HEALTHCARE AND MEDICAL SCIENCE

Pneumonia Detection Using Local Binary Pattern and Siamese Neural Network Architecture

Pratyusha Mishra, Vishesh Panghal

Department of Artificial Intelligence and Data Science, Poornima Institute of Engineering and Technology (PIET), Jaipur, Rajasthan, India

Published November 2026

Abstract

Convolutional neural networks detect pneumonia by evaluating the complete chest radiograph as a single image. For pediatric cases, such methods tend to capture aspects of patient posture, rib position, or scanning variation, instead of lung-specific abnormality. Radiologists read radiographs in a different way, namely, by comparing the left lung to the right lung to find unilateral opacities, and by looking at parenchymal texture for consolidation. In this paper, we propose a dual-domain Siamese network that matches the clinical reading process. Each radiograph is divided into left and right hemithorax crops, for which we generate a uniform Local Binary Pattern (LBP) texture map (P = 8, R = 1) and stack it with the grayscale image itself, forming a two-channel input. We then feed twin DenseNet-121 branches (using the same weights, and adapted to two-channel inputs), mapping each lung to a 128-dimensional feature space. Finally, we analyze bilateral asymmetry by calcu-lating the cosine similarity of the embeddings and training with a margin-based cosine embedding loss (m = 0.5). On the test set of 624 images from the Kermany et al. dataset, our model achieves an accuracy of 81.25% and an ROC-AUC of 0.844, which is superior to the baseline DenseNet-121 (0.609 AUC) and CheXNet (76.00%), without the need for pixel-level lung segmentations or ensemble methods of multiple models.

Keywords: bilateral symmetry; chest X-ray; DenseNet-121; local binary pattern; pneumonia detection; Siamese network

1 Introduction

Interpretation of anterior-posterior chest X-rays in pediatric emergencies is made difficult because of patient movement and rapid, shallow breathing. While adult patients can easily be asked to hold their breath, infant patients do not breathe in a regular pattern; therefore they have varying inspiratory volumes and movements across the boundaries of parenchyma [1], [2]. Normal anatomical landmarks like the sail-like thymus are located at the front part of the lungs and heart. These difficulties are easily taken care of by radiologists because of bilateral comparisons. Since pneumonia in emergency cases is usually asymmetric and occurs in one lobe only, bilateral comparison can help.

Traditional deep learning approaches treat the chest cavity as an intact image grid [3]–[7]. Deep neural networks like Inception-V3 [8] and CheXNet [9] may end up learning irrelevant aspects of images like positioning wedges, radiologic technician markings, and rib angulations [6], [10]. Global classification heads produce a scalar probability without localization, making it unclear to the clinician whether the prediction corresponds to a pathologic airspace consolidation or some other peripheral artifact [11].

In this paper, we test the effectiveness of a Siamese architecture with bilateral comparisons embedded into a weight-sharing DenseNet-121 network. The radiographic images are vertically

cropped into two images of size 224 × 112 with a 10% midline overlap to ensure inclusion of retrocardiac parenchyma. 2D uni-form Local Binary Pattern maps (P = 8, R = 1) are concatenated to grayscale images to ensure retention of texture details during the pooling operation. The two parallel DenseNet-121 networks map the left and right hemithorax into the same 128-dimensional embedding space, where cosine distance computes bilateral sim-ilarity. Setting the threshold to τ^* = 0.72 achieves an accuracy of 81.25% and an ROC-AUC of 0.844 on the Kermany bench-mark of 624 images without the use of pixel-level annotation or multi-branch ensemble methods.

2 Literature survey

Parenchymal consolidation changes the radiodensity of the lungs by replacing the alveolar air with inflammatory exudate, altering the local spatial frequencies. Early computer-aided detection algorithms captured these changes via GLCM and LBP [12]. Zaeri and Dib showed that uniform LBP captures the loss of vascular clarity within the exudative regions of the lungs. Even though these features are computationally efficient and robust to monotonic brightness shifts, their fixed neighborhood size prevents them from generalizing across different pediatric anatomy.

Convolutional neural networks learn hierarchical representations directly from the image pixels. Kermany et al. [8] used Inception-V3 for pediatric radiographs and reported more than 92% accuracy, while CheXNet [9] used a 121-layer DenseNet on the ChestX-ray14 dataset (F1 = 0.435 for all 14 pathologies). Transfer learning from adult-trained weights to the pediatric population is limited by substantial anatomical differences: high compliance of the chest wall, horizontal ribs, and relatively larger cardiothymic dimensions may cause unregularized networks to capture the geometry of the thorax rather than the parenchymal infiltration. Stephen et al. [13] constrained model capacity by using an 8-layer network with synthetic image transformations. Multi-model ensembles seek to minimize prediction variance by utilizing different network architectures. Ayan and Ünver [14] compared VGG16 with Xception for pediatric images, while variants of residual networks [15] and hybrid models [16] minimized intra-model conflicts. Chouhan et al. [17] and Liz et al. [18] used a combination of ResNet, DenseNet, and VGG networks to achieve benchmark accuracies above 96%. Running several models in parallel increases GPU memory consumption and inference time, making such models challenging to deploy in a low-performance clinical environment. Capsule networks [6] encode relative spatial relations through an iterative process called routing, but require computationally expensive training.

Comparing findings from the contralateral image provides an alternative approach. Schneider et al. [19] used Siamese DenseNets for detecting unilateral abnormality in adult chest radiographs, without including micro-textures in their network. Wang et al. [20] incorporated LBP and HOG features into convolutional layers while processing the radiograph as a single unified image. Kör et al. [21] evaluated the bilateral difference of features for distinguishing viral pneumonia from COVID-19. The approaches mentioned above are summarized in Table 1. The proposed architecture incorporates the 2D LBP map directly with the raw image input, using a shared-weight DenseNet-121 Siamese backbone to detect bilateral asymmetry while preserving low inference latency.

Table 1. Comparison of pneumonia detection studies.

Reference (Journal, Year)	Methodology	Accuracy	Drawback
Kermany et al. (Cell, 2018) [8]	Transfer learning with Inception-V3 CNN	92–93%	Limited to pediatric data; no localization provided.

