

Deep Transfer Learning with Customized Convolutional Neural Networks for Radiological Image Classification

^{*1}Dr. Sushilkumar Siddhartha Salve, ²Mr. Satish Subhash Asane, ³Dr. Prashant Rajaram Dike, ⁴Mr. Dhanajay Sudamrao Lamb, ⁵Mr. Sujit Ramesh Borey, ⁶Dr. Puja Padiya

^{*1}Assistant Professor, Sinhgad Institute of Technology, Lonavala, Maharashtra- 410401 , India
Email id: sushil.472@gmail.com Orcid ID: 0009-0008-8346-4017

²Assistant Professor, Dr. D. Y. Patil Technical Campus, Talegoan Dabhade, Varale, Pune, Maharashtra-410507, India Email id: satishsasane888@gmail.com Orcid ID: 0009-0009-5619-7950

³Assistant Professor, D.Y. Patil College of Engineering, Akurdi, Pune, Maharashtra-411044, India
Email id: prdike@dypcoeakurdi.ac.in Orcid ID: 0000-0002-3352-5642

⁴Lecturer, E&TC Dept., CSMSS College of Polytechnic, Chatrapati Sambhajinagar, Maharashtra-431011, India
Email id: dhanjay3@gmail.com Orcid ID: 0009-0009-3734-3254

⁵Research Scholar, ECE Department, Visvesvaraya National Institute of Technology (VNIT), Nagpur, Maharashtra-440010, India Email id: sujeetbore3@gmail.com Orcid ID: 0009-0009-0024-4877

⁶Associate Professor, Xavier Institute of Engineering, Mumbai, Maharashtra-400016, India
Email id: puja.padiya05@gmail.com Orcid ID: 0000-0002-3446-3340

Abstract

Radiological image interpretation is fundamental to modern clinical diagnosis, yet remains challenging due to high data dimensionality, complex anatomical variability, and the limited availability of labeled datasets. This study proposes a deep learning framework that integrates a customized convolutional neural network (CNN) with ResNet50- based transfer learning to improve diagnostic accuracy and robustness in radiology. The bespoke CNN is designed to capture hierarchical spatial representations from medical images, while ResNet50 leverages pretrained weights to accelerate convergence, preserve generic visual features, and enhance generalization to domain-specific radiological patterns. Extensive experiments conducted on multiple radiology datasets demonstrate that the proposed approach consistently outperforms CNNs trained from scratch, achieving accuracy exceeding 91% alongside notable improvements in Precision, Recall, F1-score, and Area Under the Curve (AUC). Gradient-weighted Class Activation Mapping (Grad-CAM) further reveals that the model focuses on clinically relevant regions such as lesions and abnormal tissue structures, thereby providing interpretable predictions and strengthening clinical trust. The results confirm that transfer learning significantly mitigates overfitting, stabilizes model training, and delivers reliable performance in limited-data scenarios commonly encountered in medical imaging. In addition, the proposed framework exhibits faster convergence and improved consistency across heterogeneous datasets, highlighting its robustness and adaptability. Overall, this work presents a scalable and explainable solution for automated radiological diagnosis, demonstrating how the integration of pretrained deep architectures with task-specific CNN design can advance AI-assisted clinical decision support and contribute toward more accurate, efficient, and reliable medical image analysis.

Keywords– Radiology, ResNet50, Gradient-weighted Class Activation Mapping, CNN

I. INTRODUCTION

Medical imaging has become a fundamental pillar of modern healthcare, supporting the detection, diagnosis, and long-term monitoring of a wide range of diseases, including neurological disorders, cardiovascular conditions, and cancer. Imaging modalities such as magnetic resonance imaging (MRI), computed tomography (CT), ultrasound, and X-ray provide detailed insights into anatomical structures and functional abnormalities, enabling clinicians to make informed and timely decisions [1]. Despite these advances, manual interpretation of radiological images remains a demanding and time-consuming task, particularly as imaging datasets continue to grow in size, resolution, and complexity. Diagnostic outcomes can vary depending on clinician experience, workload, and fatigue, leading to inter- and intra-observer variability. Previous studies have shown that psychological fatigue and subjective judgment can contribute to inconsistent or inaccurate diagnoses [2], highlighting the growing need for reliable automated systems that can assist radiologists and improve diagnostic consistency in clinical practice.

In recent years, deep learning has emerged as a powerful tool for medical image analysis by enabling models to learn meaningful representations directly from raw image data [3]. Among various deep learning architectures, convolutional neural networks (CNNs) have demonstrated strong performance in visual recognition tasks due to

their ability to extract hierarchical feature representations. These learned features range from low-level patterns such as edges and textures to higher-level abstractions that correspond to anatomical structures and pathological abnormalities [4]. Compared with traditional machine learning approaches that rely on handcrafted feature extraction, CNNs automatically derive discriminative features from pixel-level information, leading to improved performance in classification, segmentation, and detection tasks across diverse medical imaging applications.

