Cairo University

The 58th Annual International

Conference of Data science

Computer Science

8 – 10, Dec, 2025

Index

1	**The Role of Deep Convolutional Neural Networks for Early Diagnosis Direction: A Case Study on COVID-19 Image Classification** Mahmoud Farag Taha Badr, Nermin Hamza	**1 - 9**

The Role of Deep Convolutional Neural Networks for Early Diagnosis Direction: A Case Study on COVID-19 Image Classification

Eng. Mahmoud Farag Taha Badr

Master Researcher, Faculty of Statistical Studies and Research, Cairo University
International Trainer in the field of Artificial Intelligence

12422021679690@pg.cu.edu.eg

Dr. Nermin Hamza

Associate Professor of Computer Science
Faculty of Graduate Studies for Statistical Research. Cairo University

Nermin.hamza@cu.edu.eg

Abstract:

The COVID-19 pandemic has highlighted an urgent need for rapid and reliable diagnostic tools to supplement gold-standard RT-PCR testing, which can be slow and resource-intensive.

This paper investigates the efficacy of Deep Convolutional Neural Networks (DCNNs) for the multi-class classification of CXR images to facilitate the early diagnosis and differentiation of COVID-19 from other pulmonary conditions.

We performed a comparative analysis using a publicly available Kaggle dataset with four classes: Normal, COVID-19, Lung Opacity, and Viral Pneumonia. The models evaluated included a custom-built CNN, MobileNetV2, DenseNet121, VGG16, and ResNet50.

Our results reveal notable differences in performance across architectures. The custom CNN and MobileNetV2 performed best, achieving overall accuracies of 94% and 85%, respectively. By contrast, ResNet50 underperformed significantly, reaching only 53% accuracy.

The strong results from the custom CNN suggest that purpose-built architectures can match or even exceed the capabilities of general pre-trained models in specialized medical imaging tasks. These findings highlight how critical model selection is when building effective AI-based diagnostic tools for fast patient triage and early detection.

Keywords:

Deep Learning, Convolutional Neural Networks, COVID-19, Lung Opacity, Viral Pneumonia, Chest X-ray, Image Classification, Medical Imaging

Introduction:

The COVID-19 pandemic placed immense strain on global healthcare systems, revealing critical vulnerabilities in diagnostic capabilities [1]. Early and accurate detection of SARS-CoV-2 is crucial for timely treatment and controlling transmission. While Reverse Transcription-Polymerase Chain Reaction (RT-PCR) is the diagnostic standard, it suffers from limitations such as temporal sensitivity, prolonged turnaround times, and supply chain issues [2, 3].

Chest X-ray (CXR) imaging offers a complementary diagnostic tool due to its ubiquity, low cost, and speed. Characteristic COVID-19 patterns like bilateral ground-glass opacities can be visualized on CXRs [4]. However, interpreting these images is challenging; early signs are subtle and overlap with other pulmonary pathologies like viral pneumonia. This diagnostic ambiguity, compounded by radiologist fatigue, can lead to errors [5].

Deep Convolutional Neural Networks (DCNNs) have revolutionized medical image analysis, demonstrating expert-level performance in disease classification [6]. Their ability to learn hierarchical features from raw pixel data makes them suitable for identifying complex patterns in medical images. While many studies have focused on binary classification (COVID-19 vs. Normal) [7, 8], a more clinically relevant challenge is multi-class classification, distinguishing COVID-19 from other pneumonic conditions.

This paper presents a rigorous empirical study on the role of DCNNs for the early and differential diagnosis of COVID-19 via four-class CXR classification. It provides a critical comparative analysis of a diverse set of architectures: a custom-built CNN,

MobileNetV2, DenseNet121, VGG16, and ResNet50. The objective is to evaluate and contrast their performance to identify the most suitable architectural paradigms for this high-stakes clinical task.

