Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

An Efficient and Interpretable Deep Learning Framework for Pneumonia Detection and Severity Assessment from Chest Radiographs

S Kumar Reddy Mallidi*[1], Anant Wairagade[2], Dr. Gunjan Chhabra[3], Kandula Rojarani[4], Dr. Omkar S. Vaidya[5]
[1]*CSE, Sri Vasavi Engineering College, Tadepalligudem, INDIA*
[2]*Independent Researcher, Phoenix, USA*
[3]*School of Science and Technology, SRHU, Dehradun, INDIA*
[4]*CSE-AIML, Aditya University, Surampalem, INDIA*
[5]*MKSSS's Maharshi Karve Mahila Mahavidyalay, Wai, INDIA*
**Coresponding Author: satya.cnis@gmail.com*
0000-0001-5908-8316

ARTICLE INFO

Received: 19 Nov 2024

Revised: 28 Dec 2024

Accepted: 18 Jan 2025

ABSTRACT

Pneumonia presents considerable challenges to healthcare systems worldwide, necessitating swift and accurate diagnostic methods. This study introduces an integrated artificial intelligence approach combining a convolutional neural network classification with infection region segmentation to improve pneumonia diagnosis and assess severity. Our dual-function model was crafted using the QaTa-COVID-19 and COVID-19 Radiography datasets, enabling strong differentiation between COVID-19 pneumonia, non-COVID pneumonia, and standard chest X-rays. The proposed framework combines VGG16 convolutional features with DenseNet connectivity patterns, achieving a classification accuracy of 98.71% while using only about 8 million parameters, significantly surpassing traditional models like VGG16, ResNet50, and DenseNet121. We incorporated a UNet-based segmentation component to assess disease severity by effectively outlining infection regions. Gradient-weighted Class Activation Mapping (Grad-CAM) visualizations improve interpretability of model, providing clinicians with insights into decision-making. Our results indicate that this hybrid method enhances diagnostic precision, infection quantification, and algorithmic transparency. Future research will explore multimodal data fusion and investigate transformer architectures to boost pneumonia detection capabilities. This study contributes to developing an automated, explainable diagnostic framework that effectively aids clinical decision-making in pneumonia management.

Keywords: Deep Learning, Computer-Aided Diagnosis, Pneumonia Detection, Chest X-ray Analysis, Explainable AI

INTRODUCTION

COVID-19 epidemic has drawn focus to critical requirement for accurate medical imaging diagnostic techniques. Medical image analysis has been completely transformed by DL(deep learning), particularly CNNs (Convolutional Neural Networks), which automatically extract features from raw images and handle high-resolution datasets like "ChestX-ray8" with efficiency [1]. Although transformers' self-attention processes have made them popular in medical imaging, their use is limited by their high processing costs and reliance on sizable datasets, particularly in clinical contexts where data access is restricted by privacy concerns [2]. CNNs, conversely, demonstrate superior performance with smaller datasets and greater computational efficiency.

Transfer learning improves CNNs by utilizing pre-trained models, lowering processing requirements, and facilitating quick task adaptation with minimum training data—all of which are essential benefits for identification of pneumonia [3]. Additionally, CNNs offer strong frameworks for medical imaging segmentation and classification that are remarkably flexible to different imaging scenarios and pathological presentations [4]. This versatility facilitates their integration into clinical workflows, where accurate infection localization is essential for treatment planning and

Research Article

prognosis. CNN-based approaches thus represent practical solutions for enhancing diagnostic capabilities in chest X-ray analysis of COVID-19 as well as other respiratory conditions.

Multiple studies validate CNNs' efficacy in pneumonia detection, with models like CheXImageNet achieving 100% accuracy in classifying respiratory diseases from chest X-rays [5][6]. For infection segmentation, hybrid approaches combining GANs with CNNs have demonstrated high accuracy in distinguishing complex infection boundaries [7], while recurrent convolutional architectures enhance precision for subtle pathological changes [8]. Transfer learning further optimizes CNN performance even with limited labeled data, as demonstrated in studies of brain and lung disease classification [3] [9].

Hybrid models have shown remarkable versatility across medical imaging tasks. The LSTM-SNN approach addresses zero-shot learning challenges [10], while methodological frameworks from other domains emphasize architecture selection tailored to specific data types [11]. CNN-based hybrid models integrating architectures like InceptionV3 and VGG16 have proven effective in diverse applications from content moderation to pneumonia diagnosis [12][13], outperforming standard architectures in comparative studies.

Despite transformers' potential, CNNs maintain advantages in scenarios with constrained resources [14][15]. CNNs' cutting-edge accuracy in COVID-19 infection segmentation and pneumonia identification is demonstrated by models like as CheXNet and Inf-Net [16][17], with slice-level classification models further showcasing robust segmentation capabilities crucial for infection mapping [18]. While transformers offer promise for large, complex datasets, CNNs and transfer learning provide an efficient alternative, particularly for applications requiring precise segmentation and classification in medical imaging [2].

