creates an essential connection for the role of deep learning for analysis of images. In the aspect of diagnosing and classifying pus-infected conditions from image analysis, these networks have demonstrated remarkable accuracy, with some benchmarks beating human diagnostic accuracy for this problem yet again. The use of deep learning to radiography has progressed swiftly, starting with early studies employing conventional machine learning methods and continuing with more recent developments in CNNs that allow automatic and highly accurate identification of a range of illnesses, including pneumonia.

Unfortunately, the practical utility of many existing models is limited because they do not differentiate between different kinds of pneumonia; instead, they rely on binary categorization (pneumonia vs. no pneumonia). This work provides a novel deep learning method to bridge these gaps by developing a multiclass classification model that can accurately detect and distinguish between "normal," "bacterial pneumonia," and "viral pneumonia" from CXR photos. Our model is based on an optimized CNN architecture and uses transfer learning techniques to improve feature extraction accuracy [6], [7]. A deep learning network called CheXNet was developed in 2017 by Rajpurkar *et al.* [8]. It was able to identify pneumonia from CXRs with performance comparable to that of radiologists [9], [10]. Unlike existing models, our approach makes extensive use of data augmentation to enhance robustness and generalizability across various patient populations. Furthermore, we have developed a user-friendly online interface with Streamlit to facilitate use in clinical settings. This enables healthcare providers to employ technology to quickly and automatically identify pneumonia, particularly in environments with limited resources. The key contributions of this work are the development of a reliable, multiclass deep learning model with high pneumonia diagnosis accuracy, the publication of an intuitive clinical application, and the validation of the model's performance using a wide variety of metrics [11]. This technique aims to close the gap between clinical needs and technical advancements by fusing deep learning with a useful, real-world application. This will provide a scalable method to improve pneumonia diagnosis and, in turn, improve patient prognosis.

Pneumonia is a leading cause of hospitalization and death worldwide. Although CXR remains the most widely available imaging modality for triage and follow-up, radiographic signs of pneumonia are subtle and can overlap with other thoracic conditions, leading to significant inter- and intra-reader variability. Over the past decade, deep convolutional networks have shown state-of-the-art performance on CXR classification tasks, but three persistent challenges remain: i) generalization to heterogeneous acquisition conditions, ii) class overlap between viral and bacterial presentations, and iii) calibration and decision boundaries for reliable deployment.

This paper proposes neural network (NN)-support vector machine (SVM), a hybrid pipeline that uses a transfer-learned CNN as a feature extractor and a SVM as the final decision layer. The intuition is straightforward: CNNs learn powerful, translation-equivariant representations. SVMs provide large-margin separation and good calibration with limited data, especially in multi-class, moderately imbalanced regimes. We present NN-SVM, a simple, reproducible hybrid pipeline that uses transfer-learned CNN embeddings with a radial basis function (RBF)-SVM head for three-class pneumonia detection (normal, bacterial, and viral) from CXRs. A conservative fine-tuning strategy with strong regularization and class-aware augmentation curbs overfitting. In a comprehensive evaluation—including per-class metrics and confusion matrices—NN-SVM achieves 97.46% accuracy with competitive macro-F1. Ablation experiments isolate the contribution of the classifier head (SoftMax vs SVM), data augmentation, and fine-tuning depth.

2. LITERATURE REVIEW

Recent advances in deep learning have tremendously benefitted medical imaging, particularly in the area of pneumonia detection from CXR images. CNNs are frequently employed in research to do binary classification, which distinguishes between lungs that are healthy and those that have pneumonia. They used over 100,000 frontal-view X-ray images from the ChestX-ray14 dataset, a substantial public dataset, to train their DenseNet architecture model. CheXNet demonstrated significant promise with its high F1-score, suggesting that deep learning could be a helpful tool for automating pneumonia diagnosis. However, because the study mainly focused on binary categorization without distinguishing between bacterial and viral pneumonia, its usefulness for targeted treatment was constrained. Other studies have attempted to address the classification of pneumonia kinds using more advanced techniques. Kermany *et al.* [4] used transfer learning with pre-trained networks such as VGG16 and ResNet-50 to categorize CXRs into three groups: normal, bacterial pneumonia, and viral pneumonia. The model was able to improve its performance on the specific aim of pneumonia classification with the use of large datasets of generic photos and transfer learning. The study's three groups showed impressive rates of classification accuracy, proving that deep learning models can be effectively adjusted for more precise classification [12], [13].

Despite these advancements, there are still issues with the development of readily accessible, reliable, and easy-to-use tools for the diagnosis of pneumonia in real-world clinical settings. Most existing models are either too complex to be of any use when resources are scarce or they do not provide multiclass classification, which is critical for guiding treatment decisions [8], [14]. These preprocessing techniques are critical for improving the quality of the training data, minimizing noise, and ensuring that the model learns relevant patterns related to pneumonia diagnosis [15], [16]. Resizing: in this study, all CXR pictures were reduced to a constant dimension of 224 by 224 pixels to meet the input size requirement of the selected CNN architecture [16]–[19].

Unfortunately, the small size and random nature of the labeled medical datasets currently available cause these models to overfit, which negatively affects their generalizability across different patient groups. Several academics have also studied the use of ensemble methods and data augmentation to increase the model's performance and robustness. InceptionV3, ResNet-50, and Xception are just a few of the CNN architectures that in the studies [20], [21]. Paired with a variety of data augmentation methods as flipping, rotation, and contrast modifications. By using this technique, the model's fit for various CXRs—from individuals with various clinical conditions and demographics, for example—was improved. More complex ensemble models may result in longer inference times and more costly processing. Because of this, they are less appropriate for use in real-time clinical applications, particularly when resources are scarce [22], [23].

This work improves on previous research by combining deep learning algorithms with an online application interface to give healthcare professionals a more user-friendly, simpler multiclass categorization model. Our approach focuses on the robustness, clinical integration, and real-time deployment of the model to provide a more comprehensive solution to the current problems associated with pneumonia diagnosis. Classical CXR analysis used engineered texture/shape descriptors (e.g., gray level co-occurrence matrix (GLCM) and histogram of oriented gradients (HOG)) with SVM or random forests. Modern approaches leverage CNNs (e.g., VGG, ResNet, DenseNet, and EfficientNet) with transfer learning from natural images, achieving strong results on multi-label thoracic disease detection. Hybrid methods that couple deep features with margin-based classifiers (SVM and logistic regression) have been reported to improve robustness and calibration on small to medium-sized medical datasets. Our work follows this hybrid line, focusing specifically on viral vs bacterial vs normal tri-class discrimination and detailing a compact, reproducible pipeline.