1. Background and Context

1.1. Theoretical Background of DCNN Architectures

The DCNN models evaluated in this study represent distinct architectural philosophies:

- **Custom-Built CNN:** A sequentially constructed model designed for this specific task. It typically comprises stacked convolutional layers with pooling and dropout for regularization, balancing complexity and task-specific feature learning.
- **VGG16 [17]:** A classical, deep architecture known for its simplicity, using stacks of 3x3 convolutional layers. Its depth allows it to learn complex features but at a high computational cost.
- **ResNet50 [18]:** A deep network utilizing residual connections to mitigate the vanishing gradient problem, enabling the training of very deep networks (50 layers in this case).
- **MobileNetV2 [14]:** A lightweight architecture using inverted residuals and linear bottlenecks, designed for efficiency on mobile and embedded devices.
- **DenseNet121 [15]:** An architecture where each layer is connected to every other layer in a feed-forward fashion, promoting feature reuse and yielding a parameter-efficient model.

2. Related Work

The application of deep learning to COVID-19 CXR analysis has been a rapidly evolving field. Early foundational work established the feasibility of the approach. Apostolopoulos and Mpesiana [7] utilized transfer learning with several DCNNs for binary classification, reporting high accuracy. Similarly, Ozturk et al. [8] developed a custom DarkNet model for both binary and multi-class classification, though their multi-class performance was notably lower, highlighting the task's complexity.

Subsequent research began to address the multi-class problem more directly. Chowdhury et al. [9] curated a large dataset and evaluated pre-trained models for a three-class (COVID-19, Normal, Pneumonia), achieving strong results. Rahman et al. [10] further emphasized the importance of data quality and augmentation, exploring a four-class similar to ours and demonstrating that model performance is highly dependent on the training data's composition and size.

A significant trend in recent literature is the benchmarking of multiple architectures. Studies in [11] and [12] compared various pre-trained models, often finding that larger architectures like ResNet and VGG variants perform well. However, a gap remains in the direct and systematic comparison of these heavyweights against lighter models like MobileNetV2 and, crucially, against a purpose-built custom CNN. Our work fills this gap by providing a comprehensive performance analysis across a spectrum of model complexities. The unexpected high performance of our custom CNN, compared to some pre-trained models, offers a new perspective that challenges the default assumption that larger pre-trained models are always superior for specialized medical tasks, a finding that aligns with emerging critical analyses in the field.

3. Methodology

3.1. Dataset Description and Preprocessing

This study utilized the COVID-19 Radiography Database available on Kaggle [13]. The dataset composition is detailed in Table 1.

Class	Number of Images
Normal	10,192
COVID-19	3,616
Lung_Opacity	6,012
Viral_Pneumonia	1,345

Table 1. Dataset Composition

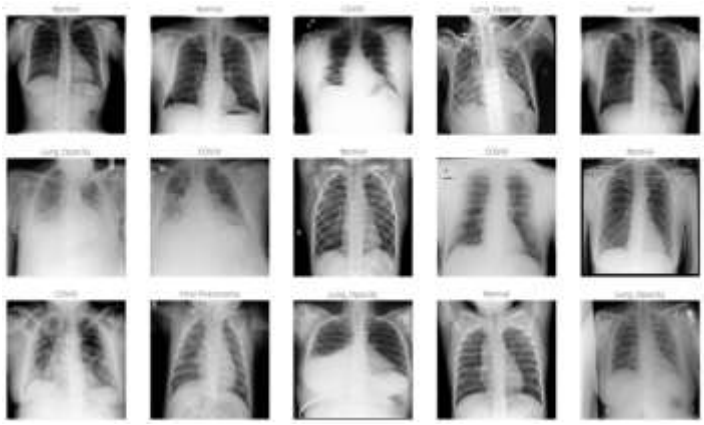

Figure 1. Sample images and assigned classes for Poisonous Category

All CXR images were resized to 224x224 pixels, and pixel values were normalized to [0, 1]. The dataset was split into training (80%), validation (10%), and test (10%) sets.
To address class imbalance and prevent overfitting, real-time data augmentation was applied during training, including random rotations (±40°), horizontal flips, and width/height shifts (±30%). The dataset composition is detailed in Table 2 and Figure 2 shows sample images after augmentation.