Although end-to-end CNN training [5,6] shows considerable promise, several practical challenges limit its effectiveness in clinical environments. High-quality annotated medical datasets are often scarce because of privacy regulations, the cost and time required for expert labelling, and the low prevalence of certain diseases. Consequently, CNNs trained from scratch are susceptible to overfitting, achieving high accuracy on training data while generalizing poorly to unseen samples. Furthermore, deep architectures typically require substantial computational resources and careful hyperparameter tuning to ensure stable convergence and optimal performance [7]. These challenges are particularly pronounced in radiology, where data heterogeneity and class imbalance are common.

To address these limitations, transfer learning has gained widespread adoption in medical image analysis. In this approach, models pretrained on large-scale generic datasets are adapted to domain-specific tasks, allowing previously learned visual representations to be reused. Pretrained architectures such as ResNet50 capture universal low-level features that can be fine-tuned to learn radiology-specific patterns, thereby reducing training time, lowering data requirements, and improving generalization. Residual learning further facilitates stable optimization by alleviating vanishing-gradient issues, enabling deeper networks to be trained effectively. As a result, transfer learning has become a practical strategy for achieving reliable performance in medical imaging scenarios characterized by limited labeled data.

Motivated by these considerations, this study proposes a hybrid framework that integrates a customized CNN with ResNet50-based transfer learning to enhance radiological diagnostic accuracy [8]. The tailored CNN is designed to extract hierarchical spatial features from radiological images, enabling detailed characterization of anatomical structures and pathological indicators. ResNet50 provides pretrained representations and residual connections that support efficient learning and faster convergence. Domain adaptation is achieved by fine-tuning higher layers of ResNet50 on radiology-specific datasets while preserving generic visual features learned from large-scale data. In addition, dynamic data augmentation techniques—including rotation, flipping, and normalization—are employed to increase data diversity and improve model robustness and generalization [9].

The primary objectives of this study are to improve diagnostic accuracy, reduce overfitting, accelerate training convergence, and deliver interpretable results suitable for clinical use. To systematically evaluate the benefits of transfer learning, the performance of CNNs trained from scratch is compared with ResNet50-based fine-tuned models across multiple radiology datasets [10]. Furthermore, Gradient-weighted Class Activation Mapping (Grad-CAM) is incorporated to highlight regions of interest within medical images, enabling clinicians to visually assess model predictions and increasing transparency. This interpretability component is particularly important for clinical acceptance, as it allows radiologists to verify that decisions are based on physiologically meaningful regions. The contributions of this work are threefold. First, a hybrid CNN–ResNet50 architecture is developed to improve feature learning and classification accuracy in radiological image analysis. Second, extensive experimental evaluation demonstrates the advantages of transfer learning over training CNNs from scratch, especially in limited-data settings commonly encountered in medical imaging. Third, visualization techniques are integrated to enhance interpretability and support clinical validation. Overall, this study aims to advance reliable and efficient computer-assisted diagnostic systems that support radiologists, improve diagnostic consistency, and facilitate the practical integration of deep learning technologies into modern healthcare workflows. This work presents a hybrid diagnostic framework that integrates task-specific convolutional feature learning with residual transfer learning. Unlike conventional pipelines that employ pretrained networks solely as fixed feature extractors, the proposed approach selectively fine-tunes higher residual layers to adapt to radiology-specific patterns while preserving robust low-level representations. In addition, Grad-CAM-based visual explanations are incorporated to support interpretability of model predictions. This integrated design aims to jointly improve classification performance, generalization capability, and transparency within a unified learning framework suitable for limited-data medical imaging environments. The novelty of this work lies in the unified integration of task-specific CNN feature learning, ResNet50 residual transfer learning, and Grad-CAM interpretability within a single framework, enabling improved accuracy, convergence stability, and transparent decision support under limited radiological data conditions.

II. LITERATURE SURVEY

Recent advances in deep learning have substantially transformed medical image analysis, particularly within diagnostic imaging, where automated interpretation has become increasingly important for improving efficiency and reducing clinician workload. Jeba Prasanna Idas et al. [11] demonstrated that convolutional neural networks (CNNs) are highly effective for detecting, classifying, and segmenting abnormalities across diverse imaging modalities, including nuclear medicine, magnetic resonance imaging (MRI), and ultrasound. Their work highlights the strong capability of CNN-based architectures to model complex spatial dependencies and heterogeneous radiological patterns, confirming their suitability for analysing high-dimensional medical images.