Reference (Journal, Year)	Methodology	Accuracy	Drawback
Rajpurkar et al. (arXiv, 2017) [9]	CheXNet – 121-layer DenseNet	76% AUROC	Trained on adult data; possible overfitting to biases.
Stephen et al. (J Healthc Eng, 2019) [13]	Custom 8-layer CNN (no TL)	93.7%	Requires extensive augmentation; lower accuracy than deeper nets.
Jaiswal et al. (Measurement, 2019) [22]	CNN with Grad-CAM (Mask R-CNN)	94.6%	Complex pipeline; needs bounding box annotations.
Chouhan et al. (Appl. Sci, 2020) [17]	Ensemble of ResNet, DenseNet, VGG	96.4%	High complexity ensemble; marginal gain per model.
Elshennawy et al. (Diagnostics, 2020) [2]	Deep-Pneumonia: AlexNet + VGG ensemble	94.8%	Uses two large CNNs; slow inference.
Wang et al. (Comp. Math. Meth. Med, 2022) [20]	FC-VGG: VGG16 with LBP & HOG	92.19%	High parameter count; risk of overfitting.
Bodapati et al. (Measurement, 2022) [6]	ChxCapsNet – single CapsNet	93.1%	Lower accuracy than ensemble; computationally heavy.
Liz et al. (FGCS, 2021) [18]	Stacked ensemble of 5 CNNs	96.0%	Resource-intensive ensembling.
Ayan et al. (IEEE EBBT, 2019) [14]	VGG16 vs Xception transfer learning	87.0% (VGG16)	Lower accuracy; no advanced augmentation.

3 Background

3.1 Radiographic texture and local binary patterns (LBP)

Healthy lungs display branching bronchovascular patterns, tapering toward the thoracic wall, against radiolucent alveolar air. In bacterial and viral pneumonia, inflammatory exudate replaces alveolar air, leading to enhanced local attenuation and attenuated peripheral vascular margins.

The LBP operator [23] detects these radiologic micro-textural changes by thresholding the intensities in a neighborhood relative to a central pixel. For a given center pixel g_c, P equidistant neighbors are selected around a circle of radius R:

$$s\left(g_p - g_c\right) = \begin{cases} 1, & g_p \geq g_c \\ 0, & g_p < g_c \end{cases} \tag{1}$$

where g_p is the pixel intensity of the neighbor with index $p \in \{0, \ldots, P-1\}$. Bit-weighting yields the integer code:

$$\mathrm{LBP}_{P,R}\left(x,\ y\right) = \sum_{p=0}^{P-1} s\left(g_p - g_c\right) 2^p \tag{2}$$

The configuration uses $P = 8$ and $R = 1$ (Figure 2), with only uniform patterns, allowing at most two circular transitions. Uniform patterns filter out noise without destroying step edges and corner junctions. Instead of collapsing spatial patterns into a 1D histogram, the 2D pattern map I_{LBP} is concatenated with the grayscale image I_{gray}. This gives the convolutional layers simultaneous access to both radiographic density and radiological micro-texture.

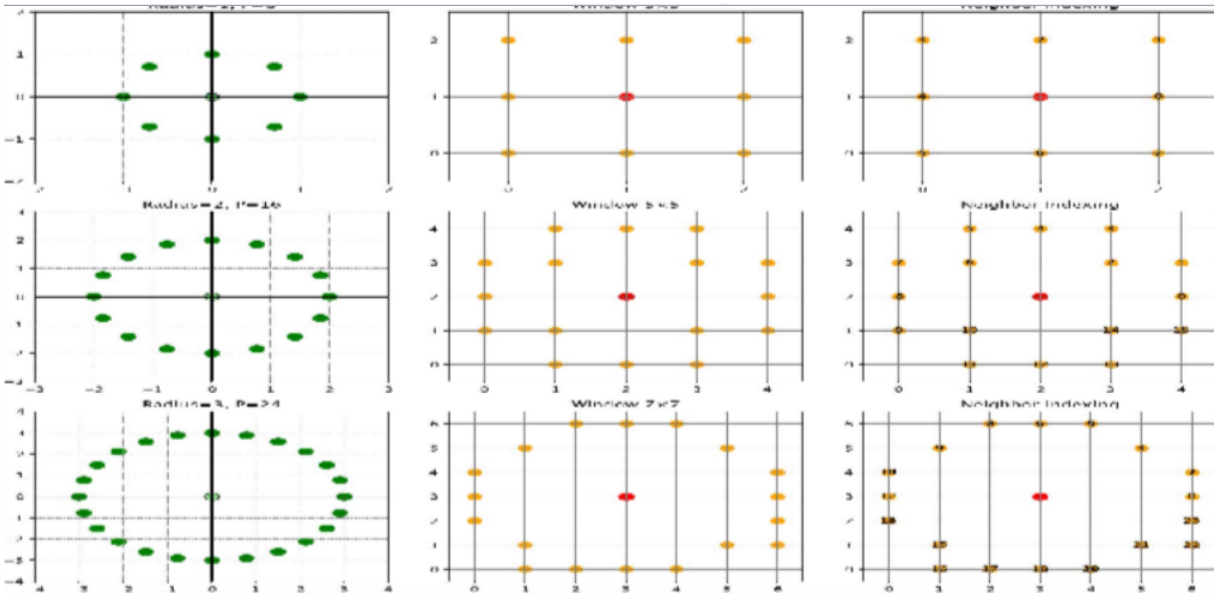


Figure 2. Visual depiction of LBP sampling with different radii (R = 1, 2, 3). Each row illustrates the sampling pattern (left), the corresponding spatial window (centre) and the index assigned to each neighbor (right) used in LBP code generation.