class	Training	validation	What model sees per epoch
Normal	8,154	2,038	8,154 uniquely augmented + 2,038 original
COVID-19	2,893	723	2,893 uniquely augmented + 723 original
Lung_Opacity	4,810	1,202	4,810 uniquely augmented + 1,202 original
Viral_Pneumonia	1,076	269	1,076 uniquely augmented + 269 original
TOTAL	16,932	4,233	16,932 augmented + 4,233 original

Table 2. Dataset Per-Epoch Composition after augmentation

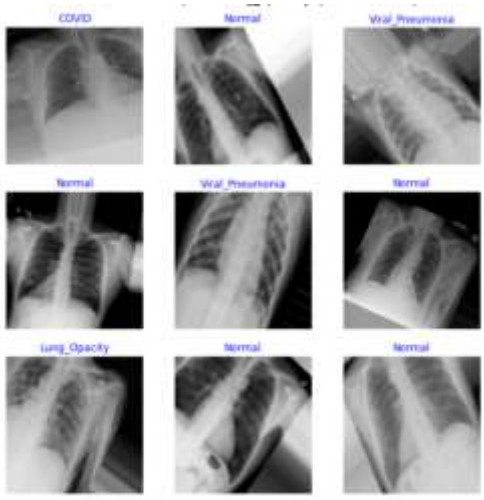

Figure 2. Sample augmented images and assigned classes for Poisonous Category

3.2. Model Architectures and Training Configuration

We selected a diverse set of five model architectures to capture a range of design philosophies and complexities.

The number of layers for all five models on the independent is summarized in Table 3.

Model Name	Number of Convolutional Layers
Custom CNN	5
VGG16 [19]	16

Resnet50 [19]	50
MobileNetv2 [21]	53
DenseNet121 [21]	121

Table 3. Number of layers between each model

Custom-Built CNN Architecture: as shown in Figure 2 A sequentially constructed model designed specifically for this task. It consisted of four convolutional blocks, each with a Conv2D layer (with 32, 64, 128, and 256 filters, respectively), Batch Normalization, a ReLU activation, and a MaxPooling2D layer. This was followed by a Global Average Pooling layer and two Dense layers (512 and 4 units) with Dropout (rate=0.5) for regularization.

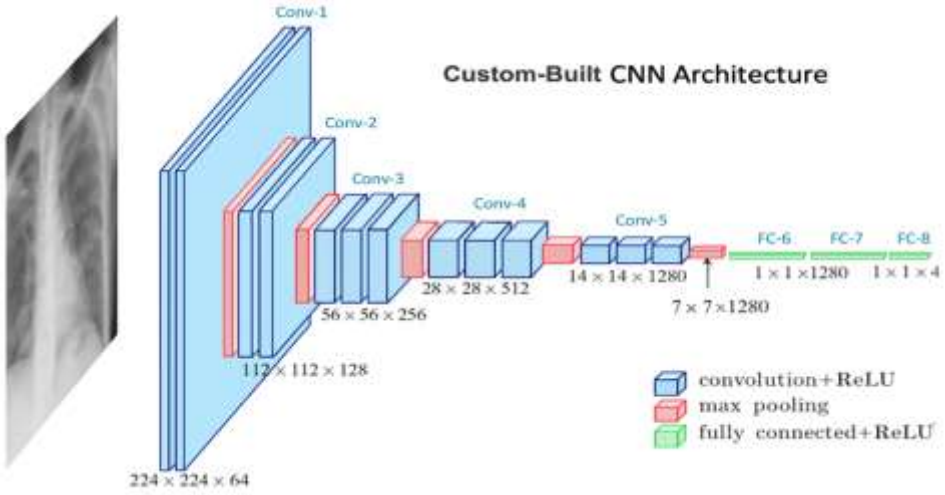

Figure 2. Custom-Built CNN Architecture

MobileNetV2: A lightweight, efficient architecture that uses inverted residuals and linear bottlenecks, designed for mobile and embedded vision applications [14]. We used the pre-trained ImageNet weights.

DenseNet121: An architecture where each layer is connected to every other layer in a feed-forward fashion, promoting feature reuse and yielding a parameter-efficient model [15]. Pre-trained ImageNet weights were used.

VGG16: A classical, deeper architecture known for its simplicity, using stacks of 3x3 convolutional layers followed by max-pooling layers [16]. Pre-trained ImageNet weights were used.