1.2 Contributions

This study offers a number of significant advancements in medical image analysis for identification of pneumonia and COVID-19:

- We developed dual-purpose CNN models, one for infection segmentation and another for pneumonia classification. Segmentation model had been trained to accurately localize infection regions in chest X-rays, and its encoder was later fine-tuned to improve classification performance by focusing on infection zones.
- We added dense blocks to the U-Net encoder's VGG16 convolution blocks. We combined DenseNet's efficient feature reuse capabilities with VGG16's robust feature extraction to improve COVID-19 and pneumonia classification accuracy.
- We employed transfer learning to modify CNN models that had already been trained to manage lack of medical imaging data. This showed CNNs' versatility in low-data settings and improved performance while lowering reliance on sizable, annotated datasets.
- We started by training only top layers and then gradually unfreezing pre-trained VGG16 layers as part of a gradual fine-tuning approach. This approach enabled efficient learning while mitigating overfitting and improving generalization with limited data.
- We enhanced model interpretability through Grad-CAM visualizations, providing greater transparency in the prediction process to support clinical decision-making.

FOUNDATIONAL DEEP LEARNING ARCHITECTURES FOR MEDICAL IMAGING

Field of medical imaging has benefited notably from DL applications, with numerous studies demonstrating effective pneumonia and COVID-19 diagnosis utilizing chest radiographs. Wang et al. [1] have pioneered DL models for COVID-19 identification utilizing chest CT images and X-ray, while semantic segmentation methods exemplified by U-Net architecture [19] have been widely implemented for identifying abnormalities in lung imaging [3][20] . Enhancing interpretability through attention mechanisms and Grad-CAM visualization [16][21], has been crucial in helping clinicians understand AI-driven diagnostic outcomes, with validation frameworks like COVID-Net [1] and ChestX-ray8 [22] providing benchmarks for performance assessment.

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Several CNN architectures have proven particularly effective in medical imaging applications. The VGG family of networks, including VGG-19 and VGG-16 developed by Zisserman and Simonyan [5][23][24], feature uniform architectures utilizing max-pooling layers after tiny 3x3 convolutional filters to allow for hierarchical feature learning from input photos. While computationally intensive, these architectures serve as foundational models influencing subsequent network designs. The InceptionV3 model by Szegedy et al. [24][25] employs "inception modules" performing parallel convolutions at different spatial scales, capturing diverse features at multiple abstraction levels while incorporating batch normalization and aggressive regularization to reduce overfitting.

ResNet-50 model presented by He et al. [26] addresses vanishing gradients through skip connections or "residual blocks" that enable training deeper networks without performance degradation. This 50-layer architecture allows gradients to flow directly through the network, facilitating complex feature learning while maintaining computational efficiency. CheXNet, discovered by Rajpurkar et al. [16], specifically targets pneumonia detection using a 121-layer DenseNet CNN that has shown precision surpassing radiologists when identifying thoracic pathologies from chest X-rays.

U-Net architecture [19] represents a fundamental component of medical image segmentation, with a symmetric expanding path that allows for accurate localization and a contracting path that records context. This method has shown remarkable effectiveness in lung tissue segmentation and pathology diagnosis with minimum training data, which makes it primarily useful in medical applications where annotated images are hard to come by. Multiple studies have extended this architecture [8][20]Implementing changes to improve efficiency and flexibility in response to novel medical imaging issues, especially for COVID-19 and pneumonia detection jobs.

METHODOLOGY

Methodology involved in this research is two-phase approach, as illustrated in Figure 1, for creating and utilizing DL models for analysis of lung infection. Employing transfer learning approaches, these phases concentrate on gradually training two interconnected models.

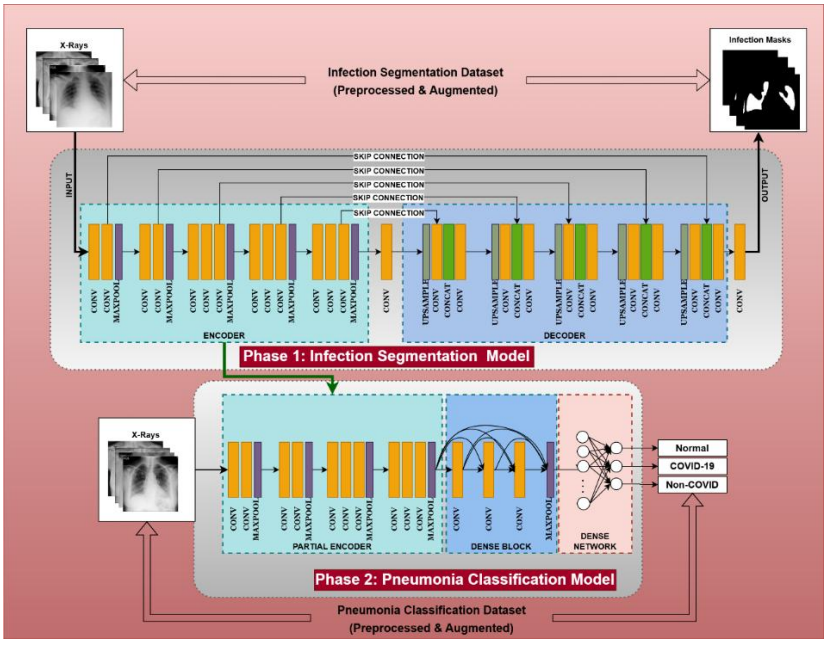

Figure 1: Two-Phase Training Approach for Pneumonia Diagnosis.