3.2 Learning of contralateral representations by Siamese networks

The Siamese neural network consists of two identical subnetworks with shared weights W, illustrated in Figure 1. Due to the shared weights, the two branches map pairs of inputs into the same latent space [24]. Comparison of the two opposite sides of the chest is used as a natural control [25]. Let I_L and I_R denote the left and right hemithorax crops. After processing both inputs with the common network $f_W(\cdot)$, we get two feature vectors:

$$x_1 = f_W(I_L) \in \mathbb{R}^{128},\ x_2 = f_W(I_R) \in \mathbb{R}^{128} \tag{3}$$

We then measure their similarity using the cosine:

$$\cos\left(x_1,\ x_2\right) = \frac{x_1 \cdot x_2}{\|x_1\|_2 \|x_2\|_2} \tag{4}$$

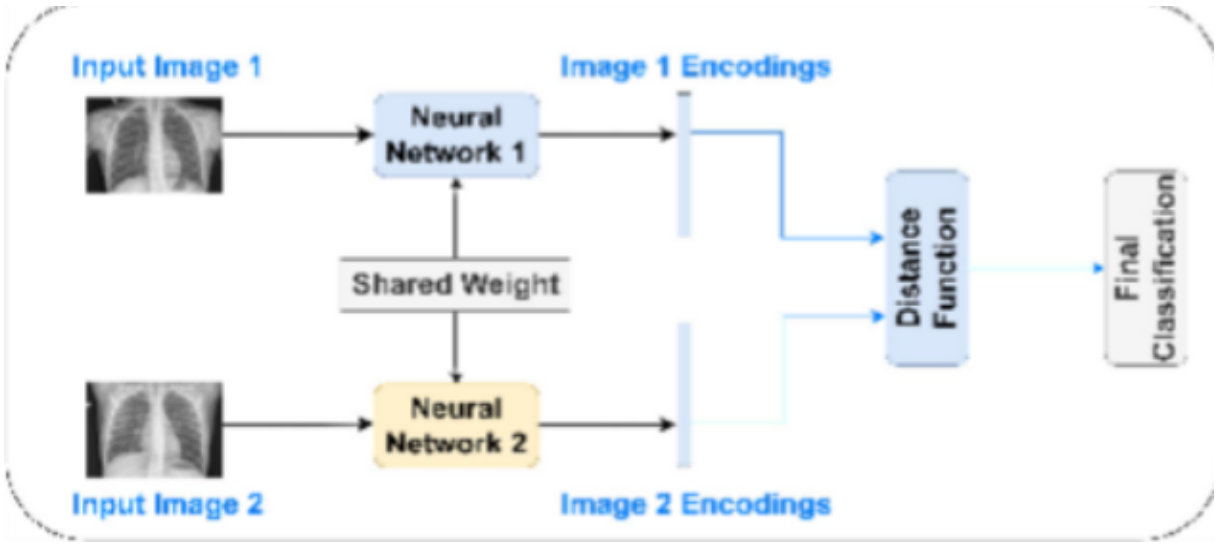


Figure 1. Siamese network configuration: paired hemithorax inputs mapped into a shared metric space for cosine similarity evaluation.

The architecture uses two-channel concatenation of the input images (I_{gray} and I_{LBP}) in the input layer and metric comparison in the output layer. Each subnetwork operates separately on its own hemithorax and computes asymmetry in the last latent layer. Figure 3 shows the DenseNet-121 backbone network [26].

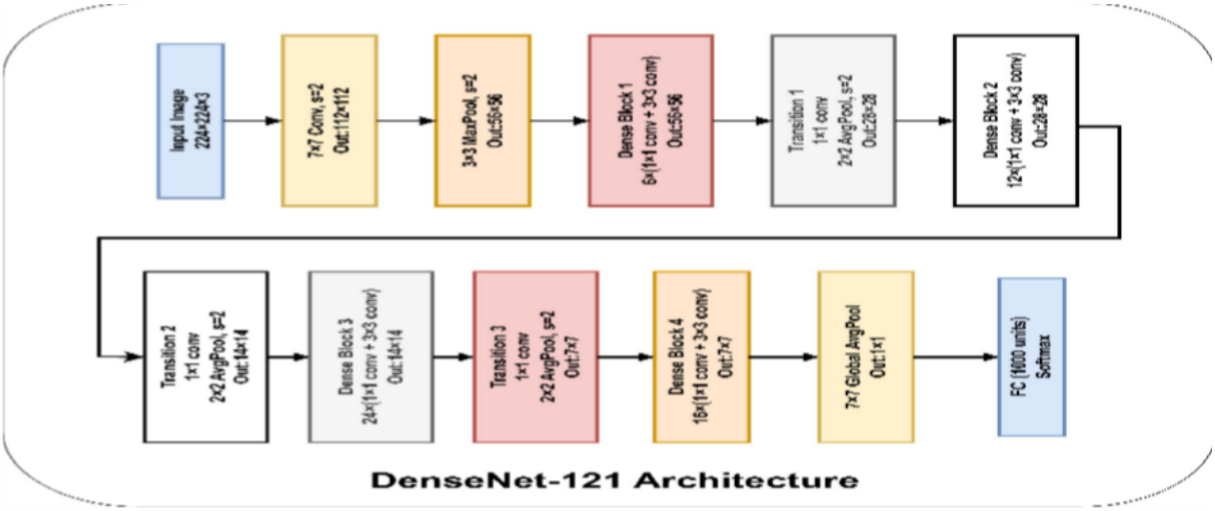


Figure 3. DenseNet-121 backbone diagram showing dense connectivity blocks, 1 × 1 bottleneck convolutions and transition downsampling layers.

4 Dataset

The experimental validation is based on the set of pediatric chest X-ray images published by Kermany et al. [8], obtained from the Guangzhou Women and Children's Medical Center. The dataset contains 5,863 anterior-posterior and posterior-anterior X-ray images of patients aged between one and five years, whose diagnoses were clinically confirmed prior to publication.

We continue to use the standard benchmark split of 5,216 training images (1,341 normal and 3,875 pneumonia cases), 16 validation images for parameter tuning (8 normal and 8 pneumonia cases), and 624 testing images (234 normal and 390 pneumonia cases). Bacterial and viral pneumonia cases are treated as a single positive class in a two-class problem. All images are rescaled to 224 × 224 pixels, and pixel intensity values are scaled to the range [0.0, 1.0].

Figure 4 compares normal lung aeration to typical consolidations.