Building upon this foundation, M. E. Naga Pavan et al. [12] examined the hierarchical feature-learning behaviour of CNNs and showed that these models automatically extract representations at multiple abstraction levels, ranging from low-level edges, textures, and intensity gradients to high-level semantic features corresponding to anatomical structures and pathological conditions. Their findings indicate that carefully optimised CNN architectures can achieve high diagnostic accuracy for conditions such as tumours and fractures, even when trained on relatively limited datasets. Furthermore, the study emphasises the importance of architectural design choices—including kernel size, network depth, and pooling strategies—in determining multi-class classification performance. These observations underline the need for carefully engineered CNN structures when addressing radiological imaging tasks. Despite these encouraging results, training CNNs from scratch remains challenging in medical imaging due to the scarcity of annotated data, privacy constraints, and the high cost of expert labelling. Subramanian R. R. et al. [13] addressed these challenges by highlighting transfer learning as a practical alternative. Their work demonstrated that pretrained networks, originally trained on large-scale natural image datasets, can be effectively adapted to medical imaging applications, significantly reducing training time while improving generalisation. In particular, residual architectures such as ResNet50 were shown to be advantageous because their residual connections mitigate vanishing-gradient issues and enable stable optimisation of deep networks. By freezing lower layers and fine-tuning higher layers, pretrained models can preserve generic visual representations while learning radiology-specific features, leading to improved accuracy, reduced overfitting, and faster convergence. More recently, hybrid modelling approaches that integrate custom CNN architectures with pretrained networks have gained attention. A. R. Singh et al. [14] reported that such hybrid frameworks consistently outperform standalone CNNs trained from scratch, especially in scenarios involving small or imbalanced datasets. In these systems, customised CNN components are responsible for extracting task-specific radiological features, while pretrained networks provide robust low-level representations and accelerate convergence. Their study also highlighted the role of advanced data augmentation strategies—including rotation, flipping, pixel shifting, scaling, and normalisation—in enhancing model robustness and generalisation.

Table 1: Summary of Literature Survey

Ref.	Authors	Method / Architecture	Dataset / Modality	Key Contributions	Limitations
[11]	Jeba Prasanna Idas et al.	CNN-based detection and segmentation	Nuclear medicine, MRI, Ultrasound	Demonstrated effectiveness of CNNs for abnormality detection and segmentation	Limited discussion on generalization and interpretability
[12]	M. E. Naga Pavan et al.	Optimized CNN architectures	Radiological images (tumors, fractures)	Showed hierarchical feature learning and importance of architectural parameters	Sensitive to architecture design and dataset size
[13]	Subramanian R. R. et al.	Transfer learning using ResNet50	Medical imaging datasets	Reduced overfitting and faster convergence using pretrained residual networks	Limited exploration of hybrid CNN structures
[14]	A. R. Singh et al.	Hybrid CNN + pretrained networks with augmentation	Small and imbalanced medical datasets	Hybrid models outperform CNNs trained from scratch; augmentation improves robustness	Interpretability not explicitly addressed
[15]	Thilagaraj M. et al.	Grad-CAM explainability	Radiological images	Improved clinician trust by highlighting pathological regions	Explainability not combined with hybrid deep models
[16]	A. D. S. et al.	CNN + Transfer learning + Explainable AI	Medical imaging	Suggested unified direction using CNNs, TL, augmentation, and explainability	No detailed unified experimental framework

These techniques are particularly valuable in medical imaging, where class imbalance and limited sample sizes are common. While predictive accuracy remains a primary objective, interpretability has emerged as a critical requirement for clinical adoption. Thilagaraj M. et al. [15] emphasised that visualisation techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM) play a crucial role in building clinician trust by highlighting diagnostically relevant regions within medical images. Such explainability approaches allow physicians to verify that model predictions are driven by physiologically meaningful areas, including lesions, tumours, and haemorrhages, thereby supporting transparent and accountable decision-making. Interpretability is especially important in radiology, where clinical acceptance depends not only on performance metrics but also on the ability to understand and validate automated predictions. Synthesising these developments, A. D. S. et al. [16] concluded

that the integration of efficient CNN architectures, transfer learning strategies, hybrid modelling frameworks, data augmentation, and explainable techniques represents a promising direction for advanced medical image analysis. Their work highlights the importance of combining accuracy, robustness, and transparency to facilitate real-world clinical deployment. Although existing studies have demonstrated the effectiveness of CNNs and transfer learning individually, many approaches focus primarily on performance metrics without fully addressing generalisation in limited-data settings or incorporating interpretability as a core component of the framework. Table 1 shows the summary of the literature survey.

Motivated by these observations, the present study proposes a unified framework that integrates a customized CNN with ResNet50-based transfer learning to enhance radiological diagnostic accuracy while maintaining model interpretability. Unlike conventional approaches that rely solely on pretrained networks or standalone CNNs, the proposed method combines task-specific feature extraction with pretrained representations to improve robustness across heterogeneous datasets. Furthermore, the incorporation of Grad-CAM provides visual explanations of model predictions, addressing an important gap in clinical applicability. By systematically comparing CNNs trained from scratch with transfer learning-based models across multiple radiology datasets, this work aims to demonstrate the practical advantages of hybrid deep learning frameworks in terms of accuracy, convergence, generalization, and interpretability, thereby contributing toward reliable and clinically meaningful AI-assisted radiological diagnosis.

III. PROPOSED MODEL

Effective application of deep learning techniques in radiology requires carefully designed network architectures, robust training strategies, and optimized hyperparameters to achieve high diagnostic accuracy and reliable generalization. Medical imaging data are typically high-dimensional, heterogeneous, and limited in annotated samples, making architectural design and training configuration critical to overall system performance. Accordingly, this section presents a detailed description of the proposed framework, which integrates a customized convolutional neural network (CNN) with ResNet50.

based transfer learning. The customized CNN is employed to learn hierarchical spatial representations from radiological images, enabling effective extraction of both low-level visual features and high-level pathological patterns. In parallel, ResNet50 utilizes pretrained weights and residual connections to facilitate efficient gradient propagation, accelerate convergence, and improve generalization, particularly under limited-data conditions.