- Phase 1: Infection Segmentation Model Training - Using a VGG16 convolutional network as its encoder, we first trained an infection segmentation model based on U-Net architecture (Figure 2). To take advantage of pre-trained weights, VGG16 blocks were first frozen. They were then gradually unfrozen, beginning with deeper

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

layers and working backwards. This method enabled model to more accurately refine its feature extraction skills, adjusting to nuances of lung infection images.

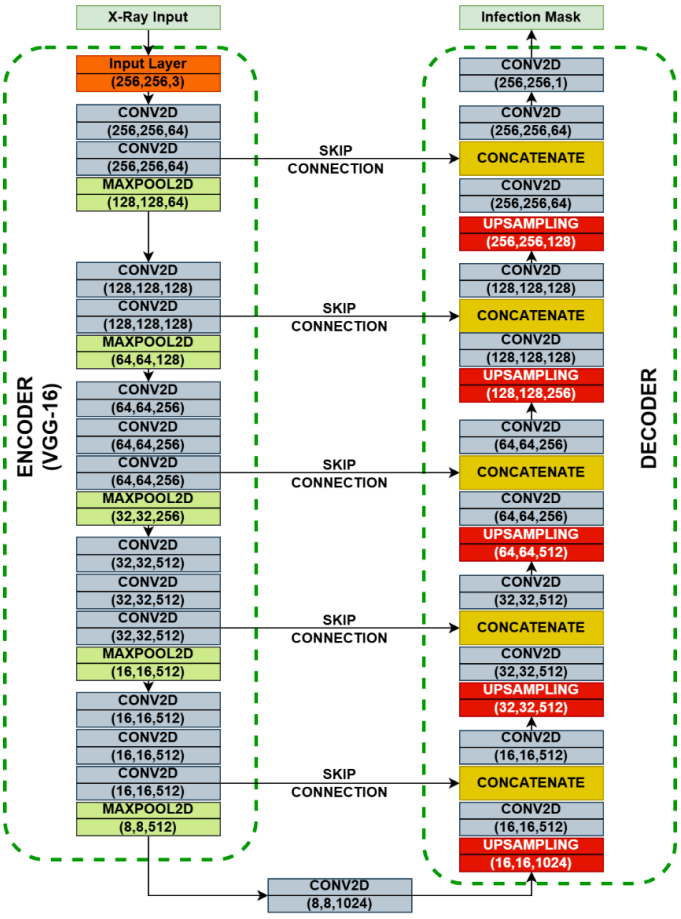

Figure 2: U-Net for Infection Segmentation.

- Phase 2: Pneumonia Classification Model Development - Using trained segmentation model's first four encoder blocks, then a dense block modeled like DenseNet121 (CheXNet), we created a classification model to divide X-ray images into three categories: Normal, COVID-19, and Non-COVID (Figure 3). By adding more dense blocks and fully linked layers to pre-trained encoder, this model improved its capacity to discriminate between specified categories. Like the segmentation model, this classification model gradually unfrozen its layers and optimized performance by phased fine-tuning.

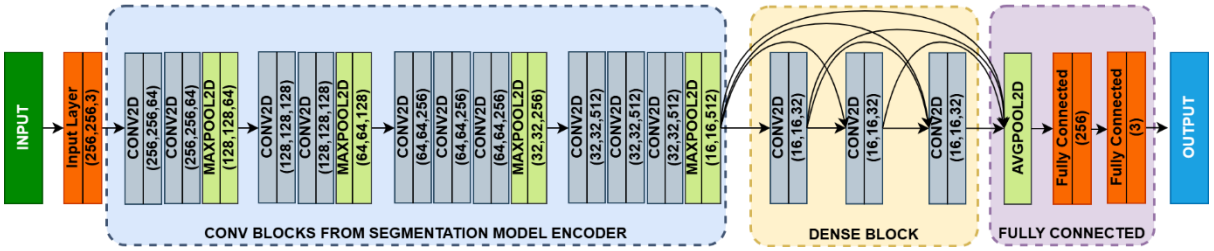

Figure 3: Pneumonia Classification Model.

Prediction Workflow Using Trained Models

Prediction workflow implements our trained deep learning models in a sequential diagnostic process, as illustrated in Figure 4. Pneumonia classification model first processes X-ray image and classifies it as Normal, COVID-19, or non-COVID. The process ends if it is categorized as Normal because no infection was discovered.