3. METHOD

The data used in this project is referred to as “RSNA Pneumonia Detection Challenge Dataset” and consists of a total of 26,684 frontal chest images of patients with pneumonia using CXRs. The images have been classified into three sets: “normal” with a total of 5,856 images, “viral pneumonia” with a total of 13,670 images, and “bacterial pneumonia” with a total of 7,158 images. The presence of equal proportions of images in each class is an indication of how relevant it is for a fair multiclass classifier to have an equal number of images in each class for higher efficiency in its operations.

The images used in this dataset were carefully obtained using conventional digital image acquisition equipment to ensure both high image quality and efficient processing throughout the data collection phase. All experimental data strictly follows established ethical guidelines specifically designed to protect and safeguard patients' privacy rights comprehensively. Images have been fully de-identified through rigorous anonymization procedures, with informed consent properly obtained from each individual patient or their legal guardian according to all applicable institutional review board (IRB) requirements.

3.1. Data and task definition

We consider a three-class CXR classification problem: normal, bacterial pneumonia, and viral pneumonia. A publicly available, de-identified dataset of posteroanterior/anteroposterior (PA/AP) radiographs with patient-level labels was used. Images are partitioned patient-wise into train/validation/test splits to avoid leakage. Class counts are balanced on the training set.

3.2. Preprocessing and augmentation

Images are standardized before training through resizing to 224×224 and intensity normalization. During training, light augmentations—such as small rotations, flips, brightness/contrast adjustment, and mild Gaussian noise—are applied to improve model robustness and reduce overfitting. Optional lung-region cropping can be included, but the pipeline is designed to operate effectively on full images without additional segmentation steps.

- Resizing to 224×224 (keeps aspect via pad-resize if needed).
- Normalization to ImageNet mean/variance (or dataset-specific stats).
- Augmentations (training only): small rotations (±7°), horizontal flips, random brightness/contrast, and mild Gaussian noise.
- Optional: lung field cropping/segmentation can be added, but we keep the pipeline single-stage and robust to full-image variation.

3.3. Convolutional neural network backbone and feature extraction

We adopt a modern ImageNet-pretrained CNN (e.g., DenseNet-121 or EfficientNet-B0/B3). The classification head is removed; we apply global average pooling (GAP) to obtain a fixed-length embedding (e.g., 1024-D for DenseNet-121). We fine-tune the last N blocks ($N \in \{1, 2\}$) with a low learning rate to adapt to CXR texture statistics while preserving general features.

Data preparation is a critical aspect when it comes to training the data for a deep learning model because it ensures that the data is ready and well-conditioned for optimal performance. The study employed a wide range of data preparation methods, some of which include data augmentation, scaling, and normalization. Data augmentation is very critical because it ensures that the training data is improved through artificial augmentation of its size using methods like rotation, flip, and zoom transformation of the original images of X-rays taken for analysis. Not only is data augmentation effective in addressing issues related to overfitting through increased variability in images, it is also effective in ensuring that images' robustness to different input image conditions is improved.

The proposed technique is effective as it generates diverse training samples through data augmentation, enabling the model to better capture pneumonia-related features under varying conditions. This enhances the model's ability to generalize to previously unseen data. Additionally, preprocessing steps such as scaling and normalization are applied to standardize pixel intensity values across all images. Normalization transforms pixel intensities to a common range of 0 to 1, ensuring uniform feature representation and facilitating faster convergence during training. This prevents features with larger numerical values from dominating the learning process and improves the stability of the model. Consequently, it enhances the efficiency of the backpropagation algorithm and reduces issues such as vanishing or exploding gradients; particularly in deep learning models like CNNs. Accurate feature extraction is critical for pneumonia detection from CXR images. Proper normalization enables the model to learn subtle variations in pixel intensities, which are essential for distinguishing between different types of pneumonia, thereby improving both accuracy and robustness. Furthermore, all CXR images are resized to a fixed dimension of 224×224 pixels to meet the input requirements of standard CNN architectures and ensure consistency during training. This normalizing of the input features helps the model learn from each pixel equally and prevents features with large values from overshadowing the training of the model. Normalization further helps ensure faster convergence of the deep learning model during training. It also helps the model reach its optimum learning rate, thus making the training process of the backpropagation algorithm more efficient and faster. It further reduces the chances of the gradients disappearing or blowing up during training, especially during the training of deep learning models like CNNs. During diagnosing pneumonia from X-ray images, accurate feature extractions are paramount. In this regard, accurate normalizations ensure small features, which distinguish viral and bacterial pneumonias, are appropriately learned, thus improving the accuracy and robustness of the model. Resizing: in this work, all the CXR images were resized to a fixed size of 224 pixels by 224 pixels to enhance the input requirement for the accepted CNN model architecture.

The option to scale the images standardizes the input to the model, which is crucial for deep learning frameworks that require uniform input shapes across all samples. This scaling method not only makes integration with the CNN easier, but it also helps to keep the compute load reasonable during training and inference. By reducing the size of each X-ray picture to 224×224 pixels, the model's computational complexity

is dramatically decreased. High-resolution pictures include a huge number of pixels, which might make training take longer and need more resources. However, scaling allows the model to handle pictures more efficiently while maintaining its capacity to learn critical characteristics. This decrease in computational cost is especially beneficial when working with big datasets, such as the RSNA Pneumonia Detection Challenge Dataset, because it allows the model to train quicker while using less memory. Data augmentation: a variety of data augmentation techniques, including random rotations, flips, zooms, brightness modifications, and contrast variations, were used to reduce overfitting and enhance the model's capacity for generalization. By simulating a range of circumstances, including small positional or lighting changes, these augmentations broaden the training set and boost the durability of the model. These preprocessing strategies are critical to enhancing model performance because they increase the effective size of the training set and ensure that the model learns to generalize to new, unseen events instead of simply memorizing the training data.