This section further discusses the selection of key architectural parameters, including kernel sizes, network depth, pooling operations, activation functions, and regularization mechanisms. Implementation aspects, such as image preprocessing, data augmentation, optimiser configuration, learning rate selection, batch size determination, and early stopping criteria, are also described. These parameters are systematically tuned to balance computational efficiency with predictive performance while mitigating overfitting. The overall system architecture and training workflow are illustrated in Fig. 1, providing an overview of the interaction between the customized CNN, the transfer learning backbone, and the final classification layers.




Fig.1. Proposed Framework

A. Convolutional Neural Networks (CNN)

The proposed convolutional neural network (CNN) is developed to automatically classify radiological images by learning hierarchical spatial patterns directly from data. The architecture consists of five main parts: an input layer, stacked convolutional layers, pooling layers, a fully connected layer, and a final classification layer. All images are

resized to $224 \times 224 \times 3$ before training, ensuring uniformity and efficient batch processing. This resolution preserves clinically relevant details—such as nodules, lesions, and hemorrhagic regions—while maintaining computational efficiency. The three-channel format allows both RGB images and grayscale scans (replicated across channels) to be processed seamlessly.

Feature extraction is performed through three sequential convolutional blocks with 32, 64, and 128 filters (3×3), each using ReLU activation. The first block captures fundamental elements like edges and intensity variations. The second models more complex textures and intermediate anatomical structures, while the third learns higher-level semantic patterns associated with disease regions. This gradual increase in depth enables multiscale learning, which is essential in radiology where abnormalities appear at varying spatial resolutions.

Each convolutional stage is followed by 2×2 max-pooling with stride 2 to reduce spatial dimensions, lower computation, and introduce translation invariance while preserving salient responses. Zero-padding maintains feature map size across layers. Batch normalization may be applied to stabilize gradients, reduce internal covariate shift, and accelerate convergence.

After convolution and pooling, feature maps are flattened and passed to a dense layer with 512 neurons and ReLU activation, converting spatial representations into a compact vector for classification. Dropout with a rate of 0.5 is used during training to reduce overfitting by encouraging more distributed and resilient feature learning. The final softmax layer outputs class probabilities according to the number of diagnostic categories.

Weights are updated through backpropagation using adaptive optimization methods, ensuring efficient and stable convergence. The architecture is intentionally balanced to provide strong representational capacity without excessive computational burden, making it practical for real-world use.

B. Using ResNet50 for Transfer Learning

To enhance diagnostic accuracy while lowering training time and computational cost, the framework adopts ResNet50 pretrained on ImageNet as its backbone. This architecture is chosen for its strong representation learning ability and residual connections, which address the vanishing-gradient issue common in deep networks. These skip connections ensure smooth gradient flow and stable optimization, enabling deeper structures to train effectively. The pretrained weights offer a solid starting point by capturing fundamental visual patterns—such as edges, contours, and textures—which are later refined to recognize radiology-specific features.

To balance knowledge retention and domain adaptation, the early layers are frozen to preserve general visual representations, while the upper layers are fine-tuned using radiological images. This targeted adjustment allows the network to learn clinically relevant patterns, including lesions, nodules, and abnormal tissue regions, without discarding valuable pretrained information. The original 1000-class output layer is replaced with a softmax classifier tailored to the number of diagnostic categories, ensuring alignment with the medical task.

This transfer learning strategy is particularly advantageous when labeled data are scarce. Starting from pretrained parameters rather than random initialization accelerates convergence, reduces the number of required epochs, and lowers the risk of overfitting. The result is improved generalization and more stable learning compared to CNNs developed entirely from scratch.

In addition, leveraging ResNet50 decreases computational burden by minimizing extensive training while maintaining strong predictive performance. This efficiency makes the framework well suited for clinical settings where resources and time are often constrained. Overall, integrating ResNet50-based transfer learning provides a dependable and adaptable foundation for automated radiological diagnosis across varied datasets. The customized convolutional network is designed to progressively learn hierarchical spatial features from radiological images. Early convolutional layers capture basic visual patterns such as edges and textures, while deeper layers encode higher-level representations associated with pathological structures. Integration with ResNet50 introduces residual connections that facilitate stable gradient propagation and deeper feature abstraction. Lower residual blocks are retained to preserve transferable visual descriptors, whereas upper layers are fine-tuned to adapt to domain-specific characteristics. This hybrid strategy balances representation stability with task specialization, leading to improved convergence behavior and enhanced generalization.

C. Setting up training environment

The training setup of the proposed framework is structured to promote stable optimization and consistent performance across diverse radiological datasets. A composite overlap-based loss function is used to penalize misclassifications across multiple categories, improving class separation and helping to manage class imbalance—an issue frequently encountered in medical imaging. By encouraging closer agreement between predictions and ground truth, this approach supports more reliable and confident outcomes.