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

For images classified as non-COVID or COVID-19, workflow continues to the infection segmentation model, which generates a detailed mask highlighting infected lung regions. In order to precisely superimpose infected areas onto the lung regions, this mask is applied to original X-ray along with a lung segmentation mask. System then determines infection's proportion of lung area, which is a crucial indicator for determining its severity.

Final output provides clinicians with a comprehensive evaluation package: classification result, visual representation of infection superimposed on original X-ray, and quantified infection percentage. This integrated approach enables both qualitative and quantitative assessment of pneumonia cases, supporting more informed clinical decision-making.

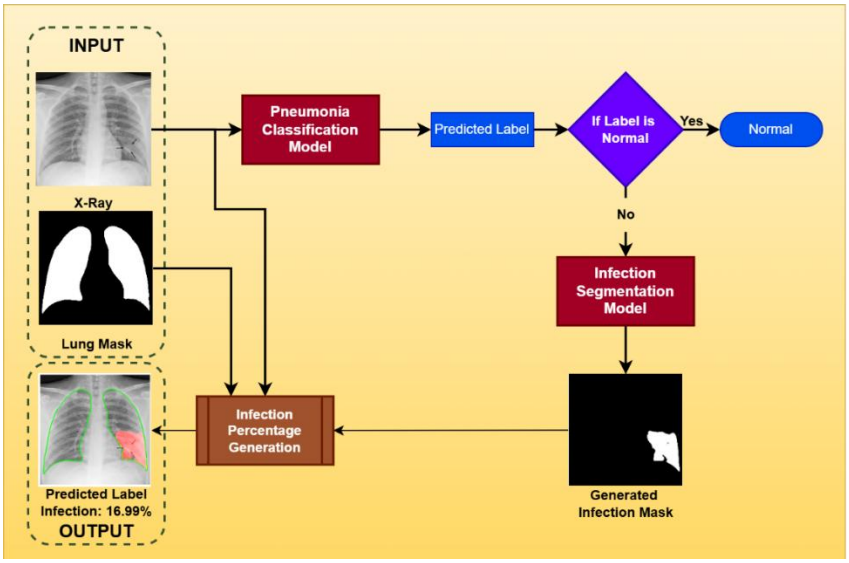

Figure 4: Workflow for making predictions using the trained models.

EXPERIMENTAL SETUP

Our experimental evaluation employed comprehensive metrics to assess the segmentation and classification models. To classify pneumonia, we used F1-Score (harmonic mean of Precision and Recall), Accuracy (ratio of correct predictions to total observations), Precision (positive predictive value), and Recall (sensitivity). Additionally, segmentation model was examined employing Dice Coefficient and Intersection over Union (IoU), which measure spatial overlap between ground truth annotations and projected infection masks and are essential for clinical relevance.

We trained our models on two complementary datasets: the QaTa-COV19 dataset [27] containing 119,316 chest X-ray images with 2,951 COVID-19 samples annotated with infection segmentation masks, and the COVID-19 Radiography Database [28] comprising 3,616 COVID-19 cases, 10,192 normal photos, and 7,357 other pneumonia cases. All images underwent preprocessing, including resizing to 256×256 pixels, normalization to zero mean and unit standard deviation, and augmentation through random rotations, zoom, and horizontal flipping to enhance model generalization.

The computational environment consisted of an AMD Ryzen 7 5800X CPU, 64GB RAM, and an NVIDIA RTX 3090 GPU with 24GB VRAM, operating on a Linux platform. This high-performance setup enabled efficient training of our deep learning models while accommodating the substantial computational demands of processing medical imaging data. The implementation framework facilitated model development and subsequent clinical deployment through the sequential prediction workflow illustrated in Figure 4.

RESULTS AND DISCUSSION

Our deep learning models demonstrated significant efficacy in infection segmentation and pneumonia classification tasks. Using QaTa-COV19 dataset and a U-Net architecture with a VGG16 encoder, infection segmentation model

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

demonstrated exceptional consistency in performance measures across 120 training epochs. Effective learning patterns and excellent generalization abilities have been demonstrated by model's gradual decline in loss values while maintaining constant accuracy and IoU scores, as demonstrated in Figure 5.

Visual assessment of model's segmentation performance is illustrated in Figure 6. It compares original X-ray images from test set with model's predictions and corresponding ground truth masks. These comparisons reveal model's precision in identifying and delineating infected regions within the lungs, with predicted infection boundaries closely aligning with expert annotations. This visual confirmation complements the quantitative metrics, demonstrating the model's practical utility for clinical infection assessment.