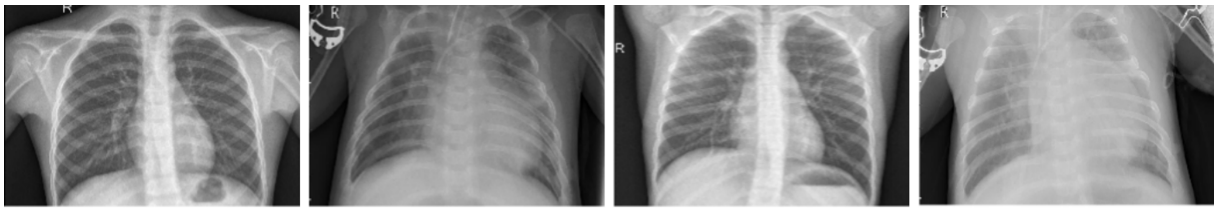


Figure 4. Sample chest radiographs from the Kermany cohort: clear aeration and preserved peripheral vascular branches in normal subjects (a, b) contrasted with focal and interstitial opacities in pediatric pneumonia cases (c, d).

5 Methodology

5.1 Pre-processing and hemithorax partitioning

Each normalized 224 × 224 radiograph is divided along the vertical midline into left and right crops of size 224 × 112 pixels. The 10% overlap along the central axis prevents truncation of retrocardiac lung parenchyma and mediastinal margins. Both crops are zero-padded to square dimensions (224 × 224). A 2D uniform $LBP_{8,1}$ texture map is calculated for each grayscale crop and stacked along the channel axis, producing a two-channel tensor:

$$I_{\text{input}} = \left[I_{\text{gray}}; I_{\text{LBP}}\right] \in \mathbb{R}^{2\times224\times224} \tag{5}$$

The training data augmentation includes random rotations within ±5°, affine shearing with 5% probability, brightness changes within ±10%, and random horizontal swapping ($I_L \leftrightarrow I_R$) to maintain invariance to the hemithorax input order.

5.2 Dual-domain Siamese network architecture

Each branch uses a DenseNet-121 backbone [26] adapted for two-channel input. The first 7 × 7 convolutional layer is adjusted from 3 to 2 input channels ($R^{64\times2\times7\times7}$) with He-normal weight initialization. The standard 1,000-class classification layer is replaced by global average pooling followed by a linear projection layer that outputs a 128-dimensional embedding vector ($x \in R^{128}$). The weights W are shared between both branches, projecting the left and right inputs into a common metric space. Figure 5 depicts the dual-domain architecture.

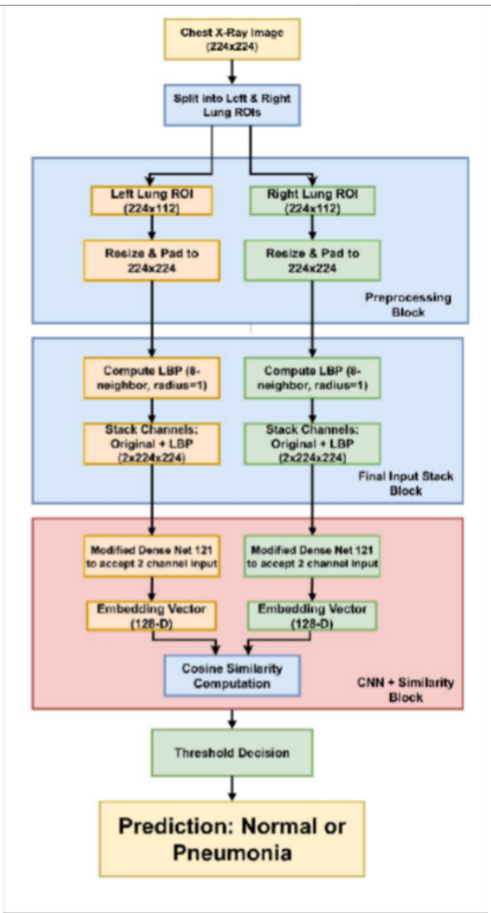


Figure 5. Dual-domain Siamese workflow: pairing raw pixel attenuation with LBP texture descriptors across weight-sharing DenseNet-121 branches for contralateral cosine metric evaluation.

5.3 Training strategy and loss function

The network is trained using Cosine Embedding Loss on paired hemithorax embeddings (x_1, x_2) with target labels $y \in \{+1, -1\}$ representing normal controls (+1) and pneumonia (−1):

$$L\left(x_1,\ x_2,\ y\right) = \begin{cases} 1 - \cos\left(x_1,\ x_2\right), & y = +1 \\ \max\left(0,\ \cos\left(x_1,\ x_2\right) - m\right), & y = -1 \end{cases} \tag{6}$$

where $\cos(x_1, x_2)$ is given by Equation (4), and the contrastive margin is set to $m = 0.5$. For normal scans, the loss penalizes the angular separation between embeddings. For pneumonia scans, pairs with cosine similarity below 0.5 incur zero loss, concentrating optimization on ambiguous pairs close to the decision boundary. Training is performed using the AdamW optimizer with a learning rate $\eta = 1 \times 10{-}4$ and weight decay $\lambda = 1 \times 10{-}4$. The mini-batches consist of 32 image pairs (64 crops). Training runs for up to 50 epochs using Automatic Mixed Precision (AMP), with early stopping on validation loss at a patience of 10 epochs.

5.4 Testing and inference protocol

For patient screening, the cosine similarity $s = \cos(x_1, x_2)$, which measures the degree of similarity between the extracted embeddings of the two hemithoraces, is evaluated. An operational threshold τ^* determines the final diagnosis:

$$\hat{y} = \begin{cases} \text{Pneumonia (1)}, & s < \tau^* \\ \text{Normal (0)}, & s \geq \tau^* \end{cases} \tag{7}$$

The threshold τ^* is chosen by sweeping $\tau \in [0.0, 1.0]$ in increments of 0.01 on the 16-sample validation set. The optimal threshold $\tau^* = 0.72$ maximizes the validation macro F1-score and is kept constant throughout all test evaluations. Algorithms 1 and 2 summarize the threshold calibration and inference procedures.

The implementation code and pre-trained model weights are publicly available at: https://github.com/Vishesh-panghal/Pneumonia-Detection-via-Radiographic-Texture.

