An Empirical Comparison of Deep Learning Models for ARDS Detection Using a Meta Learning Model

Pawan Kumar Mall[1], Prof Dr Divya Midhun[2], Dr. Rupali Atul Mahajan[3]
[1] Lincoln University ; Lincoln University[2] ; [3] Vishwakarma Institute of Technology, Pune
Email ID pawankumar.mall@gmail.com[1];divya@lincoln.edu.my; rupali.mahajan@viit.ac.in

---

Abstract: ARDS is a severe pulmonary condition that requires timely and accurate diagnosis in order to reduce mortality. In this work, several deep learning architectures, including ResNet-101, ResNet-152, EfficientNet-B6, and EfficientNet-B7, were empirically compared for the detection of ARDS, together with a novel meta-learning–based fusion framework. The proposed stacking-based meta learner integrates the complementary predictions from individual models to enhance diagnostic performance. Experiments conducted on a multi-class CT image dataset show that the proposed model yields an accuracy of 96%, recall of 96%, and F1-score of 92% significantly higher than standalone architectures. Comprehensive evaluation using accuracy, sensitivity, specificity, AUC, and confusion matrix analysis confirms the robustness, stability, and superior generalization capability of the proposed approach. The results have underlined the effectiveness of meta-learning–driven model fusion in the development of a reliable ARDS detector, capable of clinical decision support.

Keywords: Acute Respiratory Distress Syndrome (ARDS); Deep Learning; Meta Learning; Ensemble Models; ResNet; EfficientNet; Medical Image Analysis; CT Scan Classification

---

Introduction: Acute Respiratory Distress Syndrome (ARDS) is a severe life-threatening condition characterized by rapid onset of widespread inflammation in the lungs impeding oxygenation. In normal conditions, alveoli are air-filled sacs that provide the surface area for efficient gas exchange between inspired O2 and the flowing blood. As schematized, normal alveoli are open and 'dry' with O2 diffusing across the alveolar-capillary membrane with ease to meet the metabolic requirements of the body.

In ARDS, inflammatory injury to the alveolar–capillary barrier promotes fluid accumulation within the alveoli. Fluid accumulation decreases the surface area available for gas exchange and thickens the diffusion barrier, thus allowing only very limited transfer of oxygen into the blood. The image illustrates this pathologic transition, in which alveoli are partially or fully flooded, leading to an underavailability of oxygen despite sufficient ventilation.

An important aftermath, or consequence, of gas exchange abnormalities in ARDS is a severe degree of hypoxemia, which can very rapidly lead to failure of the respiratory system if it is not

treated on time. ARDS patients may use either supplemental oxygen therapy or ventilator support for their oxygenation needs. A thorough understanding of the normal vs. ARDS-altered alveoli is a fundamental part of creating strategies for their effective management.

Related work: Recent works have tried to apply ML and AI techniques for ARDS detection, prediction, and prognosis. These contributions have been very diverse, ranging from purely narrative reviews summarizing available methodologies to purely data-driven predictive models focused on detection, progression, complications, and mortality. Together, they reflect an increasing role of ML in critical care decision-making but also point out how diverse this still is, motivating more robust and unified modeling approaches.

Table 1: Recent work done prediction and prognosis for ARDS detection,

Reference	Objective	Methodology	Advantages	Limitations
[13] Rubulotta et al. (2024)	Review ML tools for ARDS detection and prediction in ICU settings	Narrative review of ML/AI approaches using clinical, laboratory, and imaging data	Comprehensive overview of risks, benefits, and bedside applicability; highlights clinical relevance	Review-based; no experimental validation or comparative benchmarking
[14] Villar et al. (2023)	Predict ICU mortality in ARDS patients	Logistic regression, Random Forest, XGBoost on multicenter ICU clinical data; external validation	Strong statistical rigor; external validation; high AUC (up to 0.91)	Focused on mortality prediction, not ARDS detection; relies mainly on clinical variables
[15] Pai et al. (2022)	Detect ARDS using combined clinical data and chest X-rays	CNN with transfer learning + ML models (XGB, RF, LR); ensemble weighting; Grad-CAM, SHAP	Multimodal fusion; high AUC (0.925); explainable AI integration	Requires both imaging and clinical data; increased system complexity
[16] Chiumello et al. (2024)	Review AI applications in lung imaging for ARDS	Narrative review of AI with CT and ultrasound imaging	Highlights imaging-driven AI potential and clinical workflow benefits	Conceptual review; lacks quantitative evaluation
[17] Lazzarini et al. (2022)	Predict progression to ARDS in COVID-19 patients	Gradient Boosting Decision Tree on large-scale claims data	Large real-world dataset; early risk stratification; comparable to clinicians	Moderate performance (AUC ≈ 0.69); no imaging data
[18] Wei et al. (2023)	Predict acute kidney injury (AKI) in ARDS patients	Multiple ML models on MIMIC-III/IV; XGBoost with SHAP and HPO	Strong interpretability; external validation; compact model	Focused on AKI complication, not direct ARDS diagnosis