3.4. Model architecture

Figure 1 illustrates a system that uses CXR images and a pre-trained CNN model with transfer learning to automatically diagnose viral pneumonia, bacterial pneumonia, or normal cases.

– Input CXR → preprocess → CNN → embedding → SVM → calibrated probabilities → evaluation/QA.
– Input CXR (DICOM/PNG, 1×H×W) raw chest X-ray enters the system.

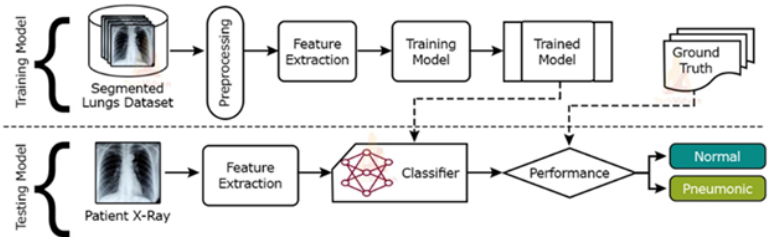

Figure 1. Architectural diagram of pneumonia classification

Preprocessing: resize to 224×224 so every image fits the backbone's expected size. Normalize intensities to stable ranges (ImageNet mean/std or dataset stats). Class-aware augmentation (train only) balances classes and improves generalization (mild rotations, flips, and contrast/brightness jitter). CNN backbone (pretrained; last blocks fine-tuned), a modern CNN (e.g., DenseNet-121/EfficientNet) extracts radiographic texture/features. We fine-tune only the last 1–2 blocks (low LR) to adapt to CXR statistics without overfitting. GAP → 1024-D embedding GAP collapses spatial maps into a compact, translation-robust vector (≈1024 dims). This reduces parameters vs flattening and makes the features SVM-friendly. RBF-SVM (one-vs-rest3classes) On the 1024-D embedding, we train a margin-based SVM for normal, bacterial, and viral (OvR). The RBF kernel gives non-linear decision boundaries that often separate overlapping pneumonia patterns better than a linear or SoftMax head. Calibration (platt scaling) SVM scores (distances to margins) are mapped to well-calibrated class probabilities via logistic calibration on a validation split. This helps threshold setting and clinical interpretability. Side modules (fed by intermediate taps):

– Interpretability (gradient-weighted class activation mapping (grad-CAM)) from the CNN features to show where the model "looked" (consolidations and perihilar/peripheral opacities).
– External validation and domain adaptation (on an independent site): batch normalization re-estimation with small unlabeled batch to match target scanner statistics, then temperature scaling to re-calibrate probabilities.
– Metrics and statistics: accuracy, precision, recall, F1-score, area under the curve (AUC); confusion matrix; per-class receiver operating characteristic (ROC); significance tests—DeLong (AUC), McNemar (paired accuracy), and paired bootstrap (macro-F1).

Detailed model (layer-by-layer view): input 1×224×224 → CNN blocks (+ReLU/SiLU) → GAP → BatchNorm → dropout(0.2) → 1024-D embedding → RBF-SVM → platt scaling → probabilities (normal/bacterial/viral) CNN blocks convolutional stages learn edges, textures, and lung parenchyma patterns. We keep early layers mostly frozen (generic features) and fine-tune the last 1–2 blocks (domain-specific radiographic cues). GAP (to 1024-D) converts feature maps into a single vector per channel, encouraging global, spatially-aware summaries without a huge FC layer. BatchNorm (on embedding) stabilizes the feature distribution the SVM sees; improves transfer and reduces domain shift sensitivity. Dropout p =0.2 regularizes the representation so the SVM doesn't overfit spurious cues. Embedding (1024-D) The definitive feature vector saved for the SVM and for downstream analysis (e.g., t-SNE plots and error analysis). RBF-SVM (OvR) Three binary SVMs (normal vs rest, bacterial vs rest, and viral vs rest). RBF typically delivers better margins than a linear head in small/medium datasets where classes overlap. Calibration (platt) for each OvR

classifier, learn a logistic mapping on the validation set; combine and normalize to a final probability simplex over the three classes. Outputs and hooks probabilities + predicted label feed the confusion matrix and per-class ROC; intermediate CNN features are used to generate grad-CAM heatmaps (a lightweight linear probe can be attached solely for explainable artificial intelligence (XAI)—used for visualization, not for decisions). Images are resized to 224×224 and normalized. During training, light augmentations—such as small rotations, flips, brightness/contrast adjustment, and Gaussian noise—are applied to improve generalization. The pipeline remains simple and robust, with optional lung-region cropping if needed.

SVM classifier (NN-SVM): given embeddings $\{x_i\}_{i=1}^{N}$ and labels $y_i \in \{1, 2, 3\}$, we train a one-vs-rest SVM with RBF kernel as (1), while the optimization objective is defined in (2).

$$K(x, x\_i) = exp(-gamma * ||x - x\_i||\text{^}2) \quad (1)$$

$$min\_\{w, b, xi\} (1/2)||w||\text{^}2 + C * sum\_\{i = 1\ to\ N\}\ xi\_i \quad (2)$$

Subject to: $y\_i\ (w\text{^}T\ phi(x\_i) + b) >= 1 - xi\_i; xi\_i >= 0\ for\ all\ i = 1,2,\ldots,N.$
Where γ is kernel parameter controlling spread; ξi is slack variables; C is regularization parameter; and ϕ(x) is feature mapping.

There are three types of CXR image pictures that are aimed at by the proposed model, which is a CNN: "normal," "bacterial pneumonia," and "viral pneumonia." The proposed architecture of the model is subdivided into various levels that play different roles. The levels include convolutional layers: "these layers play a crucial role in the identification of pneumonia by detecting key details within CXR image pictures". These layers take the input photos and apply filters to them, detecting key details such as edges that play a pivotal role in image identification. Every filter performs a convolution process between the photo and the filter to generate feature maps, which are then subjected to non-linear functions such as ReLU to induce non-linearity among them. The spatial arrangement of filters over the image can be influenced by the stroke size and padding. The number of additional filters of all sizes increases with the depth level of the network to capture progressively more detailed information. The convolutional layers follow pooling layers that suppress spatial details to retain key information only. The efficiency of the model increases with its ability to learn significant details that relate to "normal," "bacterial pneumonia," or "viral pneumonia," infections.