Algorithm 1. Validation threshold calibration

Require: left–right validation pairs $\{(x_1^i, x_2^i, y^i)\}$, $i = 1, \ldots, N$; search step δ

Ensure: selected operational cut-off τ^*

Initialize best_f1 ← 0, $\tau^* \leftarrow 0.5$

for $\tau = 0.0$ to 1.0 step δ **do**

  Initialize correct ← 0

  for each sample $i \in \{1, \ldots, N\}$ **do**

    $s^i \leftarrow \cos(x_1^i, x_2^i)$

    if $(s^i \geq \tau \wedge y^i = +1) \vee (s^i < \tau \wedge y^i = -1)$ **then**

      correct ← correct + 1

    end if

  end for

Compute validation F1-score: $F1(\tau)$

if $F1(\tau) >$ best_f1 **then**

  best_f1 ← $F1(\tau)$

  $\tau^* \leftarrow \tau$

end if

end for

return τ^*

Algorithm 2. Operational Siamese inference protocol

Require: radiograph input I, Siamese parameters W, decision cut-off τ^*

Ensure: predicted class $\hat{y} \in \{\text{Normal } (0), \text{Pneumonia } (1)\}$, similarity s

Split I into hemithoracic regions: I_L, I_R

Compute uniform LBP texture maps: $L_L = \text{LBP}(I_L)$, $L_R = \text{LBP}(I_R)$

Construct 2-channel tensors: $T_L = [I_L; L_L]$, $T_R = [I_R; L_R]$

Extract 128-D latent embeddings: $x_1 \leftarrow f_W(T_L)$, $x_2 \leftarrow f_W(T_R)$

Calculate cosine similarity: $s \leftarrow \cos(x_1, x_2)$

if $s < \tau^*$ **then**

7: **return** Pneumonia (1), s

8: **else**

9: **return** Normal (0), s

10: **end if**

6 Results and evaluation

We test the model on the 624-image held-out test set from Kermany et al.

6.1 Evaluation metrics and confusion matrix

We measure the performance of the screening process on the held-out test cohort of 624 images using Accuracy, Sensitivity (Recall), Specificity, Precision, and macro F1-score. The resulting test confusion matrix, evaluated at $\tau^* = 0.72$, is shown in Figure 6.

On the 624 test radiographs, the network correctly flagged 240 of 324 pneumonia scans (80.00% sensitivity) while ruling out 246 of 300 healthy controls (82.00% specificity). This left 54 false positives and 60 false negatives, resulting in an overall test accuracy of 81.25%, precision of 81.63%, and a macro F1 of 0.8081.

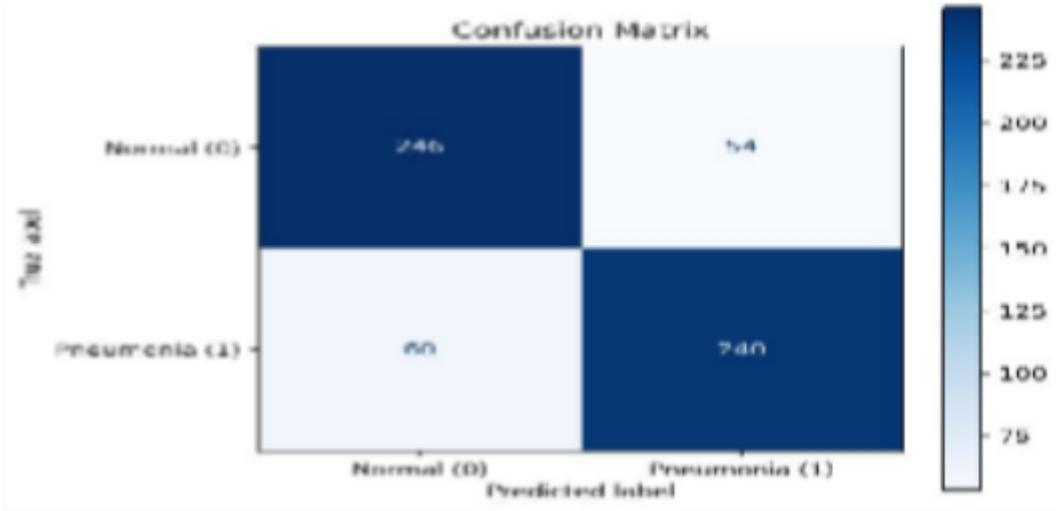


Figure 6. Confusion matrix on the held-out test set. Class 0 = normal, class 1 = pneumonia.

6.2 Training curves and ROC analysis

The training and validation loss trajectories are plotted in Figures 7 and 8. The loss decreased steadily throughout the first 25 epochs. The validation accuracy reached about 80% by epoch 30 and stayed steady until epoch 50, showing stable convergence with no late divergence.

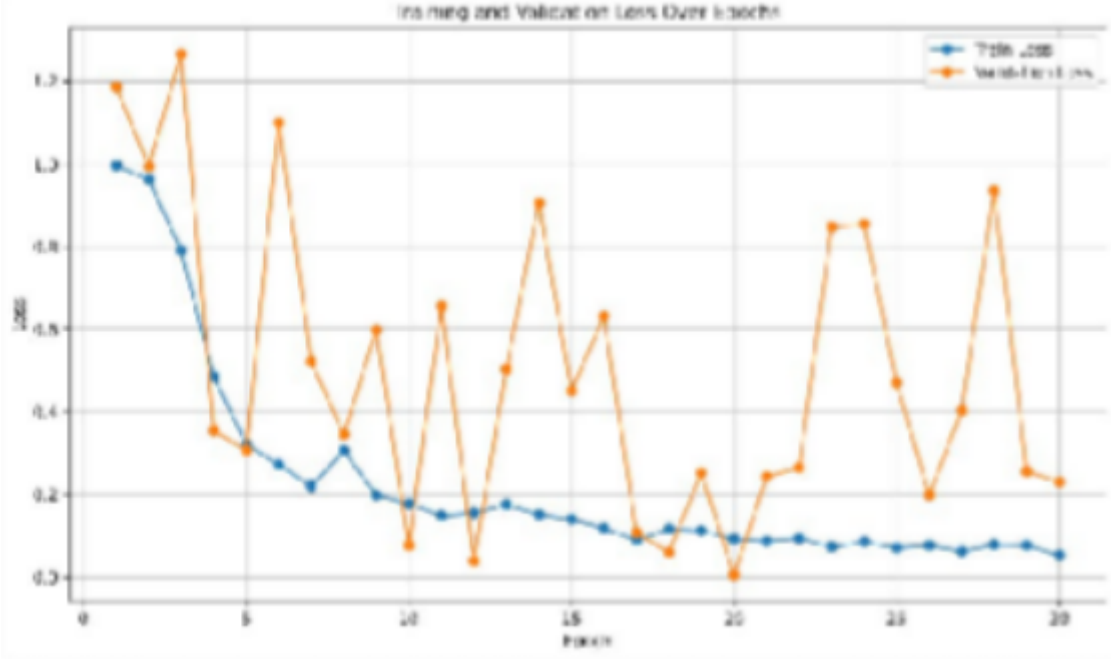


Figure 7. Optimization loss curves across 50 epochs.