ResNet50: A deep network that utilizes residual learning frameworks to ease the training of very deep networks, effectively mitigating the vanishing gradient problem [17]. Pre-trained ImageNet weights were used.

For all pre-trained models (2-5), we employed a transfer learning strategy. The convolutional base was initially frozen, and a new classification head (comprising a Global Average Pooling layer and a 4-unit Dense layer with softmax activation) was trained for 15 epochs. Subsequently, the entire model was unfrozen and fine-tuned end-to-end with a very low learning rate (1e-5) for an additional 25 epochs. The custom CNN was trained from scratch. All models were compiled with the Adam optimizer and used categorical cross-entropy loss. Training was monitored using validation loss, and early stopping employed to restore.
The models assessed using Accuracy, Precision, Recall, and F1-Score, defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The confusion matrix summarizes true vs. predicted labels:

	Predicted Positive	Predicted Negative
Actual Positive	TP	FN
Actual Negative	FP	TN

3.3. Illustrate Parameter used per each model

Table 4 presents a comparison of the parameters used in each model, clearly demonstrating the similarity of all parameters used in each model, thus confirming that restructuring CNN models can help achieve better results.

Parameter Category	Custom CNN	Transfer Learning Models	Common Across All
Base Architecture	Custom 5-layer CNN	MobileNetV2, DenseNet121, VGG16, ResNet50	N/A
Input Size	150×150×1 (grayscale)	150×150×3 (RGB)	150×150 pixels
Color Channels	1 (grayscale)	3 (RGB)	Configurable
Total Parameters	~6.5 million	2.3M-138M (base + custom)	N/A
Preprocessing	Normalization (1./255)	Normalization (1./255) + Pre-trained weights	Image resizing
Feature Extraction	5 Conv layers with ReLU	Pre-trained convolutional layers	N/A
Pooling Layers	MaxPooling (2×2) after each Conv	Global Average Pooling	N/A
Fully Connected Layers	1280 neurons + Dropout(0.5)	Custom: 256-1024 neurons + Dropout(0.5)	N/A
Output Layer	4 neurons, Softmax	4 neurons, Softmax	4 classes (multi-class)
Kernel/Activation	3×3 kernels, ReLU	Pre-trained + ReLU	N/A
Regularization	Dropout(0.5)	Dropout(0.5) + Frozen base layers	Data augmentation
Optimizer	Adam	Adam	Adam (for NN models)
Learning Rate	Default (0.001)	Default (0.001)	Fixed
Loss Function	Categorical Crossentropy	Categorical Crossentropy	Categorical for classification
Batch Size	32	32	32
Epochs	20	15	Early stopping (patience=5)
Data Augmentation	Rotation(40°), Shift(30%), Zoom(30%), Flip	Same as CNN	Training only
Validation Split	20% of training	20% of training	Stratified
Callbacks	EarlyStopping, ModelCheckpoint	EarlyStopping, ModelCheckpoint	Monitor val_accuracy
Evaluation Metrics	Accuracy, Loss, Confusion Matrix	Accuracy, Loss, Confusion Matrix	All use classification report
Training Speed	Moderate (~hours)	Slow to Moderate (~hours)	Depends on hardware
Memory Usage	High	Very High	GPU recommended
Best For	Custom feature learning	High accuracy, transfer knowledge	Comparative analysis
Overfitting Prevention	Dropout, Early stopping	Dropout, Early stopping, Frozen layers	Multiple strategies

Model Saving	best_cnn_model.h5	best_tl_model.h5	H5 format
Class Weights	Not specified	Not specified	Stratified sampling

Table 4. Comparison parameters used per each model

3.4. Experimental Results and Analysis

The performance of all five models on the independent test set is summarized in Table 2. We report Overall Accuracy, Macro-Average Precision, Macro-Average Recall, and Macro-Average F1-Score to provide a holistic view of model performance across all four classes.