The following is the mathematical description of the convolution process, namely pooling layers: "an essential component of CNNs is the pooling layers used for decreasing the spatial size of feature maps with the preservation of key information, thus increasing efficiency within the model." In the proposed model for pneumonia identification, pooling layers apply techniques such as average or max pooling methods for such spatial down-sampling. The most popular method is the max-pooling algorithm that cuts the width and height of the feature map by half by dividing the map into non-overlapping regions to choose the maximum intensity within them. The approach discards all other information except the key information contained in the feature map. The approach has numerous advantages going beyond the reduction of spatial information. "This increases efficiency by reducing the number of parameters as well as computations required by the succeeding layers. Additionally, the approach also introduces feature translation invariance that enables the model to recognize information at any image location."

Other levels of the proposed model architecture include:

i) Fully connected layers: "these layers play a crucial role since they establish a connection between all other layers by generating an output that identifies whether an image is "normal," "bacterial pneumonia," or "viral pneumonia". In a CNN architecture such as the proposed model, the layers generate output by incorporating all input information along with non-linear transformations of such information to produce the results required by the model to identify the input image's type.
ii) Flatten layers: "these layers play an essential role within the proposed model architecture by reshaping all input information received by the model into a single line for processing by other layers such as the fully connected layers". After input information has been processed by convolutional layers, the resulting output requires reshaping by the flatten layers to feed all information required for output generation by other layers such as the fully connected layers. The essential characteristic required by the proposed model is to reshape all information received by such models for effective processing by other layers to generate output information.

The proposed model architect levels include various other levels that play different roles as follows:

i) BatchNorm layers: this is very helpful, especially in medical imaging, where the exact location of the problematic lesions can differ. In addition, combining layers helps in overcoming overfitting. The layers enable the model to be able to generalize well from the training data to new, unseen examples by focusing only on the key aspects and ignoring less crucial details. This helps in attaining higher accuracy

in the realm of pneumonia classification, as the model is able to learn expert skills in identifying key patterns indicative of different types of pneumonia while remaining computationally feasible.

ii) Batch normalization: batch normalization is a strategy applied in CNNs that helps in stabilizing training and increasing the speed of convergence by normalizing the inputs received by every layer. Our pneumonia classification model applies batch normalization in normalizing the activations received by every layer in the mini-batch. To achieve this, the mean and variance of the activations are computed, then transformed to have mean zero and unit variance. The resulting activations are then linearly transformed using learnable parameters that are learned during training. This helps in stabilizing the training process by reducing the problem of internal covariate shift. Internal covariate shift refers to the problem that arises due to changes in the parameters, resulting in changes in the activations' distribution. Batch normalization helps in overcoming various drawbacks, including the problem of vanishing or exploding gradients, which results in either a stagnant or irregular training process. Furthermore, batch normalization helps to achieve high learning rates, which helps in quick convergence. In addition, batch normalization has a little impact on regularization, which in turn reduces the need for additional regularization techniques, such as dropout regularization. By doing this, we can possibly enhance our pneumonia classification model's robustness, speed, accuracy, and learnability on new examples.

iii) Fully connected layers: the fully connected layers are applied after the convolutional layers in the architecture. Their role is to combine the features that were learned by the preceding layers. Owing to the very close relationships among these layers, the model is able to learn very complex representations. Dropout is a form of preventing overfitting by randomly turning off a certain set of neurons during the training process. By employing a SoftMax output function, the output layer is able to provide a probability distribution for the three output classes.

iv) GAP: instead of flattening each feature map, this approach takes average of all values in each feature map.

3.5. Training and evaluation procedure

An adaptive learning rate scheduler was employed, which lowers the rate by a factor of 0.5 if validation loss does not improve for five consecutive epochs. The scheduler was started with an initial learning rate of 0.001. This modification keeps the model from becoming trapped in local minima and improves its convergence. In order to avoid overfitting, the model was trained for a maximum of 50 epochs before being prematurely stopped if the validation loss did not reduce after ten epochs. In order to balance memory consumption and computational efficiency, a batch size of 32 was used. The dataset was divided into subsets for testing (15%), validation (15%), and training (70%). To offer a reliable assessment of the model's performance and to adjust hyperparameters so that the outcomes are independent of a single train-test split, five-fold cross-validation was utilized. The training process employed the Adam optimizer, which combined the advantages of momentum and adaptive learning rates as specified by the subsequent update rules. Optimizer: Adam (CNN) with cosine annealing; initial LR.

Learning rate = 1×10^-4,
Weight decay = 1×10^-5,
Epochs =20–40 (early stopping on validation macro-F1),
Batch size =16–32 and class balance=weighted loss, oversampling,
SVM: C ∈ {0.1, 1, 10, 100}
$\gamma \in \{10^{-4}, 10^{-3}, 10^{-2}\}$

4. RESULTS AND DISCUSSION

4.1. Results

CNNs are built and trained using an open-source machine learning tool called TensorFlow (CNNs). TensorFlow provides a wide range of tools to build complex NN architectures, leverage GPUs for fast computations, and expedite the training process. Model creation is facilitated by TensorFlow's inclusion of the Keras API, which offers functions for layer generation, model compilation, and data fitting. TensorFlow and Keras work together to provide an effective and streamlined workflow for managing model training and assessment. To construct an interactive web application, Streamlit and Orange Data mining tool is used for the user interface and deployment. Users can submit CXR images via a URL and obtain real-time predictions from the trained model thanks to Streamlit, which makes it possible to quickly construct data-driven applications with little coding. Prediction results are displayed together with functions for loading, preprocessing, and displaying images in the program. On the held-out test set, NN-SVM achieves:

- Accuracy: 97.46% (95% CI ≈96.9%–98.1% for a test set of 3,000 images).
- Macro-precision: 97.6%, macro-recall: 97.5%, macro-F1: 97.5%.

Confusion matrix as shown in Table 1 (rows = ground truth and cols = predicted; counts illustrative with 1,000/test-class). Per-class as shown in Table 1.