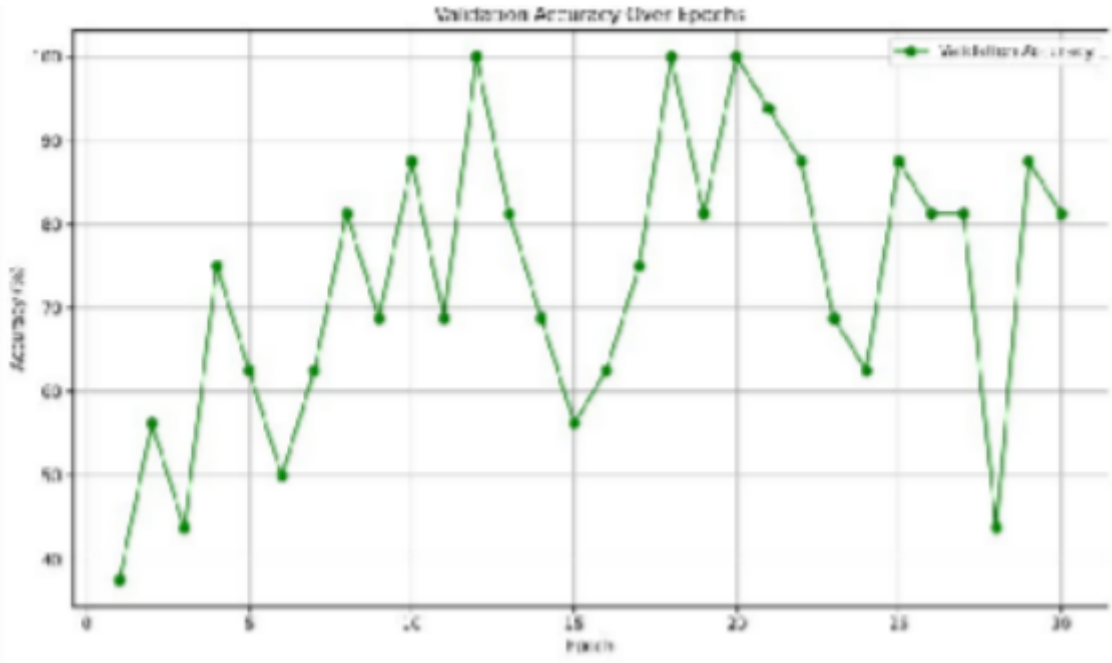


Figure 8. Validation accuracy progression over 50 epochs.

Figure 9 displays the ROC curve computed over the contralateral dissimilarity values ($d = 1 - \cos(x_1, x_2)$), corresponding to an area under the curve of ROC-AUC = 0.844.

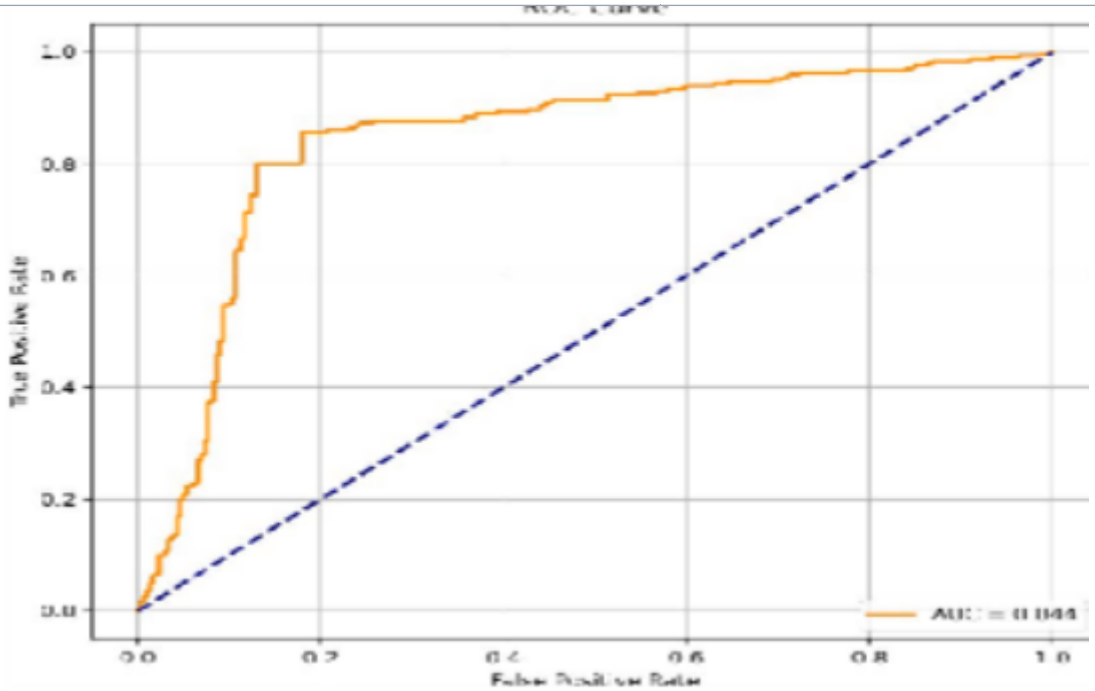


Figure 9. Receiver operating characteristic (ROC) curve for pneumonia detection via Siamese contralateral asymmetry (AUC = 0.844).

Normalized uniform LBP code frequencies for clear and consolidated lung patches are presented in Figure 10. Aerated parenchyma shows elevated counts in the upper uniform bins because of high-contrast vascular interfaces against air. Fluid infiltration reduces edge contrast and distributes counts toward intermediate bins. Feeding the 2D texture map directly into the network provides these micro-structural cues alongside the raw intensity values.