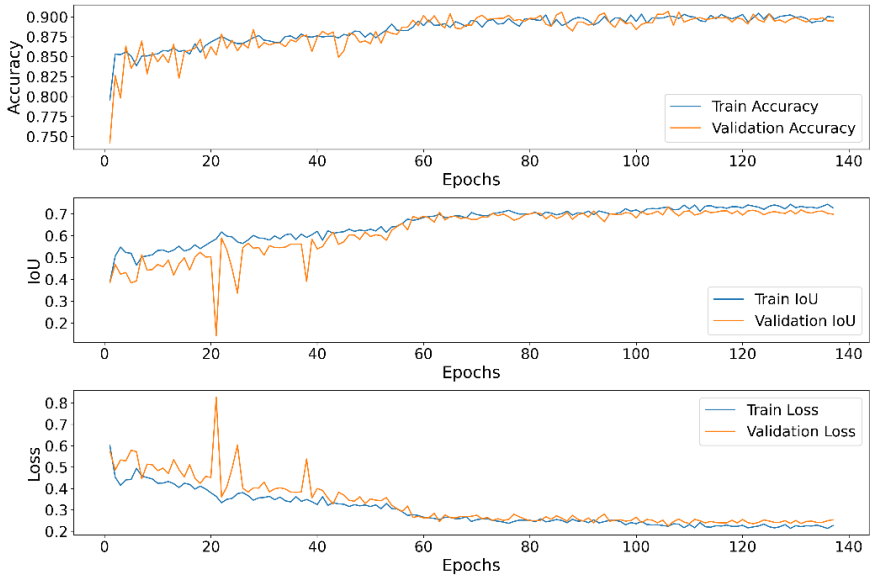

Figure 5: Training performance of the infection segmentation model.

Classification model for pneumonia performed exceptionally well in differentiating between normal, COVID-19, and non-COVID patients, with comprehensive results presented through confusion matrices and class-specific metrics. Notably, its interpretability through Grad-CAM visualizations enhances model's effectiveness, highlighting specific regions influencing classification decisions. This transparency feature provides clinicians valuable insights into model's reasoning process, facilitating greater trust and adoption in clinical settings. The combined workflow of both models enables accurate diagnosis and quantitative severity assessment through infection percentage calculation, representing a significant advancement in AI-assisted pneumonia management.

Pneumonia Classification Performance Analysis

During pneumonia classification phase, several CNNs were assessed to determine which model would best categorize chest X-rays into Normal, COVID-19, and non-COVID categories. We conducted a comprehensive analysis using accuracy, F1-score, and parameter count as key performance indicators, with results summarized in Table 1. Our proposed hybrid model, combining VGG16 convolutional blocks with DenseNet-inspired dense connections, demonstrated superior performance with exceptional accuracy while maintaining significantly lower parameter counts than standard architectures. Our model is especially well-suited for clinical deployment in resource-constrained situations, where processing speed and memory restrictions are crucial factors, because of its ideal balance among high diagnostic accuracy and computational efficiency.

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

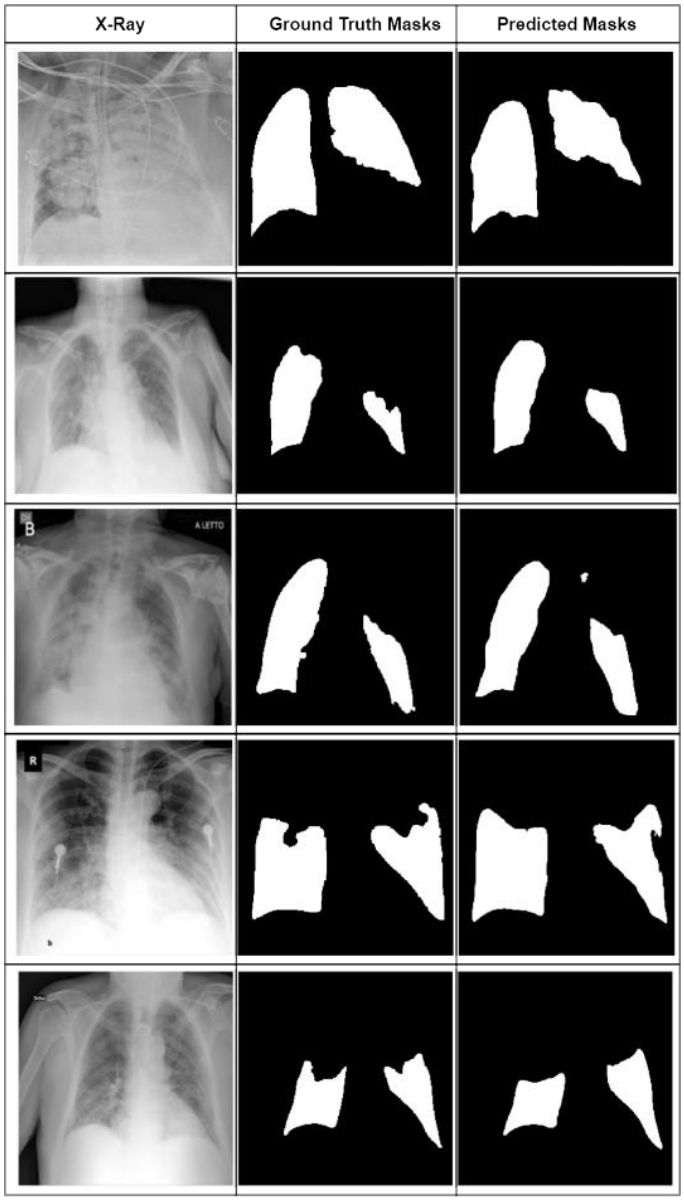

Figure 6: Lung Segmentation Accuracy Comparison.