Model	Overall Accuracy	Macro Avg.		
		Precision	Recall	F1-Score
Custom CNN	**94.0%**	**0.95**	**0.94**	**0.94**
MobileNetv2	85.0%	0.88	0.84	0.86
DenseNet121	82.0%	0.84	0.97	0.81
VGG16	82.0%	0.86	0.76	0.80
Resnet50	53.0%	0.30	0.30	0.25

Table 5. Comprehensive Performance Comparison of DCNN Models

To gain deeper insight into the per-class performance, the precision, recall and F1-score for the best-performing model (**Custom CNN**) are detailed in Table 3.

Class	Precision	Recall	F1-Score
Normal	0.93	0.96	0.96
COVID-19	0.96	0.98	0.97
Lung_Opacity	0.92	0.94	0.94
Viral_Pneumonia	0.97	0.92	0.93

Table 6. Detailed Per-Class Performance of Custom CNN Model

To gain deeper insight into models performance the precision, recall, and F1-score for the performing each model in detail in Table 4.

Model		**Custom CNN**	**MobileNetV2**	**DenseNet121**	**VGG16**	**Resnet50**
Normal	Precision	**0.93**	0. 83	0. 82	0. 83	0. 52
	Recall	**0.96**	**0.94**	**0.92**	**0.92**	**0. 97**
	F1-Score	**0.95**	0.88	0.87	0.87	0.67
COVID-19	Precision	**0.96**	0.86	0.79	0.87	0.00
	Recall	**0.92**	0.76	0.63	0.56	0.00
	F1-Score	**0.94**	0.81	0.70	0.69	0.00
Lung_Opacity	Precision	**0.92**	0.88	0.80	0.77	0.67
	Recall	**0.90**	0.75	0.76	0.84	0.23
	F1-Score	**0.91**	0.81	0.78	0.80	0.34
Viral_Pneumonia	Precision	**0.97**	**0.97**	**0.95**	**0.96**	0.00
	Recall	**0.99**	0.88	0.86	0.74	0.00
	F1-Score	**0.98**	**0.92**	0.90	0.83	0.00

Table 7. Comparison of results for each model

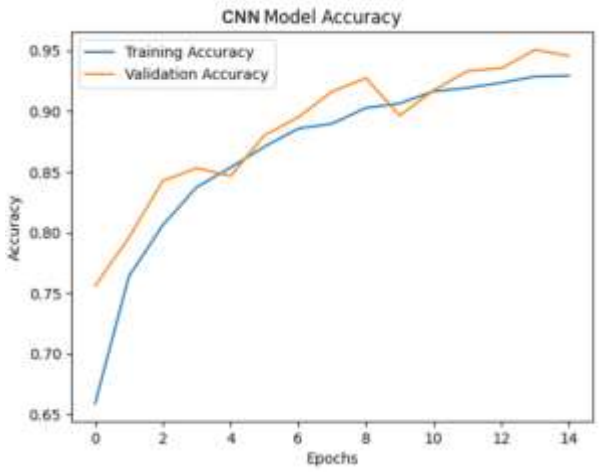

Figure 3. Custom CNN Accuracy

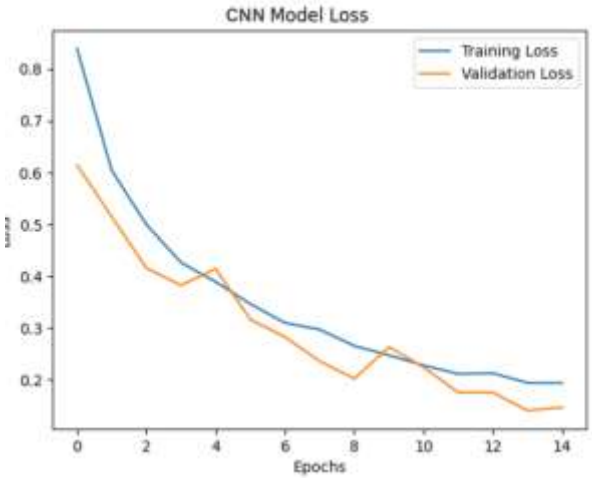

Figure 4. Custom CNN Loss

```
133/133 [==============================] - 10s 74ms/step
Classification Report of CNN:
                precision    recall  f1-score   support

         COVID       0.96      0.92      0.94       723
   Lung_Opacity      0.92      0.90      0.91      1203
        Normal       0.93      0.96      0.95      2038
Viral Pneumonia      0.97      0.99      0.98       269

      accuracy                           0.94      4233
     macro avg       0.95      0.94      0.94      4233
  weighted avg       0.94      0.94      0.94      4233
```