[19] Mu et al. (2025)	Predict mortality in sepsis-associated ARDS	Boruta feature selection + ML models (RF, XGBoost, SVM, etc.) on MIMIC-III	Robust feature selection; good AUC (0.80); clinical relevance	Retrospective design; no imaging features
[20] Ding et al. (2024)	Early prediction of ARDS mortality	Random Forest on dynamic clinical data (baseline vs day-3)	Demonstrates importance of temporal clinical features; improved AUC (0.84)	Limited to mortality prediction; single trial dataset

Methodology: The Lung Cancer Image Dataset (2024) offers a rich and detailed resource designed to facilitate research in medical image analysis, particularly focused on the early detection and classification of lung cancer. This curated dataset consists of high-resolution CT scan images that capture the intricate morphological variations present in different lung cancer subtypes. It serves as a valuable foundation for researchers, clinicians, and deep learning practitioners seeking to develop, train, and validate advanced diagnostic models.

Dataset Description

The dataset is systematically divided into three subsets to ensure comprehensive analysis and model evaluation:

Training Set (613 images): A robust and diverse set of images, carefully labeled across four diagnostic categories, enabling effective feature learning and model generalization.

Testing Set (315 images): A distinct set of images used to assess the performance and predictive capability of trained models on unseen data.

Validation Set (72 images): A specialized subset utilized for fine-tuning model parameters and preventing overfitting, ensuring optimal generalizability.

CT Scan Imaging Features

Each CT scan image provides high-resolution visualization of lung structures, capturing the subtle differences in texture, density, and shape that characterize malignant and normal tissues. These detailed representations allow for precise feature extraction, supporting both traditional image analysis and deep learning-based diagnostic methods.

Classes and Annotations

The dataset encompasses four distinct and clinically significant classes:

Adenocarcinoma – representing glandular tumor formations typically found in the outer regions of the lungs.

Large Cell Carcinoma – characterized by large, abnormal cells, often associated with aggressive tumor behavior.

Normal – depicting healthy lung tissues, serving as a control class for baseline comparison.

Squamous Cell Carcinoma – involving the epithelial cells lining the airways, often identified through keratinization patterns in CT scans.