Parameter updates are performed using the Adam optimizer with an initial learning rate of 0.0001. Adam is chosen for its adaptive mechanism, which adjusts learning rates individually using gradient moment estimates. This leads to faster convergence and greater stability than traditional stochastic gradient descent, especially when working with deep networks and limited data. A batch size of 32 balances GPU memory efficiency with stable gradient estimation. Training runs for up to 50 epochs, with early stopping applied to monitor validation loss and halt learning if no improvement is observed over five consecutive epochs. This prevents overfitting while preserving generalization.

The preprocessed dataset is loaded via PyTorch DataLoader or TensorFlow's tf.data pipeline for efficient batch handling and parallel processing. During training, real-time augmentation—such as random rotations ($\pm 15^\circ$), horizontal and vertical flips, pixel shifts, and intensity normalization—enhances robustness and reduces sensitivity to image orientation or contrast variations. Shuffling ensures that each mini-batch contains diverse samples, minimizing bias.

For CNNs trained from scratch, weights are initialized using standard strategies to support stable learning. Each mini-batch passes through convolutional, pooling, and dense layers, followed by categorical cross-entropy loss computation and backpropagation with Adam-based updates. Validation accuracy and loss are assessed after every epoch to track generalization.

In the ResNet50 fine-tuning strategy, pretrained ImageNet weights provide initialization. The first 40 layers remain frozen to retain general visual representations, while higher layers are adapted to capture radiology-specific features. The original classifier is replaced with a softmax layer matching the number of target classes. Fine-tuning typically converges within 20–30 epochs, enabling efficient domain adaptation without compromising pretrained knowledge. Training progress is visualized using TensorBoard or Matplotlib, and checkpoints are saved according to peak validation accuracy. The proposed framework was evaluated on multiple radiological datasets comprising chest X-ray and brain MRI images. Subject-wise data partitioning was adopted to avoid information leakage, with approximately 70% of samples used for training, 15% for validation, and 15% for testing. All images were resized to 224×224 pixels and normalized prior to training. Class imbalance was addressed through weighted learning and data augmentation. Augmentation strategies included controlled rotations, geometric shifts, intensity normalization, and horizontal flipping to improve robustness while preserving anatomical consistency. The proposed framework was evaluated on three publicly available radiological datasets: NIH ChestX-ray14 (Dataset A), Brain MRI Tumor Dataset (Dataset B), and a combined radiology dataset comprising chest X-ray and brain MRI images (Dataset C). All datasets were split into 70% training, 15% validation, and 15% testing using subject-wise partitioning to avoid data leakage. Images were resized to 224×224 pixels and normalized prior to training, while data augmentation was applied to mitigate class imbalance.

IV. RESULTS AND DISCUSSIONS

Model performance was evaluated using standard metrics including Accuracy, Precision, Recall, F1-score, and AUC. To reduce sensitivity to random initialization and dataset splitting, experiments were repeated across multiple runs, and average performance values were reported. This evaluation strategy provides a more reliable estimate of generalization capability across heterogeneous radiological datasets. Model training is conducted on GPU hardware (e.g., NVIDIA Tesla V100/A100) to accelerate computational efficiency. For large batch processing, advanced memory optimization techniques such as gradient accumulation and mixed-precision training are employed to maximize GPU utilization while maintaining numerical stability. To ensure experimental reproducibility, random seeds are fixed across NumPy, PyTorch/TensorFlow, and CUDA environments.

Hyperparameters—including learning rate, batch size, and dropout probability—are optimized using grid search or Bayesian optimization to identify optimal configurations. These strategies enhance model robustness, reproducibility, and suitability for clinical deployment.

Empirical results demonstrate that CNN models trained from scratch typically achieve approximately 85% accuracy on limited datasets. In contrast, the integration of ResNet50-based transfer learning improves performance to nearly 92%, while simultaneously reducing overfitting and accelerating convergence. Furthermore, Grad-CAM visualizations highlight diagnostically relevant regions within medical images, improving interpretability and increasing clinician confidence in model predictions. Grad-CAM visualizations were used to examine spatial increasing patterns produced by the model. Activation maps consistently highlighted clinically relevant regions such as lesions and abnormal tissue areas, indicating that predictions are guided by meaningful anatomical features rather than background artifacts. This qualitative alignment between network attention and pathological regions enhances interpretability and supports the potential clinical applicability of the proposed framework.

Overall, the combination of carefully designed architectures, optimized training strategies, and transfer learning yields reliable and accurate diagnostic outcomes, supporting the practical adoption of deep learning frameworks in radiological decision support systems.

A. Configuring the Experiment

To improve diagnostic performance in medical imaging, this study adopts a hybrid deep learning approach that combines a customized CNN with ResNet50-based transfer learning. The CNN automatically learns hierarchical spatial features from radiological images, which are resized to $224 \times 224 \times 3$ to accommodate both RGB images and grayscale scans replicated across three channels, ensuring compatibility across imaging modalities.