Table 1: Comparison of CNN architectures for pneumonia classification.

Model	Accuracy (%)			F1-Score (%)			No. of Parameters (in millions)
	Training	Validation	Difference	Training	Validation	Difference	
ResNet50	97.74	91.39	6.35	97.72	91.10	6.62	~292
InceptionV3	99.78	97.95	1.83	99.78	97.86	1.92	~172
VGG16	99.66	97.95	1.71	99.67	97.98	1.69	~165
VGG19	99.32	97.89	1.43	99.31	97.85	1.46	~171
ChesXNet	99.68	98.22	1.46	99.68	98.21	1.47	~141
Proposed	99.46	98.71	0.75	99.46	98.71	0.75	~8

System Operation and Final Output

The integrated lung infection detection system leverages our models that are trained to comprehensively analyze chest X-rays, offering both qualitative visualization and quantitative assessment of infection severity. As demonstrated in Figure 7, the system generates detailed outputs that include lung boundary delineation (green

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

contours), infection region identification (red highlights), and precise infection percentage calculations. The visual results clearly distinguish between COVID-19 cases (top row), non-COVID pneumonia cases (middle row), and normal cases (bottom row), with infection percentages ranging from 0% in healthy lungs to over 87% in severe cases.

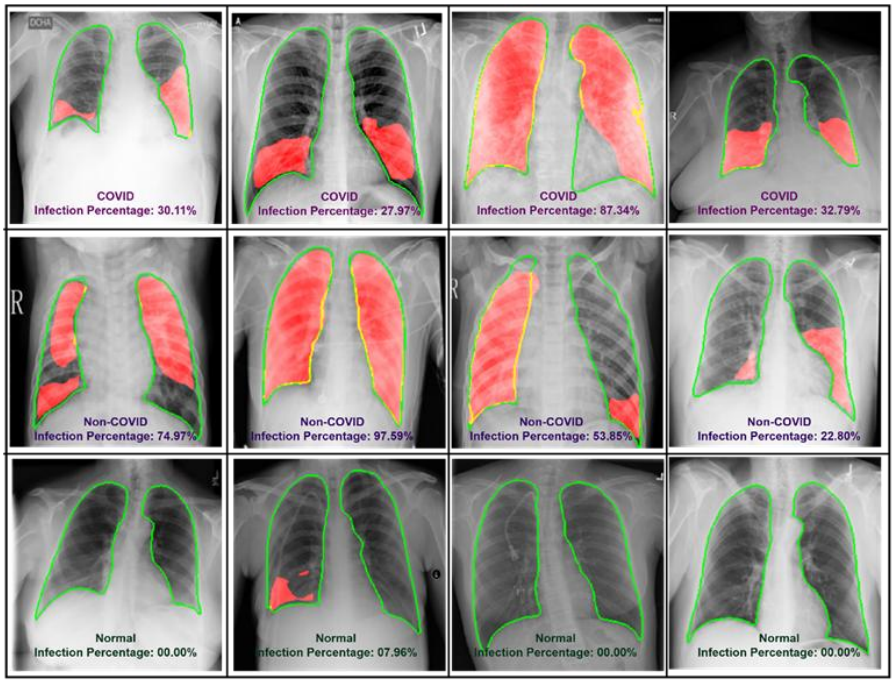

Figure 7: Pneumonia Severity Assessment with Infection Quantification.

This output format provides clinicians with an intuitive, at-a-glance assessment tool that enhances diagnostic capabilities in three critical ways: accurate classification of the pneumonia type, precise localization of affected lung regions, and quantitative severity metrics through infection percentage calculations. Notably, the system demonstrates high specificity by correctly identifying normal cases without generating false infection masks, while maintaining sensitivity in detecting varying degrees of infection across both COVID-19 and non-COVID pneumonia categories. This comprehensive visualization approach supports more informed clinical decision-making, potentially expediting treatment planning and improving patient results through early and correct diagnosis.

Model Explainability

A key component of our strategy is guaranteeing model interpretability and transparency, which is particularly essential in medical applications where clinical adoption of AI-driven decisions depends on a comprehension of reasons behind them. Figure 8 demonstrates how we utilized Grad-CAM to give visual justifications for classification choices made by our model. This technique generates heatmap visualizations highlighting regions of chest X-rays that most significantly influence model's diagnostic predictions.

Comparison between infection-highlighted X-rays (top row) and their corresponding Grad-CAM heatmaps (bottom row) demonstrates strong alignment among model's attention areas and clinically relevant pathological features. For COVID-19 cases, the heatmaps show intense activation over regions exhibiting characteristic COVID-19 opacities, while for non-COVID pneumonia, the model appropriately focuses on areas showing distinct non-COVID infection patterns. This visualization confirms that our model bases its classifications on medically relevant features rather than incidental image characteristics or artefacts.