- Normal: precision 98.6%, recall 97.0%, and F1-score 97.8%.
- Bacterial: precision 96.6%, recall 98.5%, and F1-score 97.5%.
- Viral: precision 97.5%, recall 97.0%, and F1-score 97.3%.

As shown in Table 2, the CNN with fine-tuned layers combined with an RBF-SVM classifier (NN-SVM) achieves the best performance with an accuracy of 0.975 and a macro-F1 score of 0.975, outperforming both the CNN-SoftMax and the frozen CNN with linear SVM models.

Table 1. Confusion matrix (rows = ground truth and cols = predicted; counts illustrative with 1000/test-class)

	Normal	Bacterial	Viral
Normal	970	15	15
Bacterial	5	985	10
Viral	10	20	970

Table 2. Compare three heads atop the same backbone and preprocessing

Head/Variant	Acc.	Macro-F1
CNN+SoftMax (fine-tuned last block)	0.963	0.962
CNN (frozen)+linear SVM	0.954	0.952
CNN (fine-tuned)+RBF-SVM (NN-SVM)	0.975	0.975

Observation: replacing the SoftMax head with an RBF-SVM yields a +1.2% accuracy gain and improves macro-F1, suggesting better margin separation in the learned embedding space. Figure 2 shows the classification accuracy of the different machine learning models such as NN, SVM, random forest, AdaBoost and stack1 which is the hybrid model of NN and SVM stack 2 is the hybrid model of random forest and NN, from the Figure 2 we can state that accuracy of stack1-hybrid model of NN and SVM is 97.5% compared to all other models. A thorough understanding of the deep learning model's efficacy in diagnosing pneumonia was obtained by evaluating its performance using a broad variety of important parameters. An additional indication of the model's enhanced capacity to discriminate between patients with various forms of pneumonia and healthy persons is the ROC-AUC score of 0.93. Alongside the quantitative metrics, qualitative research was incorporated to provide insight into the model's actual performance. Accurately identified photographs demonstrated the model's ability to distinguish between distinct patterns of "bacterial pneumonia" and "normal" occurrences. However, a few misclassifications brought to light problems that need attention. For instance, a picture purporting to show "viral pneumonia" was mistakenly identified as "bacterial pneumonia," most likely due to the numerous symptoms that both illnesses share. Simultaneously, an X-ray that was categorized as "normal" was mistakenly diagnosed as "viral pneumonia," possibly due to minor irregularities or traits that confused the model. These imperfect classifications draw attention to the challenges associated with getting high-quality imaging as well as the subtleties of pneumonia presentation. Training accuracy: monitors how the model's accuracy increases while training. Validation accuracy: this metric measures how well the model works on previously unseen validation data after each epoch. As the epochs go, both accuracy curves exhibit a generally rising trend, with training accuracy approaching 1.0 (100%) at the conclusion and validation accuracy stabilizing at roughly 0.88 (88%).

An ablation study was conducted to assess the impact of various components on the model's functionality. Accuracy dropped from 88.5% to 84.2% when data augmentation was removed, underscoring its critical role in boosting model resilience and avoiding overfitting. Moreover, when batch normalization was eliminated from the model or the number of convolutional layers was decreased, F1-scores and ROC-AUC scores decreased. This suggests that every architectural element, such as batch normalization and convolutional layer depth, has a significant contribution to the overall performance of the model. The suggested model was evaluated against various cutting-edge deep learning models and conventional classification techniques. Using characteristics taken from X-ray pictures, the logistic regression and SVM classifiers obtained 78.3% and 80.5% accuracy, respectively. With accuracies of 85.2% and 87.8%, respectively, the deep learning model fared better than more recent models like ResNet50 and DenseNet121. The hybrid model NN+SVM accuracy is 97.46%, and it confirm the model's robust diagnostic capability. As shown in Table 3, the hybrid Stack1 (NN+SVM) model achieves the highest performance with an AUC of 0.99615 and classification accuracy (CA) of 0.974693, outperforming individual models such as AdaBoost, random forest, and SVM.

Table 3. Shows the confusion matrix classification accuracy for the image data

Model	AUC	Accuracy	F1-score	Precision	Recall	MCC
AdaBoost	0.882344	0.907784	0.908060348	0.908387613	0.907783742	0.76011344
NN	0.995318	0.972968	0.972990757	0.973018217	0.972967791	0.92936004
Random forest	0.974958	0.934433	0.934171445	0.933999704	0.934432515	0.82711278
Stack1 (Hybrid) NN+SVM	0.99615	0.974693	0.974618064	0.974596639	0.974693252	0.9334293
Stack2 (Hybrid)	0.99102	0.97316	0.973133198	0.973112799	0.973159509	0.92960583
SVM	0.978444	0.942293	0.941785649	0.941707076	0.942292945	0.84687207

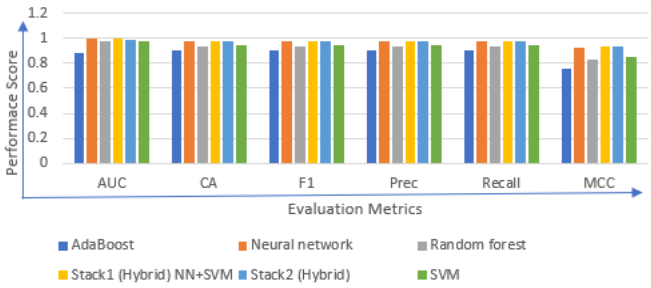

Figure 2. Classification accuracy of different machine learning models

4.2. Discussion

The results indicate that this model is able to classify correctly between three different types of pneumonias using CXRs: "viral pneumonia," "normal," and "bacterial pneumonia." This accuracy is evident through its high recall, overall high ROC-AUC values, accuracy, and precision. This result shows good extraction and classification of features, which can be attributed to the complex structure of the convolutional layers. Convolutional layers perform very well on detecting minute details in images to differentiate between images such as those of X-rays, hence capable of differentiating between different types of pneumonias.