The network includes three convolutional blocks with 32, 64, and 128 filters, each followed by 2×2 max-pooling to reduce spatial dimensions, lower computational load, and enhance translation invariance. ReLU activation is applied throughout, with optional batch normalization to stabilize gradients and speed up convergence. A fully connected layer of 512 neurons converts extracted feature maps into a compact representation for classification,

followed by dropout (0.5) to reduce overfitting. The final softmax layer produces class probabilities aligned with the target categories.

To further enhance accuracy and shorten training time, transfer learning is implemented using ResNet50, a 50-layer residual network pretrained on ImageNet. It skips connection address vanishing-gradient issues and enable deeper, more stable learning. Early layers are frozen to retain general visual features, while upper layers are fine-tuned to capture domain-specific radiological patterns. The original 1000-class output layer is replaced with a softmax classifier tailored to the medical dataset, allowing effective adaptation to healthcare imaging tasks.

Training settings are carefully configured. Categorical cross-entropy is used for multi-class prediction, optimized with Adam at a learning rate of 0.0001. A batch size of 32 balances memory efficiency and stable gradient updates. Early stopping monitors validation loss and halts training if no improvement occurs for five consecutive epochs. Models trained from scratch typically run up to 50 epochs, whereas ResNet50 fine-tuning converges within 20–30 epochs due to pretrained initialization.

Data preprocessing is managed through PyTorch DataLoader or TensorFlow tf.data pipelines. On-the-fly augmentation—such as rotations, flipping, pixel shifts, and normalization—improves robustness and reduces reliance on fixed image orientations. Backpropagation updates model weights, while validation metrics are tracked after each epoch. Progress is visualized using TensorBoard or Matplotlib, and checkpoints are saved based on peak validation performance.

All experiments are executed on GPU platforms (e.g., NVIDIA Tesla V100 or A100) to accelerate computation. Techniques like gradient accumulation and mixed-precision training support efficient memory use. Reproducibility is ensured by fixing random seeds across NumPy, PyTorch/TensorFlow, and CUDA. Hyperparameters—including learning rate, batch size, and dropout—are tuned through grid search or Bayesian optimization. Altogether, this framework delivers a reliable, reproducible, and clinically relevant system for high-accuracy radiological diagnosis.

B. Analyzing the Results

The proposed framework was evaluated on multiple radiology datasets using standard performance metrics, including Accuracy, Precision, Recall, F1-score, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). To assess the effectiveness of transfer learning, the performance of the customized CNN trained from scratch was compared with that of the ResNet50-based fine-tuned model. This comparative analysis highlights the advantages of leveraging pretrained architectures for medical image classification, demonstrating improved accuracy, enhanced generalization, and reduced overfitting when adapting pretrained models to radiological datasets. The proposed framework was compared with commonly used deep learning architectures, including VGG16, DenseNet121, EfficientNetB0, and MobileNetV2, under consistent preprocessing and optimization settings. While lightweight models offered faster inference, they demonstrated reduced sensitivity to subtle pathological features. In contrast, the hybrid CNN-ResNet50 architecture achieved more consistent discrimination across datasets, highlighting the benefit of combining residual transfer learning with customized convolutional refinement.

Table 2. Performance Metrics

Dataset	Model	Accuracy	Precision	Recall	F1-Score	AUC
Dataset A	CNN	85.2	83.5	84.1	83.8	88.5
	ResNet50	91.8	90.5	91.2	90.8	94.3
Dataset B	CNN	84.0	82.7	83.2	82.9	87.6
	ResNet50	92.1	91.0	91.5	91.2	94.7
Dataset C	CNN	83.5	82.0	82.6	82.3	87.0
	ResNet50	91.5	90.2	90.8	90.5	94.0

Compared with standard architectures such as VGG16, DenseNet121, EfficientNetB0, and MobileNetV2 under identical preprocessing, the proposed hybrid CNN-ResNet50 model demonstrated superior consistency in detecting subtle pathological features.. To understand the impact of individual components, qualitative comparisons were performed by examining model behavior with and without transfer learning and data augmentation. Results indicate that transfer learning improves convergence stability, while customized convolutional layers enhance sensitivity to localized pathological patterns. Data augmentation further contributes to reducing overfitting. These observations suggest that performance gains arise from the combined influence of architectural design, pretrained representations, and data enrichment strategies. The quantitative results summarized in Table 2 demonstrate that the CNN trained from scratch achieves reasonable classification performance, with accuracy values ranging between approximately 83% and 85% across the evaluated datasets. However, a consistent and substantial improvement is observed when employing ResNet50-based transfer learning, with overall accuracy exceeding 91% in all cases. This performance gain can be attributed to the utilization of pretrained weights, which preserve generic low-level visual representations—such as edges and

textures—while allowing higher network layers to be fine-tuned for domain-specific radiological patterns.

In addition to accuracy, the Precision, Recall, and F1-score metrics exhibit comparable improvements, indicating that the transfer learning model not only increases correct predictions but also enhances sensitivity and classification balance across classes. These results suggest reduced false positives and false negatives, which is particularly important in clinical diagnostic scenarios. Moreover, AUC values surpassing 94% highlight the strong discriminative capability of the proposed framework, confirming its effectiveness in separating pathological from non-pathological cases.