Journal of Information Systems Engineering and Management
2025, 10(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

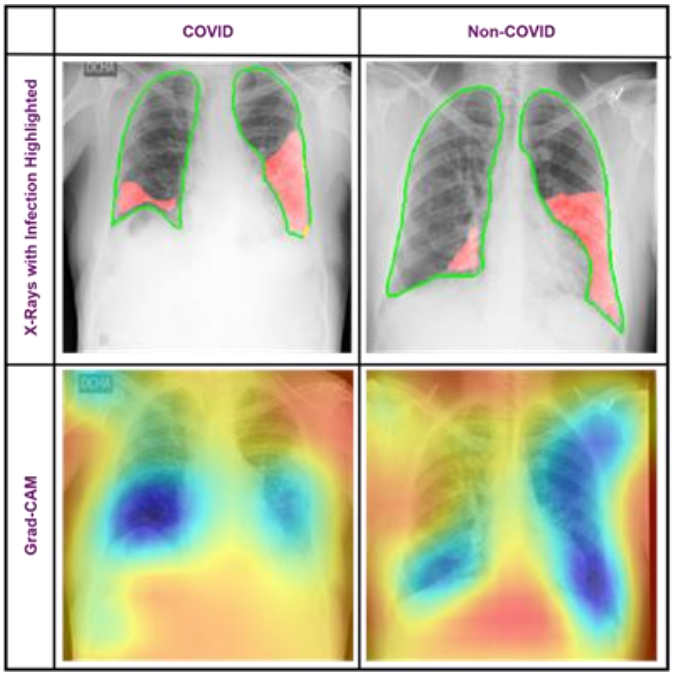

Figure 8: Model Interpretability through Grad-CAM Visualizations.

These explainability features address a common concern in medical AI applications by providing clinicians with insight into decision-making process of model. By demonstrating that the model attends to pathologically significant regions, we enhance trust in the system's outputs and provide an additional validation mechanism for clinical users. This transparency component represents an important step toward responsible AI deployment in healthcare settings, where interpretability is as critical as raw performance metrics in determining a system's practical utility.

CONCLUSION AND FUTURE WORK

In our study, we have shown potential of DL techniques for pneumonia diagnosis and severity assessment using chest X-rays, with particular emphasis on COVID-19 related cases. Our approach achieved high diagnostic accuracy and effective localization of affected regions by integrating CNNs for classification and U-Net-based architectures for infection segmentation. The combined framework facilitates timely and informed clinical decisions by offering categorical predictions and visual insights into infection severity. These results validate the utility of DL in medical imaging and highlight its promise in improving diagnostic workflows and enhancing patient care results.

Several future research directions can further enhance this work. Integrating multimodal clinical data—for instance, patient history, lab results, and demographics—with imaging could yield more comprehensive diagnostic models. Transfer learning may reduce training data requirements, while federated learning can ensure data privacy across institutions. Developing real-time diagnostic systems and embedding them into telemedicine platforms will improve accessibility, particularly in resource-constrained settings. Finally, ongoing model refinement and validation with new clinical data will be essential to maintain performance and adaptability to evolving healthcare demands.

REFERENCES

[1] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-ray: Hospital-Scale Chest X-ray Database and Benchmarks on Weakly Supervised Classification and Localization of Common Thorax Diseases," in Deep Learning and Convolutional Neural Networks for Medical Imaging and Clinical Informatics, L. Lu, X. Wang, G. Carneiro, and L. Yang, Eds. Cham: Springer International Publishing, 2019, pp. 369–392.

[2] F. Shamshad et al., "Transformers in Medical Imaging: A Survey." 2022, doi: https://doi.org/10.48550/arXiv.2201.09873.

Research Article

[3] A. W. Salehi et al., "A Study of CNN and Transfer Learning in Medical Imaging: Advantages, Challenges, Future Scope," Sustainability, vol. 15, no. 7, 2023, doi: 10.3390/su15075930.

[4] G. Tong, Y. Li, H. Chen, Q. Zhang, and H. Jiang, "Improved U-NET network for pulmonary nodules segmentation," Optik (Stuttg)., vol. 174, pp. 460–469, 2018, doi: https://doi.org/10.1016/j.ijleo.2018.08.086.

[5] S. Sharma and K. Guleria, "A Deep Learning based model for the Detection of Pneumonia from Chest X-Ray Images using VGG-16 and Neural Networks," Procedia Comput. Sci., vol. 218, pp. 357–366, 2023, doi: https://doi.org/10.1016/j.procs.2023.01.018.

[6] S. Shastri, I. Kansal, S. Kumar, K. Singh, R. Popli, and V. Mansotra, "CheXImageNet: a novel architecture for accurate classification of Covid-19 with chest x-ray digital images using deep convolutional neural networks," Health Technol. (Berl)., vol. 12, no. 1, pp. 193–204, Jan. 2022, doi: 10.1007/s12553-021-00630-x.