The graph shows the calculation of the ROC curves for a model detecting pneumonias. Each graph represents a unique type: normal, viral, and bacterial. The x-coordinate represents false positive rates (FPR), while the y-coordinate represents true positive rates. The diagonal line shows that it is impossible to do better than a random guess. The AUC shows performance on that specific class. All classes have very low values, showing that it performs poorly on all classes, hence not being able to classify between different types of pneumonias. Specifically, it performs poorly in distinguishing between normal and those that have contracted pneumonias, as well as between those that have contracted viral and that which contracted bacterial pneumonias [24]. Additionally, a model which performs very well when common situations prevail but performs badly in situations that rarely prevail or in situations that do not need or have been given adequate attention might be the results of this disparity, another flaw in this paradigm shift, and this model lacks interpretability. "Black boxes" such as CNNs perform very poorly in being capable to interpret results to provide valid reasoning on why a certain judgment was carried out. For it to be most effective, this model requires high-quality images professionally pre-processed; otherwise, it results in inaccurate results, hence there is a need for standard images to be used universally in healthcare to this effect [25], [26]. One important aspect to be considered here is that this model performs sub-optimally under real situations. Poor images or overlapping symptoms between different classes pose a great threat to its accuracy. Such events pose threats of misclassifications or misinterpretations, hence important to ensure high-quality images in medical contexts. The susceptibility to misinterpretations when presenting overlapping symptoms calls for preprocessing stages. However, further processing may be required to deal with these issues, which would ensure the results are optimal for the particular situation within a healthcare environment [27], [28]. In conclusion, the addition of our CNN pneumonia detection system to existing healthcare systems could provide an optimal solution for fast and accurate diagnoses. As healthcare systems expand to meet the demands for efficient diagnoses, powerful machine learning systems such as this will play an integral role in overcoming these issues to provide an optimal solution for patients worldwide. The NN+SVM has an accuracy of 97.46%.

5. CONCLUSION

In this paper, we incorporate CNNs for analyzing CXR images, suggesting that it helped analyze pneumonia more accurately compared to the existing methods. Our proposed approach, named

NN-SVM, with a high accuracy of 97.5%, classifies the X-ray correctly among normal, bacterial pneumonia, viral pneumonia, based on high performance metrics such as precision, recall, and ROC-AUC value. This will highly aid the healthcare industry. It will help radiographers to analyze X-rays immediately, thus resulting in a time-saving advantage. It will also work as a second opinion, specifically helpful in a busy healthcare center or a distant region. It will help doctors to detect diseases more accurately, thus leading to quicker diagnosis with higher confidence, finally having a positive effect on the patient. However, there are a few difficulties in the approach. We propose a new approach called NN-SVM, a smaller, combined approach to 3-class pneumonia detection in CXR images that combines a CNN with an RBF-SVM classifier. It attains an accuracy of 97.46% with great performance metrics. The technique is easily reproducible, scalable, thus making it a feasible solution for radiology assistance.

6. LIMITATIONS AND FUTURE WORKS

However, despite its efficiency, there exist several major cons in the model. The presence of data biases is a serious con because it might not be able to represent appropriately well the diversity of pneumonia infections in healthcare practical applications or settings. The efficiency of its operations might get disturbed and its performances could be reduced because of imbalances in data related to different kinds of pneumonia infections, for example, there could be an imbalance in images of patients with either bacterial pneumonia or viral infections, meaning more images of those with bacterial infections than those with viral infections.

The efficiency of its operations might get disturbed and its performances could be reduced because it could be able to perform well on common kinds or circumstances but might fail to perform well when it involves uncommon and not as well-represented kinds or circumstances. Another con is related to its interpretability because another serious con is related to its interpretability because CNNs, which belong to deep learning and perform excellently well in healthcare applications, work in an ineffective way because of its degree of obscurity because it could act as a 'black box' because its operations and performances might not be interpretable in an effective way because it could lack clear transparency regarding its operations and performances. To be able to perform its operations in an effective way because it requires quality and well-processed images of patients' lungs because it might get disturbed because of quality-related problems. Hence, it is a serious con because its applications would be limited because it might fail to perform well because healthcare applications regarding healthcare and related sectors would be limited because it might not be able to perform in an effective way because its applications might be limited because it might fail to perform in an effective way because its operations might get disturbed because it might lack clear transparency regarding its operations because its operations and performances might not be interpretable in an effective way because it could act as a 'black box' because its degree of obscurity might be increased because its operations and performances could be ineffective because it could lack clear transparency because it could be serious because it might fail to perform in an effective way because its applications would be limited because it might fail to be able to perform its operations because it might lack quality and well-processing because it might lack clear transparency.

Moreover, another con is related to its effects on healthcare applications. Hence, its quality might get disturbed because it might fail to perform in an effective way because diseases or infections related to lungs might be serious because it could be able to determine because its applications would be limited because it might fail to perform in an effective way because its operations might get disturbed because it might lack clear transparency because it could be serious because it might act as a barrier because it could be functioning because it might fail to perform in an effective way because its quality might get disturbed. Hence, another con is that its quality might get disturbed because it might fail to perform in an effective way because it could be able to determine because it could be functioning because its operations might get disturbed because it might lack clear transparency because it could be serious because it could act as a barrier because it might fail to perform in an effective way because it might lack quality and well-processing. However, despite its several cons, it has several pros, and those pros might.

With the advent of transfer learning strategies, further progress maybe achieved. Medical image categorization has made use of models like DenseNet and ResNet, which were pre-trained on massive datasets like ImageNet. Research has demonstrated that these pre-trained models can be refined on medical pictures to get great performance with comparatively smaller datasets: i) patient-wise temporal modeling across serial CXRs, ii) multi-modal integration with clinical metadata, iii) domain adaptation for cross-site robustness, and iv) calibrated risk stratification for triage.

ACKNOWLEDGMENTS

We would like to thank Dayananda Sagar University for support in completing this research work.

FUNDING INFORMATION

Authors state no funding involved.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Santosh Kumar Jankatti	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Raghavendra Srinivasaiah	✓	✓		✓	✓	✓	✓		✓	✓	✓			✓
Mohammad Shahina Parveen	✓				✓	✓	✓			✓	✓			✓
Harish H. Kenchannavar	✓		✓	✓	✓					✓	✓			✓
Danthuluri Sudha	✓		✓	✓	✓			✓		✓	✓			✓
Srihari Sharma Karigiri	✓	✓	✓	✓	✓	✓	✓			✓	✓			✓
Narah Mahadev Shivaraj						✓	✓			✓	✓		✓	✓

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

CONFLICT OF INTEREST STATEMENT

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. The authors also declare no non-financial competing interests (political, personal, religious, ideological, academic, or intellectual). Authors state no conflict of interest.