Overall, these findings validate the superiority of transfer learning over training CNNs from scratch, especially in medical imaging environments characterized by limited labeled data. The ResNet50-based approach demonstrates improved generalization, reduced overfitting, and faster convergence, making it a reliable and clinically viable solution for automated radiological diagnosis.

C. Loss versus Epoch

The training process was monitored by visualizing loss curves across epochs for both training and validation datasets. Figure 2 illustrates the convergence behavior of the proposed models.




Fig 2. Training and Validation Loss vs Epoch (Dataset A)




Fig 3. Training and Validation Loss vs Epoch (Dataset B)




Fig 4. Training and Validation Loss vs Epoch (Dataset C)

Figures 2–4 present the training and validation loss curves for both the CNN trained from scratch and the ResNet50-based transfer learning model. In both cases, loss decreases steadily over epochs, confirming effective learning. However, their convergence patterns differ notably. The baseline CNN converges more slowly, shows higher final validation loss, and displays noticeable oscillations, suggesting instability and sensitivity to limited radiological data. These fluctuations point to weaker generalization and a greater risk of overfitting when the network relies solely on small datasets.

In contrast, the ResNet50-based framework converges faster and produces smoother loss curves, indicating more stable optimization and stronger generalization. This improved behavior stems from residual connections that support efficient gradient flow and pretrained weights that provide a solid initialization, helping the model avoid poor local minima. Early stopping further limits overfitting by terminating training once validation performance plateaus, preserving transferable representations for unseen cases.

Performance across datasets reinforces these findings. On Dataset A, ResNet50 achieves 91.8% accuracy, effectively identifying lesions and pathological features. Grad-CAM visualizations confirm that predictions are guided by clinically relevant regions, such as nodules and hemorrhagic areas, enhancing interpretability and diagnostic confidence. On Dataset B, the model maintains strong performance despite distribution shifts, whereas the scratch-trained CNN shows a slight accuracy drop, reflecting its vulnerability to dataset heterogeneity. Even in Dataset C, which contains fewer samples, ResNet50 sustains accuracy above 91%, while the baseline CNN exhibits clear signs of overfitting.

Overall, transfer learning boosts classification accuracy from approximately 83–85% with the baseline CNN to over 91%, alongside consistent gains in Precision, Recall, F1-score, and AUC values exceeding 94%. These improvements arise from ResNet50's ability to retain general low-level visual features—such as edges, contours, and textures—while adapting higher layers to capture radiology-specific patterns, including lesions and bleeding regions. Continuous augmentation strategies, including rotations, flipping, pixel shifts, and normalization, further strengthen robustness by exposing the network to diverse variations. Residual learning, dropout, and early stopping collectively enhance optimization efficiency and reduce overfitting.

Beyond numerical metrics, interpretability stands out as a key advantage. Grad-CAM maps offer transparent insight into the decision-making process, allowing radiologists to verify that predictions are grounded in meaningful anatomical features. This transparency supports clinician trust and facilitates integration into real-world diagnostic workflows. High AUC and F1-scores, rapid convergence, and manageable computational demands make the framework suitable for time-sensitive clinical settings, where attention maps can aid collaborative human-AI decision-making.

Although the results are encouraging, validation across additional imaging modalities such as MRI and PET is necessary to confirm broader applicability. Addressing class imbalance for rare conditions and refining hyperparameters remain important considerations. Future work may explore multimodal image fusion, lightweight architectures for edge deployment, and semi-supervised or self-supervised approaches to reduce reliance on large, annotated datasets.

V. CONCLUSION

This study demonstrates that integrating customized convolutional modeling with residual transfer learning improves radiological image classification in limited-data scenarios. The proposed framework achieves consistent performance gains across multiple datasets, supported by improved convergence behavior and enhanced interpretability through Grad-CAM visualizations. By combining transfer learning with task-specific feature extraction, the approach offers a practical and scalable solution for AI-assisted radiological diagnosis. Future work will focus on extending this framework to additional imaging modalities and exploring semi-supervised learning

strategies to further reduce dependence on annotated data. Although the results are promising, further validation on larger multi-institutional datasets is required to confirm generalizability across diverse clinical settings.