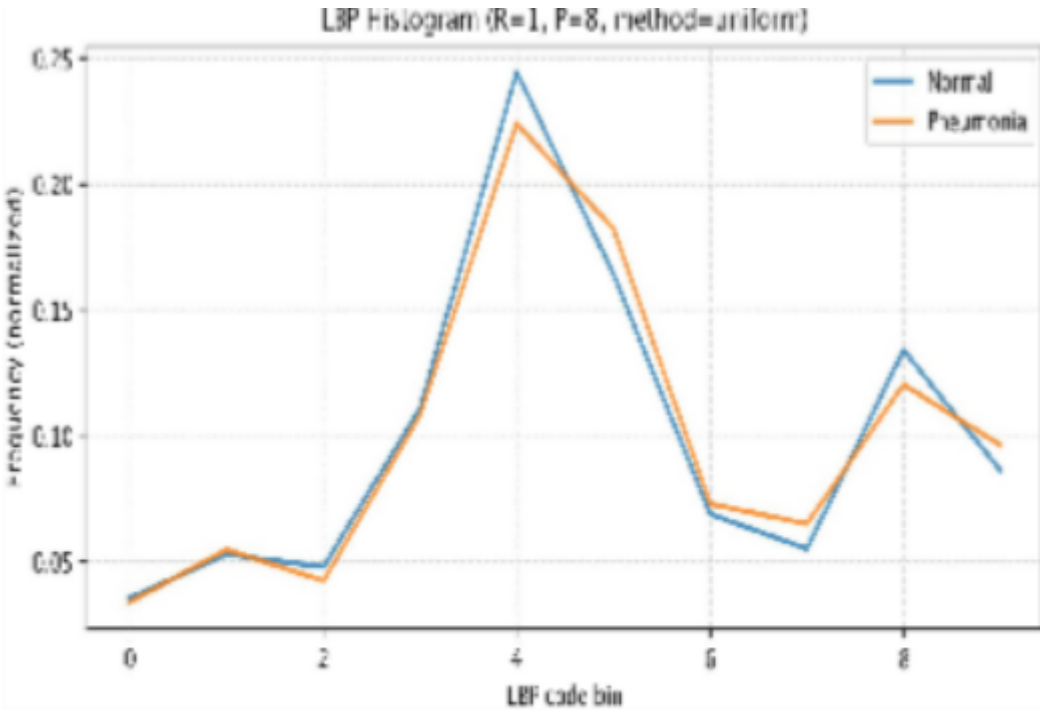


Figure 10. Normalized LBP code distributions for normal versus pneumonic tissue (P = 8, R = 1).

6.3 Comparison with existing models

Table 2 compares our test metrics with published models evaluated on the same Kermany benchmark.

Table 2. Performance comparison of pneumonia detection models on pediatric chest radiographs.

Author (year)	Core Methodology / Architecture	Accuracy (%)	ROC-AUC
Liang et al. (2020) [27]	Transfer Learning via VGG16 Residual Network	74.20	–
Sirazitdinov et al. (2019) [28]	Mask R-CNN + RetinaNet Ensemble	75.80	–
Rajpurkar et al. (2017) [9]	CheXNet – 121-layer DenseNet	76.00	–
Antin et al. (2017) [10]	DenseNet-121 Baseline	–	0.6090
Odaibo et al. (2019) [29]	Weakly Supervised Inception-v4	78.20	–
Hofmeister et al. (2024) [30]	Monolithic CNN Classifier	72.60	0.7400
Khaniki et al. (2023) [31]	Vision Transformer (ViT) with Re-attention	79.63	–
Varshni et al. (2019) [32]	DenseNet-169 Feature Extractor + SVM	80.02	0.8002
Gu et al. (2018) [33]	CAD System + Fully Convolutional Net	80.40	0.8234
Proposed Method	Dual-Domain Siamese DenseNet-121 (LBP + Asymmetry)	81.25	0.8440

The proposed dual-domain Siamese pipeline reaches 81.25% accuracy and an ROC-AUC of 0.844. This performance surpasses standard monolithic convolutional networks, including VGG16 residual transfer learning [27] (74.20%), CheXNet [9] (76.00%), weakly-supervised Inception-v4 [29] (78.20%), Vision Transformers [31] (79.63%), and DenseNet-169 with SVM classification [32] (80.02%). It likewise outperforms the detection ensemble (Mask R-CNN and RetinaNet) of Sirazitdinov et al. [28] (75.80%).

While complex multi-network ensembles (such as Chouhan et al. [17] at 96.4%) achieve higher raw scores through parallelization, they do so at the expense of several high-capacity backbones, multiplied inference time, and substantial GPU memory demands. In contrast, our pipeline uses a single DenseNet-121 backbone operating without region bounding-box supervision, keeping inference lightweight and practical for primary clinic hardware.

7 Conclusion and future directions

In this study, a dual-domain Siamese network was evaluated for the detection of pediatric pneumonia using bilateral asymmetry and Local Binary Patterns. By dividing each X-ray into contralateral hemithoraces and merging raw grayscale intensities with uniform LBP maps, twin weight-sharing DenseNet-121 backbones achieve 81.25% accuracy and 0.844 ROC-AUC on the 624-image Kermany benchmark without bounding-box supervision or multi-model ensembles. While symmetrical bilateral consolidations are an inherent constraint of contralateral comparison, integrating relative metric learning with normative anatomical atlases outlines the direction for future work.

References

[1] Zhang D, Ren F, Li Y, Na L, Ma Y. Pneumonia detection from chest X-ray images based on convolutional neural network. *Electronics*. 2021;10(13):1512.

[2] Elshennawy NM, Ibrahim DM. Deep-Pneumonia framework using deep learning models based on chest X-ray images. *Diagnostics*. 2020;10(9):649.

[3] Abe H, MacMahon H, Shiraishi J, Li Q, Engelmann R, Doi K. Computer-aided diagnosis in chest radiology. *Seminars in Ultrasound, CT and MRI*. 2004;25(5):432–437.

[4] Khawaja MA, Aziz MZ, Iqbal N. Effectual lung segmentation for CAD systems using CT scan images. In: *Proceedings of the 8th International Multitopic Conference (INMIC)*. IEEE. 2004:49–54.

[5] Awai K, Murao K, Ozawa A, Komi M, Hayakawa H, Hori S, Nishimura Y. Pulmonary nodules at chest CT: effect of computer-aided diagnosis on radiologists' detection performance. *Radiology*. 2004;230(2):347–352.