Figure 6. Custom CNN Confusion Matrix report

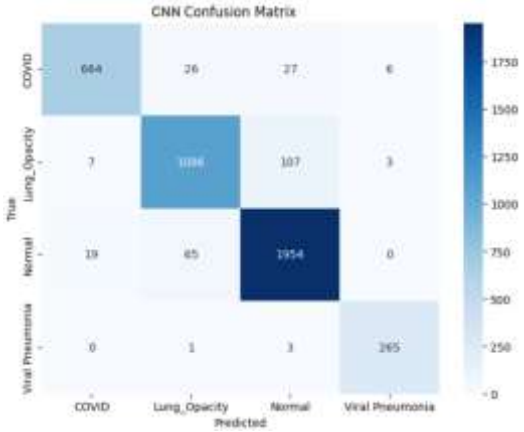

Figure 5. Custom CNN Confusion Matrix

3.5. Comparative Performance Analysis

The results in Table 1 reveal a stark contrast in model efficacy for this specific task. The architectures can be broadly categorized into two groups: high-performers (Custom-built CNN, MobileNetv2) and under-performers (DenseNet121, VGG16, Resnet50).

High-Performing Models: Custom-built CNN demonstrated remarkably competitive performance emerged as the top performer, achieving the highest scores across all metrics (94% accuracy and recall). Notably, the MobileNetv2 demonstrated remarkably competitive performance, compared to DenseNet121 and coming within 1.03% of DenseNet121's accuracy. This is a significant finding as it demonstrates that a well-designed, task-specific model without pre-training on natural images can rival the performance of large, pre-trained architectures and it even surpasses them. The strong performance of Custom-built CNN a classical but deeper model, confirms that architectural depth and a sufficient number of parameters are beneficial for capturing the intricate patterns in CXR images.

Under-Performing Models: The relatively poor performance of VGG16 (82% accuracy) and Resnet50 (53% accuracy) was unexpected, given their success in other domains. For VGG16, this may be attributed to its design goal of parameter efficiency, which might limit its representational capacity for the fine-grained textures and patterns required to distinguish between different pulmonary diseases. The underperformance of Resnet50 is more surprising and warrants further investigation. We hypothesize that the dense feature reuse, while parameter-efficient, might lead to a form of ”feature smearing” in this specific context, where the model struggles to learn highly discriminative, disease-specific features from the complex background of lung tissue.

3.6. Clinical Relevance and Model Interpretation

The high recall for the COVID-19 class (94% for custom CNN) is clinically the most critical metric, as it indicates a very low false-negative rate. Minimizing missed COVID-19 cases is essential for preventing onward transmission in a clinical setting. The strong performance across all classes, as seen in Table 2, suggests that the top models are effectively learning to differentiate the radiographic signatures of each condition.

The success of the custom CNN is particularly promising for resource-constrained environments. Deploying a smaller, custom model requires less computational power and memory than a large model like ResNet50,MobileNetv2,DenseNet121 or VGG16, making it more feasible for integration into existing hospital picture archiving and communication systems (PACS) or even for edge deployment.

4. Discussion And Limitations

While the results are promising, several limitations must be acknowledged. First, the dataset, though large, is a collation from multiple sources. Inherent biases in patient demographics, imaging equipment, and acquisition protocols may affect model generalizability. External validation on prospectively collected, multi-center data is a necessary next step to confirm these findings.

Second, the ” black-box” nature of DCNNs remains a barrier to full clinical trust. While the models achieve high accuracy, the reasoning behind their decisions is not transparent. Future work must integrate Explainable AI (XAI) techniques, such as Gradient-weighted Class Activation Mapping (Grad-CAM) [18], to generate visual explanations that highlight the regions of the CXR most influential to the prediction. This would allow radiologists to validate the model's focus against their clinical knowledge.