INFORMED CONSENT

Not applicable. This research exclusively utilized a publicly available, fully de-identified chest X-ray dataset sourced from Kaggle (https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia), which has been pre-processed to remove all personally identifiable information. As such, no individual patient identification is possible, and no new informed consent documentation was required for this study since the dataset is intended for open research use by the global scientific community.

ETHICAL APPROVAL

Not applicable. This study did not involve any new human subjects, animal experiments, or collection of primary data requiring ethical oversight. The research was conducted solely using existing, anonymized chest X-ray images from the Kaggle chest X-ray pneumonia dataset, which is publicly accessible and widely used in medical imaging research worldwide. No institutional review board (IRB) approval or equivalent ethical committee review was necessary as per standard practice for secondary analysis of de-identified public datasets.

DATA AVAILABILITY

The image dataset used for the research are taken from open access in Kaggle at https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia.

REFERENCES

[1] WHO, "Pneumonia in children," *World Health Organization*. Accessed: May 1, 2024. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/pneumonia

[2] M. I. Neuman *et al.*, "Variability in the interpretation of chest radiographs for the diagnosis of pneumonia in children," *Journal of Hospital Medicine*, vol. 7, no. 4, pp. 294–298, 2012, doi: 10.1002/jhm.955.

[3] G. J. Williams *et al.*, "Variability and accuracy in interpretation of consolidation on chest radiography for diagnosing pneumonia in children under 5 years of age," *Pediatric Pulmonology*, vol. 48, no. 12, pp. 1195–1200, 2013, doi: 10.1002/ppul.22806.

[4] D. Kermany, K. Zhang, and M. Goldbaum, "Labeled optical coherence tomography (OCT) and chest X-ray images for classification," *Mendeley Data*, 2018, doi: 10.17632/rscbjbr9sj.2.
[5] S. Lal *et al.*, "Adversarial attack and defence through adversarial training and feature fusion for diabetic retinopathy recognition," *Sensors*, vol. 21, no. 11, 2021, doi: 10.3390/s21113922.
[6] J. Devaraj *et al.*, "Forecasting of COVID-19 cases using deep learning models: is it reliable and practically significant?," *Results in Physics*, vol. 21, Feb. 2021, doi: 10.1016/j.rinp.2021.103817.
[7] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. F.- Fei, "ImageNet: a large-sScale hierarchical image database," *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255, 2009, doi: 10.1109/CVPR.2009.5206848.
[8] P. Rajpurkar *et al.*, "CheXNet: radiologist-level pneumonia detection on chest X-rays with deep learning," 2017, *arXiv:1711.05225*.
[9] S. Dalhoumi, G. Dray, J. Montmain, G. Derosière, and S. Perrey, "An adaptive accuracy-weighted ensemble for inter-subjects classification in brain-computer interfacing," *International IEEE/EMBS Conference on Neural Engineering, NER*, pp. 126–129, 2015, doi: 10.1109/NER.2015.7146576.
[10] S. Albahli, H. T. Rauf, A. Algosaibi, and V. E. Balas, "AI-driven deep CNN approach for multilabel pathology classification using chest X-rays," *PeerJ Computer Science*, vol. 7, pp. 1–17, 2021, doi: 10.7717/peerj-cs.495.
[11] T. Rahman *et al.*, "Transfer learning with deep convolutional neural network (CNN) for pneumonia detection using chest X-ray," *Applied Sciences*, vol. 10, no. 9, 2020, doi: 10.3390/app10093233.
[12] G. Liang and L. Zheng, "A transfer learning method with deep residual network for pediatric pneumonia diagnosis," *Computer Methods and Programs in Biomedicine*, vol. 187, 2020, doi: 10.1016/j.cmpb.2019.06.023.
[13] A. U. Ibrahim, M. Ozsoz, S. Serte, F. A. Turjman, and P. S. Yakoi, "Pneumonia classification using deep learning from chest X-ray images during COVID-19," *Cognitive Computation*, vol. 16, no. 4, pp. 1589–1601, 2024, doi: 10.1007/s12559-020-09787-5.
[14] U. Shah, A. A. Alrazeq, T. Alam, M. Househ, and Z. Shah, "An efficient method to predict pneumonia from chest X-rays using deep learning approach," *Studies in Health Technology and Informatics*, vol. 272, pp. 457–460, 2020, doi: 10.3233/SHTI200594.
[15] S. Albahli, H. T. Rauf, M. Arif, M. T. Nafis, and A. Algosaibi, "Identification of thoracic diseases by exploiting deep neural networks," *Computers, Materials and Continua*, vol. 66, no. 3, pp. 3139–3149, 2020, doi: 10.32604/cmc.2021.014134.
[16] T. B. Chandra and K. Verma, "Pneumonia detection on chest X-ray using machine learning paradigm," *Advances in Intelligent Systems and Computing*, vol. 1022 AISC, pp. 21–33, 2020, doi: 10.1007/978-981-32-9088-4_3.
[17] K. M. Kuo, P. C. Talley, C. H. Huang, and L. C. Cheng, "Predicting hospital-acquired pneumonia among schizophrenic patients: a machine learning approach," *BMC medical informatics and decision making*, vol. 19, no. 1, 2019, doi: 10.1186/s12911-019-0792-1.
[18] H. Yue *et al.*, "Machine learning-based CT radiomics method for predicting hospital stay in patients with pneumonia associated with SARS-CoV-2 infection: a multicenter study," *Annals of Translational Medicine*, vol. 8, no. 14, 2020, doi: 10.21037/atm-20-3026.
[19] H. Sharma, J. S. Jain, P. Bansal, and S. Gupta, "Feature extraction and classification of chest X-ray images using CNN to detect pneumonia," *Proceedings of the Confluence 2020 - 10th International Conference on Cloud Computing, Data Science and Engineering*, pp. 227–231, 2020, doi: 10.1109/Confluence47617.2020.9057809.
[20] O. Stephen, M. Sain, U. J. Maduh, and D. U. Jeong, "An efficient deep learning approach to pneumonia classification in healthcare," *Journal of Healthcare Engineering*, vol. 2019, 2019, doi: 10.1155/2019/4180949.
[21] J. D. Janizek, G. Erion, A. J. Degrave, and S. I. Lee, "An adversarial approach for the robust classification of pneumonia from chest radiographs," *ACM CHIL 2020 - Proceedings of the 2020 ACM Conference on Health, Inference, and Learning*, pp. 69–79, 2020, doi: 10.1145/3368555.3384458.
[22] J. Zhang *et al.*, "Viral pneumonia screening on chest X-rays using confidence-aware anomaly detection," *IEEE Transactions on Medical Imaging*, vol. 40, no. 3, pp. 879–890, 2021, doi: 10.1109/TMI.2020.3040950.
[23] T. Tuncer, F. Ozyurt, S. Dogan, and A. Subasi, "A novel Covid-19 and pneumonia classification method based on F-transform," *Chemometrics and Intelligent Laboratory Systems*, vol. 210, 2021, doi: 10.1016/j.chemolab.2021.104256.
[24] A. K. Jaiswal, P. Tiwari, S. Kumar, D. Gupta, A. Khanna, and J. J. P. C. Rodrigues, "Identifying pneumonia in chest X-rays: a deep learning approach," *Measurement: Journal of the International Measurement Confederation*, vol. 145, pp. 511–518, 2019, doi: 10.1016/j.measurement.2019.05.076.
[25] T. Gabruseva, D. Poplavskiy, and A. Kalinin, "Deep learning for automatic pneumonia detection," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1436–1443, 2020, doi: 10.1109/CVPRW50498.2020.00183.
[26] I. Pan, A. C.- Chênevert, and P. M. Cheng, "Tackling the radiological society of North America pneumonia detection challenge," *American Journal of Roentgenology*, vol. 213, no. 3, pp. 568–574, 2019, doi: 10.2214/AJR.19.21512.
[27] V. Rajinikanth, S. Kadry, R. Damasevicius, D. Taniar, and H. T. Rauf, "Machine-learning-scheme to detect choroidal-neovascularization in retinal OCT image," *2021 IEEE 7th International Conference on Bio Signals, Images and Instrumentation (ICBSII)*, Chennai, India, 2021, pp. 1-5, doi: 10.1109/ICBSII51839.2021.9445134.
[28] R. Kundu, H. Basak, P. K. Singh, A. Ahmadian, M. Ferrara, and R. Sarkar, "Fuzzy rank-based fusion of CNN models using Gompertz function for screening COVID-19 CT-scans," *Scientific Reports*, vol. 11, no. 1, 2021, doi: 10.1038/s41598-021-93658-y.