[6] Antin B, Kravitz J, Martayan E. *Detecting pneumonia in chest X-rays with supervised learning*. Stanford University Machine Learning Technical Report. 2017.

[7] Tilve A, Nayak S, Vernekar S, Turi D, Shetgaonkar PR, Aswale S. Pneumonia detection using deep learning approaches. In: *Proceedings of the 2020 International Conference on Emerging Trends in Information Technology and Engineering (ic-ETITE)*. IEEE. 2020:1–8.

[8] Kermany DS, Goldbaum M, Cai W, Valentim CCS, Liang H, Baxter SL, McKeown A, Yang G, Wu X, Yan F, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*. 2018;172(5):1122–1131.

[9] Rajpurkar P, Irvin J, Zhu K, Yang B, Mehta H, Duan T, Ding D, Bagul A, Langlotz C, Shpanskaya K, et al. CheXNet: radiologist-level pneumonia detection on chest X-rays with deep learning. *arXiv* preprint arXiv:1711.05225. 2017.

[10] Antin B, Kravitz J, Martayan E. *Evaluating deep architectures for chest radiograph pathology classification*. Stanford University CS229 Project Report. 2017.

[11] Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2017:618–626.

[12] Zaeri N, Dib RK. *X-ray pneumonia detection using angular and radial local binary patterns fusion*. Journal of Engineering Research. 2024.

[13] Stephen O, Al-antari MA, Lee SM, Ryu JG, Hossain MS, Muhammad G. An efficient deep learning approach to pneumonia classification in healthcare. *Journal of Healthcare Engineering*. 2019;2019:4180949.

[14] Ayan E, Ünver HM. Diagnosis of pneumonia from chest X-ray images using deep learning. In: *Proceedings of the 2019 Scientific Meeting on Electrical-Electronics and Biomedical Engineering and Computer Science (EBBT)*. IEEE. 2019:1–5.

[15] He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016:770–778.

[16] Rajaraman S, Sornapudi S, Alderson PO, Folio LR, Antani SK. Analyzing inter-reader variability affecting deep ensemble learning for COVID-19 detection in chest radiographs. *PLOS ONE*. 2020;15(11):e0242301.

[17] Chouhan V, Singh SK, Khamparia A, Gupta D, Tiwari P, Moreira C, Damaševičius R, de Albuquerque VHC. A novel transfer learning based approach for pneumonia detection in chest X-ray images. *Applied Sciences*. 2020;10(2):559.

[18] Liz H, Sánchez-Montañés M, Tagarro A, Domínguez-Rodríguez S, Dagan R, Camacho D. Ensembles of convolutional neural network models for pediatric pneumonia diagnosis. *Future Generation Computer Systems*. 2021;122:220–233.

[19] Schneider H, Yildiz EC, Biesner D, Layer YC, Wulff B, Nowak S, Theis M, Sprinkart AM, Attenberger UI, Sifa R. Symmetry-aware Siamese network: exploiting pathological asymmetry for chest X-ray analysis. In: *Proceedings of the International Conference on Artificial Neural Networks (ICANN)*. Springer. 2023:165–176.

[20] Wang HN, Zheng LX, Pan SW, Yan T, Su QL. Image recognition of pediatric pneumonia based on fusion of texture features and depth features. *Computational and Mathematical Methods in Medicine*. 2022;2022:1973508.

[21] Kör H, Erbay H, Yurttakal AH. Diagnosing and differentiating viral pneumonia and COVID-19 using X-ray images. *Multimedia Tools and Applications*. 2022;81(27):39041–39057.

[22] Jaiswal AK, Tiwari P, Kumar S, Gupta D, Khanna A, Rodrigues JJPC. Identifying pneumonia in chest X-rays: a deep learning approach. *Measurement*. 2019;145:511–518.

[23] Ojala T, Pietikäinen M, Mäenpää T. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2002;24(7):971–987.

[24] Bromley J, Guyon I, LeCun Y, Säckinger E, Shah R. Signature verification using a "Siamese" time delay neural network. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 6. 1993:737–744.

[25] Yoshida H. Local contralateral subtraction for lung nodule detection in digital chest radiographs. *Academic Radiology*. 2004;11(3):278–287.

[26] Huang G, Liu Z, van der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017:4700–4708.

[27] Liang G, Zheng L. A transfer learning method with deep residual network for pediatric pneumonia diagnosis. *Computer Methods and Programs in Biomedicine*. 2020;187:104964.

[28] Sirazitdinov I, Kholiavchenko M, Mustafaev T, Yuan Y, Kuleev R, Ibragimov B. Deep neural network ensemble for pneumonia localization from a large-scale chest X-ray database. *Computers & Electrical Engineering*. 2019;78:388–399.

[29] Odaibo D, Zhang Z, Skidmore F, Tanik M. Detection of visual signals for pneumonia in chest radiographs using weak supervision. In: *Proceedings of the 2019 IEEE SoutheastCon*. IEEE. 2019:1–5.

[30] Hofmeister J, Garin N, Montet X, Scheffler M, Platon A, Poletti PA, Stirnemann J, Debray MP, Claessens YE, Duval X, et al. Validating the accuracy of deep learning for the diagnosis of pneumonia on chest X-ray against a robust multimodal reference diagnosis. *European Radiology Experimental*. 2024;8(1):20.

[31] Khaniki MAL, Mirzaeibonehkhater M, Manthouri M. Enhancing pneumonia detection using vision transformer with dynamic mapping re-attention mechanism. In: *Proceedings of the 2023 13th International Conference on Computer and Knowledge Engineering (ICCKE)*. IEEE. 2023:144–149.

[32] Varshni D, Thakral K, Agarwal L, Nijhawan R, Mittal A. Pneumonia detection using CNN based feature extraction. In: *Proceedings of the 2019 IEEE International Conference on Electrical, Computer and Communication Technologies (ICECCT)*. IEEE. 2019:1–7.

[33] Gu X, Pan L, Liang HY, Yang R. Classification of bacterial and viral childhood pneumonia using deep learning in chest radiography. In: *Proceedings of the 3rd International Conference on Multimedia and Image Processing (ICMIP)*. ACM. 2018:88–93.