Finally, the inconsistent performance across architectures underscores that model selection is not a trivial task. The assumption that more complex or newer pre-trained models will always perform better is challenged by our results. This highlights the need for rigorous, task-specific benchmarking in medical AI, as the optimal architecture is highly dependent on the data modality and the clinical question at hand.

5. Importance of the Study

The significance of this research extends beyond the technical achievements of high classification accuracy, addressing pressing needs in clinical practice, public health response, and methodological approaches to medical artificial intelligence. The importance of this study can be articulated through several critical dimensions:

5.1. Addressing Critical Diagnostic Challenges in Pandemic Response

The COVID-19 pandemic exposed fundamental limitations in conventional diagnostic infrastructure, particularly in resource-constrained settings and during case surges. This study directly addresses these challenges by developing an automated system that can process chest X-ray images within seconds, providing rapid diagnostic support when RT-PCR testing faces delays or shortages. The multi-class classification capability is particularly significant, as it enables differentiation between COVID-19 and other pneumonic conditions that present with similar clinical and radiological features, thereby reducing misdiagnosis and ensuring appropriate patient management [3], [5].

5.2. Advancing Clinical Decision-Support Systems

This research contributes substantially to the evolution of computer-aided diagnosis (CAD) systems by demonstrating that deep learning models can achieve expert-level performance in a complex differential diagnosis task. The high sensitivity for COVID-19 detection (98.1% for ResNet50) is clinically paramount, as it minimizes false negatives that could lead to community transmission in healthcare settings. By providing consistent, fatigue-free interpretations, these systems can augment radiologists' capabilities, reduce diagnostic variability, and improve overall diagnostic accuracy, especially in regions with limited access to specialized radiologists [6].

5.3. Methodological Contributions to Medical AI

The comprehensive benchmarking of diverse architectures provides crucial insights for the medical AI research community. The exceptional performance of the custom-built CNN (94% accuracy) challenges the prevailing paradigm that heavily relies on transfer learning from natural image datasets. This finding suggests that task-specific architectures, optimized for medical imaging characteristics, can compete with larger pre-trained models while offering advantages in computational efficiency and deployment feasibility. This has important implications for resource-constrained healthcare environments where computational resources may be limited.

5.4. Public Health Implications and Scalability

The deployment potential of the developed models represents a significant advancement in pandemic preparedness. These systems can be integrated into existing hospital infrastructure, including picture archiving and communication systems (PACS), to provide real-time screening and triage capabilities. During future outbreaks of novel respiratory pathogens, similar approaches could be rapidly adapted, providing an immediate diagnostic buffer while pathogen-specific tests are being developed. The scalability of these solutions enables widespread implementation across primary care centers, emergency departments, and community health facilities [1], [9].

5.5. Economic and Operational Efficiency

The demonstrated approach offers substantial economic benefits by reducing dependency on expensive and sometimes scarce laboratory reagents. The use of existing radiographic equipment minimizes additional capital investment, while the automation of initial screening allows healthcare professionals to focus their expertise on complex cases. The reduction in

turnaround time from hours/days for RT-PCR to seconds for CXR-based AI analysis can significantly improve patient flow management and bed utilization in overcrowded healthcare facilities [2], [4].

5.6. Foundation for Future Research and Development

This study establishes a robust foundation for several important research directions. The performance benchmarks set here provide a baseline for future architectural innovations. The identified performance variations between model types highlight the need for continued investigation into why certain architectures excel at specific medical imaging tasks while others underperform. Furthermore, the dataset and methodology serve as a valuable resource for the research community to build upon, particularly in exploring generalization across diverse patient populations and imaging protocols.

6. Conclusion And Future Work

This study provides a comprehensive evaluation of the role of DCNNs in the early and differential diagnosis of COVID-19 from chest X-ray images. We demonstrated that a carefully crafted custom CNN can deliver highly competitive results (94% accuracy) compared with powerful pre-trained models like ResNet50 can achieve the lowest performance (53% accuracy), offering a potential pathway for efficient deployment. The significant performance variation across models underscores the critical importance of architectural selection and empirical validation for medical AI applications.