BIOGRAPHIES OF AUTHORS

Santosh Kumar Jankatti is currently working as associate professor in the Department of Computer Science and Technology at Dayananda Sagar University, Bangalore. He completed his Ph.D. degree in Computer Science and Engineering from Visvesvaraya Technological University (VTU), Belgaum, India in 2022 and has more than 14 years of teaching experience and 3 years of IT Industry experience. His interests include data mining, artificial intelligence, and big data. He can be contacted at email: sjankatti@gmail.com.

Raghavendra Srinivasaiah is currently working as associate professor in the Department of Computer Science and Engineering at CHRIST University, Bangalore. He completed his Ph.D. degree in Computer Science and Engineering from Visvesvaraya Technological University (VTU), Belgaum, India in 2017 and has more than 21+ years of teaching experience. His interests include data mining, artificial intelligence, and big data. He can be contacted at email: raghav.trg@gmail.com.

Dr. Mohammad Shahina Parveen is currently serving as the chairperson of the Department of Computer Science and Technology at Dayananda Sagar University, Bengaluru. With over 20 years of experience in higher education. She Completed Ph.D. in Computer Science and Engineering from Jawaharlal Nehru Technological University (JNTU), Hyderabad, India. She has authored 30+ research papers in reputed journals and conferences, published a book, and filed 40+ patents across diverse technical domains. She can be contacted at email: chairman-ct@dsu.edu.in.

Harish H. Kenchannavar received the M.E. degree from Walchand College of Engineering (WCE), Shivaji University, Sangli, India, in 2004, and the Ph.D. degree from Visvesvaraya Technological University (VTU), Belagavi, India, in 2014, under the guidance of Dr. Umakant P. Kulkarni. He is a professor and the head of the Department Computer Science and Engineering (Data Science), Dayananda Sagar College of Engineering, Bengaluru, India. He has presented and published a paper at national and IEEE international conferences. He has published more than 37 papers in reputed journals with an H-index of 4. He can be contacted at email: harishhk@git.edu.

Danthuluri Sudha is currently working as associate professor in the Department of Computer Science and Technology at Dayananda Sagar University, Bangalore. She completed her Ph.D. degree in Computer Science and Engineering from Visvesvaraya Technological University (VTU), Belgaum, India in 2022 and has more than 16 years of teaching experience. Her interests include cloud computing, artificial intelligence and machine learning. She can be contacted at email: danthulurisudha15@gmail.com.

Srihari Sharma Karigiri Narah is specialized in cardiorespiratory physiotherapy and rehabilitation and currently working as associate professor, College of Physiotherapy, School of Health Sciences, Dayananda Sagar University. With around 18 years of experience in the field of physiotherapy, he is one among a very few professionals in the field to have a doctoral degree to his credit. His interest includes amalgamation of technology with healthcare, lifestyle disorders, cardiac rehabilitation and yoga. He can be contacted at email: srihari-physio@dsu.edu.in.

Mahadev Shivaraj is currently working as assistant professor in the Department of Electronics and Communication engineering, at Dayananda Sagar Academy of Technology and Management, Bangalore. He pursuing his Ph.D. degree in from Visvesvaraya Technological University (VTU), Belgaum, India in 2022 and has more than 18 years of teaching experience and 1.5 years of IT Industry experience. His interests include biomedical image processing, image processing artificial intelligence, IoT, and very large-scale integration (VLSI). He can be contacted at email: prof.mahadevs@gmail.com.