REFERENCES

- [1] Senthil Pandi, S., Kumaragurubaran, T., Jaeyalakshmi, M., & Sheriff, M. A. (2024). A Comparative Analysis on the Prediction of Heart Failure using Machine Learning Algorithms. In *Proceedings - 2024 5th International Conference on Mobile Computing and Sustainable Informatics, ICMCSI 2024* (pp. 206–211). Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICMCSI61536.2024.0003>
- [2] Saravana Gokul, G., Senthil Pandi, S., Deepak Kumar, K., Sheriff, M. A., & Krishnamurthy, R. (2024). Optimizing Event Management Application through Data Analytics Driver and Dynamic Exploration. In *2nd IEEE International Conference on Advances in Information Technology, ICAIT 2024 - Proceedings*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICAIT61638.2024.10690407>
- [3] Jesuraj, Y. A., & Hemalatha, K. (2020). A prototype model of unmanned automatic level crossing system using piezoelectric sensor. *Microprocessors and Microsystems*, 79. <https://doi.org/10.1016/j.micpro.2020.103265>
- [4] Mohamed Yacin Sikkandar, K. Hemalatha, M. Subashree, S. Srinivasan, Seifedine Kadry, Jungeun Kim, Keejun Han, Leveraging Multimodal Ensemble Fusion-Based Deep Learning for COVID-19 on Chest Radiographs, *Computer Systems Science and Engineering*, 47(1),2023, 873-889, <https://doi.org/10.32604/csse.2023.035730>.
- [5] Afrah Fathima, S., Sheriff, M. A., Vyshnavi, A. R., & Jananee, V. (2023). INTEGRATING FUZZY LOGIC WITH SUPPLY CHAIN MANAGEMENT. In *2023 IEEE International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering, RMKMATE 2023*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/RMKMATE59243.2023.10368961>
- [6] Senthil Pandi, S., Kumar, P., Nandha Kumar, P., & Karthick, R. (2025). Machine Unlearning for Grid SearchCV. In *2025 International Conference on Computing Technologies and Data Communication, ICCTDC 2025*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICCTDC64446.2025.11158784>
- [7] Salve, S. S., & Narote, S. P. (2024). Performance evaluation of efficient segmentation and classification based iris recognition using sheaf attention network. *Journal of Visual Communication and Image Representation*, 103. <https://doi.org/10.1016/j.jvcir.2024.104262>
- [8] Senthil Pandi, S., Bala Subramanian, C., Praveen, S. D., & Ezhil Dyana, V. M. (2024). Improving Skin Lesion Diagnosis Employing Attention-Based LSTM Networks. In the *2024 International Conference on Integration of Emerging Technologies for the Digital World, ICIETDW 2024*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICIETDW61607.2024.10939196>
- [9] M. Thilagaraj, N. Arunkumar, S. Ramkumar, S. Hariharasitaraman, Electrooculogram signal identification for elderly disabled using Elman network, *Microprocessors and Microsystems*, Volume 82, 2021, 103811, ISSN 0141-9331, <https://doi.org/10.1016/j.micpro.2020.103811>.
- [10] Arockia Jesuraj Yagappan, Hemalatha Karuppiiah, Maragatharajan Muthusamy, Sathish Kumar Kannaiah, Optimizing lung cancer classification using transformer and Gooseneck Barnacle Optimization, *Expert Systems with Applications*, Volume 282, 2025, 127413, ISSN 0957-4174, <https://doi.org/10.1016/j.eswa.2025.127413>.
- [11] Jeba Prasanna Idas, S., Hemalatha, K., Naveenkumar, J. et al. Recent trends on mammogram breast density analysis using deep learning models: neoteric review. *Artif Intell Rev* 58, 240 (2025). <https://doi.org/10.1007/s10462-025-11232-8>
- [12] Naga Pavan, M. E., Reddy, U. B. K., Ghantasala, C. M., Devi, M. M. Y., & Nesarani, A. (2023). Enhanced Skin Cancer Classification with AlexNet and Transfer Learning. In *International Conference on Sustainable Communication Networks and Application, ICSCNA 2023 - Proceedings* (pp. 1020–1024). Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICSCNA58489.2023.10370110>
- [13] S. S. Salve. (2024). Efficient Iris Recognition Using Hybrid Preprocessing and Sheaf Attention Network-Based Segmentation and Classification. *Educational Administration: Theory and Practice*, 30(1), 8319–8329. <https://doi.org/10.53555/kuey.v30i1.11437>
- [14] Singh, A. R., Athisayamani, S., Hariharasitaraman, H. S., Karim, F. K., Varela-Aldas, J., & Mostafa, S. M. (2025). Depth-Enhanced Tumor Detection Framework for Breast Histopathology Images by Integrating Adaptive Multi-Scale Fusion, Semantic Depth Calibration, and Boundary-Guided Detection. *IEEE Access*, 13, 58735–58745. <https://doi.org/10.1109/ACCESS.2025.3554342>
- [15] Sushilkumar S. (2023). Hybrid Image Enhancement and Dense Segmentation for Robust Iris Recognition Using Attention Mechanism. *Journal for ReAttach Therapy and Developmental Diversities*, 6(10s(2), 2398–2408. [https://doi.org/10.53555/jrtdd.v6i10s\(2\).3846](https://doi.org/10.53555/jrtdd.v6i10s(2).3846)
- [16] Alex David, S., Begum, A., Kalilulah, S. I., Ruth Naveena, N., Rajathi, K., & Hemalatha, D. (2024). Deep Learning Advances in Brain Tumor Classification: Leveraging VGG16 and MobileNetV2 for Accurate MRI Diagnostics. In *4th International Conference on Power, Energy, Control and Transmission Systems: Harnessing Power and Energy for an Affordable Electrification of India, ICPECTS 2024*. Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ICPECTS62210.2024.10780014>