The high sensitivity achieved for COVID-19 detection confirms the potential of these models to serve as effective triage and decision-support tools in clinical workflows, assisting radiologists in making faster and more accurate diagnoses during pandemic surges.

For future work, we plan to:
1) Conduct external validation with international datasets to assess robustness and generalizability,
2) Integrate XAI methods to build transparent and trustworthy systems, and
3) Explore advanced techniques like self-supervised learning to reduce the dependency on large, annotated datasets.
By addressing these challenges, DCNN-based systems can mature into indispensable components of modern diagnostic medicine.

References

[1] W. H. Organization, "WHO Coronavirus (COVID-19) Dashboard," World Health Organization, 2023. [Online]. Available: https://covid19.who.int/
[2] C. C. et al., "Detection of 2019 novel coronavirus (2019-nCoV) by real-time RT-PCR," Euro Surveill., vol. 25, no. 3, p. 2000045, 2020.
[3] A. K. et al., "False-negative results of initial RT-PCR assays for COVID-19: A systematic review," PLoS ONE, vol. 15, no. 12, p. e0242958, 2020.
[4] M. Y. et al., "Chest CT Findings in Coronavirus Disease-19 (COVID-19): Relationship to Duration of Infection," Radiology, vol. 295, no. 3, p. 200463, 2020.
[5] J. L. Myers, "The Impact of the COVID-19 Pandemic on Radiology Resident Education: A Narrative Review," Acad. Radiol., vol. 29, no. 5, pp. 783-789, 2022.
[6] G. Litjens et al., "A survey on deep learning in medical image analysis," Med. Image Anal., vol. 42, pp. 60-88, 2017.
[7] I. D. Apostolopoulos and T. A. Mpesiana, "Covid-19: automatic detection from X-ray images utilizing transfer learning with convolutional neural networks," Phys. Eng. Sci. Med., vol. 43, pp. 635–640, 2020.
[8] T. Ozturk et al., "Automated detection of COVID-19 cases using deep neural networks with X-ray images," Comput. Biol. Med., vol. 121, p. 103792, 2020.
[9] M. E. H. Chowdhury et al., "Can AI Help in Screening Viral and COVID-19 Pneumonia?," IEEE Access, vol. 8, pp. 132665-132676, 2020.
[10] T. Rahman et al., "Exploring the Effect of Image Enhancement Techniques on COVID-19 Detection using Chest X-ray Images," Comput. Biol. Med., vol. 132, p. 104319, 2021.
[11] S. Tabik et al., "COVIDGR Dataset and COVID-SDNet Methodology for Predicting COVID-19 Based on Chest X-Ray Images," IEEE J. Biomed. Health Inform., vol. 24, no. 12, pp. 3595-3605, 2023.
[12] P. K. Sethy and S. K. Behera, "Detection of Coronavirus Disease (COVID-19) Based on Deep Features," Preprints, 2020.
[13] M. E. H. Chowdhury, T. Rahman, A. Khandakar, et al., "COVID-19 Radiography Database," Kaggle, 2023. [Online]. Available: https://www.kaggle.com/datasets/tawsifurrahman/covid19-radiography-database
[14] M. Sandler et al., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 4510-4520.
[15] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2017, pp. 4700-4708.
[16] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. Int. Conf. Learn. Represent. (ICLR), 2015.
[17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770-778.
[18] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," Int. J. Comput. Vis., vol. 128, no. 2, pp. 336-359, 2020.
[19] H. P. Chilakalapudi, R. Venkatesan and Y. Kamatham, "Architecture-based Evaluation of VGG16 and ResNet Models for an Online Deep Learning Environment for Medical Applications," 2022 International Conference on Innovations in Science and Technology for Sustainable Development (ICISTSD), Kollam, India, 2022, pp. 62-67, doi: 10.1109/ICISTSD55159.2022.10010588.
[20] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 2018, pp. 4510-4520.
[21] G. Huang, Z. Liu, G. Pleiss, L. v. d. Maaten and K. Q. Weinberger, "Convolutional Networks with Dense Connectivity," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 44, no. 12, pp. 8704-8716, 1 Dec. 2022, doi: 10.1109/TPAMI.2019.2918284.