[7] A. Negi, A. N. J. Raj, R. Nersisson, Z. Zhuang, and M. Murugappan, "RDA-UNET-WGAN: An Accurate Breast Ultrasound Lesion Segmentation Using Wasserstein Generative Adversarial Networks," Arab. J. Sci. Eng., vol. 45, no. 8, pp. 6399–6410, Aug. 2020, doi: 10.1007/s13369-020-04480-z.

[8] M. Z. Alom, C. Yakopcic, T. M. Taha, and V. K. Asari, "Nuclei Segmentation with Recurrent Residual Convolutional Neural Networks based U-Net (R2U-Net)," in NAECON 2018 - IEEE National Aerospace and Electronics Conference, 2018, pp. 228–233, doi: 10.1109/NAECON.2018.8556686.

[9] A. Manickam, J. Jiang, Y. Zhou, A. Sagar, R. Soundrapandiyan, and R. Dinesh Jackson Samuel, "Automated pneumonia detection on chest X-ray images: A deep learning approach with different optimizers and transfer learning architectures," Measurement, vol. 184, p. 109953, 2021, doi: https://doi.org/10.1016/j.measurement.2021.109953.

[10] Y. Li, Z. Yang, H. Zheng, and X. Zhang, "A hybrid LSTM-SNN approach for robust multimodal zero-shot learning," Int. J. Complex. Appl. Sci. Technol., vol. 1, no. 2, pp. 195–209, 2024, doi: 10.1504/ijcast.2024.143881.

[11] V. Sammartino, "A comparative study of machine learning models for hate speech and stereotype detection in Italian texts," Int. J. Complex. Appl. Sci. Technol., vol. 1, no. 2, pp. 155–176, 2024, doi: 10.1504/ijcast.2024.143880.

[12] S. K. R. Mallidi, S. Bellapukonda, N. Gollapudi, K. Malaka, P. Sreeshanth, and S. S. H. Polisetti, "Deep Learning in Cartoon Moderation: Distinguishing Child-Friendly Content with CNN Architectures," in Cognitive Computing and Cyber Physical Systems, 2025, pp. 241–256.

[13] S. K. R. Mallidi, R. Vanapalli, D. S. P. Jella, T. Nimmakayala, D. Narni, and C. Attili, "Improving Pneumonia Diagnosis via Deep Learning: A Comprehensive Approach Incorporating CNN Classification," in Cognitive Computing and Cyber Physical Systems, 2025, pp. 206–219.

[14] S. Singh, M. Kumar, A. Kumar, B. K. Verma, K. Abhishek, and S. Selvarajan, "Efficient pneumonia detection using Vision Transformers on chest X-rays," Sci. Rep., vol. 14, no. 1, p. 2487, Jan. 2024, doi: 10.1038/s41598-024-52703-2.

[15] L. Gai, M. Xing, W. Chen, Y. Zhang, and X. Qiao, "Comparing CNN-based and transformer-based models for identifying lung cancer: which is more effective?," Multimed. Tools Appl., vol. 83, no. 20, pp. 59253–59269, Jun. 2024, doi: 10.1007/s11042-023-17644-4.

[16] P. Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning." 2017, [Online]. Available: https://arxiv.org/abs/1711.05225.

[17] D.-P. Fan et al., "Inf-Net: Automatic COVID-19 Lung Infection Segmentation From CT Images," IEEE Trans. Med. Imaging, vol. 39, no. 8, pp. 2626–2637, 2020, doi: 10.1109/TMI.2020.2996645.

[18] P. Afshar et al., "COVID-CT-MD, COVID-19 computed tomography scan dataset applicable in machine learning and deep learning," Sci. Data, vol. 8, no. 1, p. 121, Apr. 2021, doi: 10.1038/s41597-021-00900-3.

[19] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2015, 2015, pp. 234–241.

Research Article

[20] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation," in Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2016, 2016, pp. 424–432.

[21] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 618–626, doi: 10.1109/ICCV.2017.74.

[22] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases," in 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 3462–3471, doi: 10.1109/CVPR.2017.369.

[23] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition." 2015.

[24] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 2818–2826, doi: 10.1109/CVPR.2016.308.

[25] R. Jain, P. Nagrath, G. Kataria, V. Sirish Kaushik, and D. Jude Hemanth, "Pneumonia detection in chest X-ray images using convolutional neural networks and transfer learning," Measurement, vol. 165, p. 108046, 2020, doi: https://doi.org/10.1016/j.measurement.2020.108046.

[26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.

[27] A. Degerli et al., "COVID-19 infection map generation and detection from chest X-ray images," Heal. Inf. Sci. Syst., vol. 9, no. 1, p. 15, Apr. 2021, doi: 10.1007/s13755-021-00146-8.

[28] M. E. H. Chowdhury et al., "Can AI Help in Screening Viral and COVID-19 Pneumonia?," IEEE Access, vol. 8, pp. 132665–132676, 2020, doi: 10.1109/ACCESS.2020.3010287.