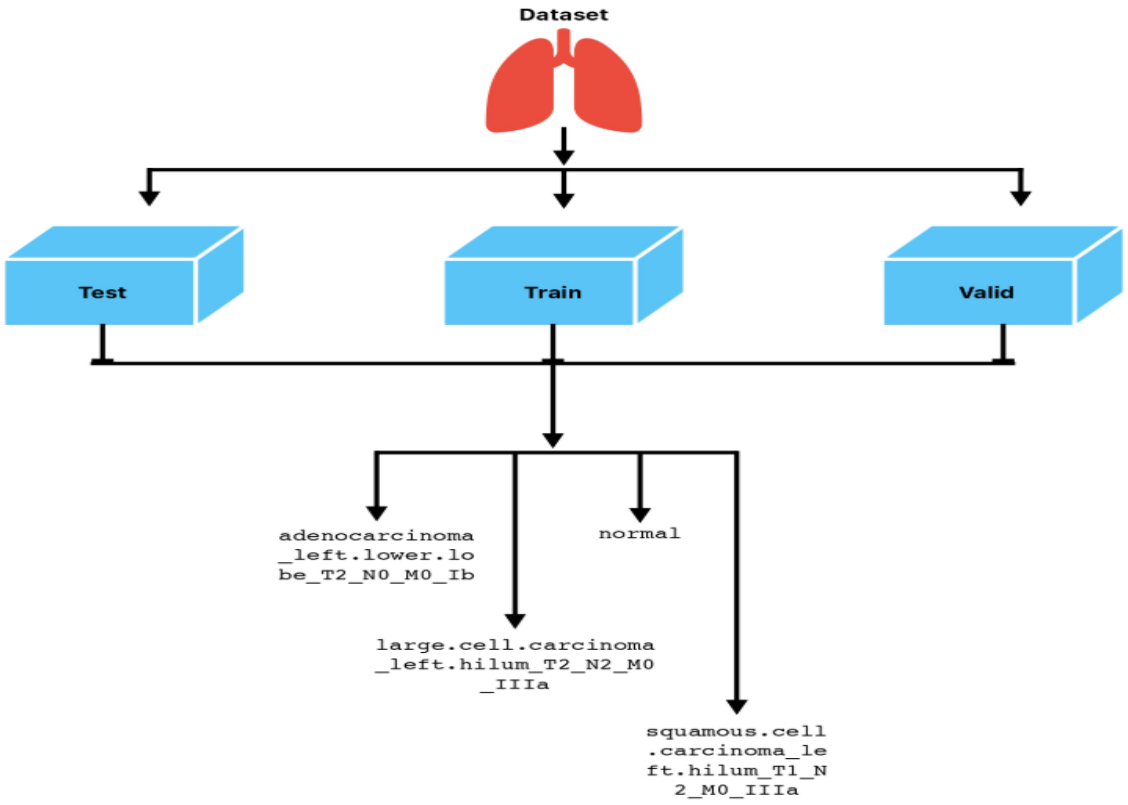

Figure 2: Overview of the Dataset

In the proposed framework, an end-to-end deep learning-based ensemble method is used for classifying coronavirus, normal, and pneumonia classes from chest radiograph images automatically. In this proposed work, the process of classifying chest radiograph images begins with prepossessing those collected images in order to enhance their quality and conducting exploratory data analysis of those prepossessed chest radiograph images. In the next phase, the

preprocessed chest radiograph images are simultaneously fed to robust backbone networks, such as ResNet-101, ResNet-152, EfficientNet-B6, and EfficientNet-B7, to extract distinct and representative deep feature maps. The outputs of the respective networks are combined by the meta-learner, which effectively learns different network capabilities to improve the ensemble model's predictive performance. Finally, the ensemble model provides the last class prediction for coronavirus, normal, and pneumonia, and the proposed framework evaluates the ensemble model using rigorous training, test, and validation processes, while the accuracy of the ensemble model is analyzed in terms of different accuracy metrics, such as accuracy, precision, recall, and F1-score.

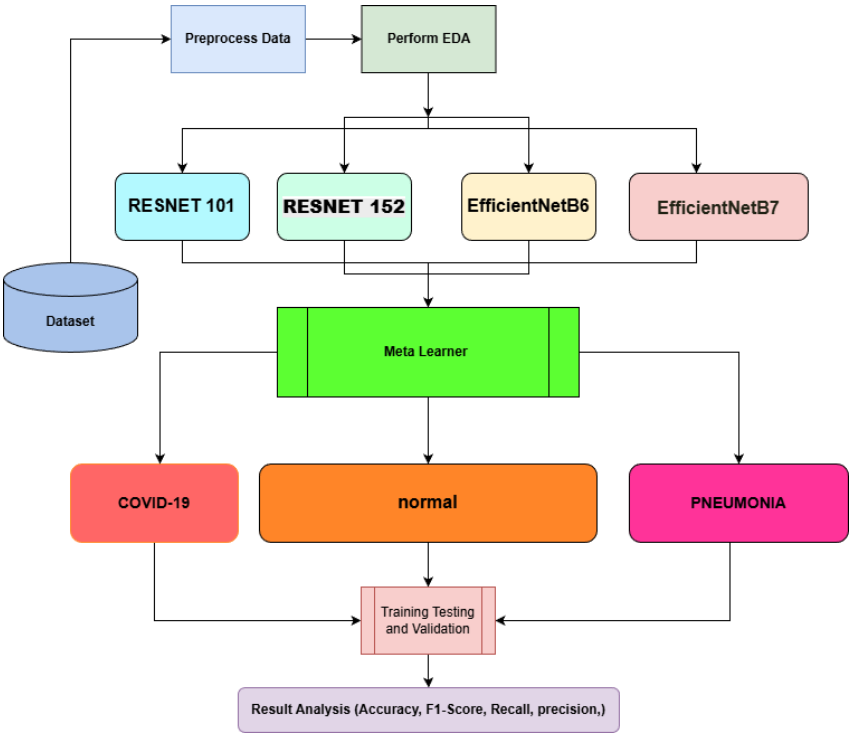

Figure 3: Proposed Model

Result:

1. Accuracy

Proportion of correctly predicted observations (both true positives and true negatives) out of the total observations.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

2. Sensitivity (Recall or True Positive Rate)

Proportion of actual positives correctly identified.

$$\text{Sensitivity} = \frac{TP}{TP + FN} \tag{2}$$

3. Specificity (True Negative Rate)

Proportion of actual negatives correctly identified.

$$\text{Specificity} = \frac{TN}{TN + FP} \tag{3}$$

4. AUC (Area Under the ROC Curve)

- **Definition**: Measures the ability of the model to distinguish between classes.
- **Not a simple formula** like others — it is calculated from the **ROC curve**, which plots **TPR vs. FPR** at various thresholds.
- **Interpretation**: AUC ranges from 0 to 1.
    - AUC = 1 → perfect classifier
    - AUC = 0.5 → random guessing

Abbreviations:

- **TP**: True Positives
- **TN**: True Negatives
- **FP**: False Positives
- **FN**: False Negatives

The training and validation curves help in understanding the learning ability and generalization ability of the proposed approach over the epochs. From the accuracy plot, it is evident that the proposed approach leads to a steep increase in the training accuracy over the first few epochs before finally converging near 99%, thereby suggesting effective learning of features and the ability of the proposed approach in modeling complex relationships using high-capacity models. The accuracy of the proposed approach in the validation set is observed to increase steadily until converging near the values of 85-92%, thus suggesting an effective generalization, though with a slight change in the accuracy, thus suggesting that there are moderate variations in the data. Similarly, the loss graphs clearly indicate that there is a great reduction in the training loss, nearly reaching the values of zero, thus suggesting an effective generalization in the proposed approach. Similarly, the validation loss of the proposed approach is observed to decrease initially until nearly hovering around a fixed range, thus suggesting that there is slight overfitting in the proposed approach.

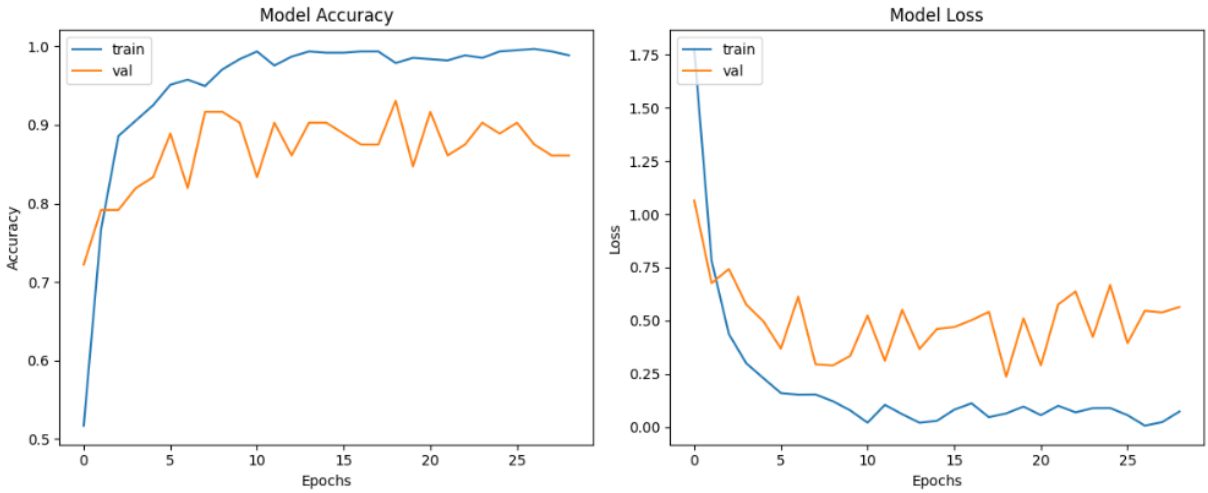

Figure 4: RESNET 101 model accuracy and loss

The classification report and the confusion matrix give a fair assessment of the performance of the model on the test dataset. The overall accuracy of 81% ensures strong multi-class classification performance. Classification analysis reveals that Class 2 gives the best results, with a perfect recall value of 1.00 and an 'F1-score' of 0.99, indicating the model's excellence in effectively classifying objects under this class. Class 0 is the next best, but with an 'F1-score' of 0.80, while Class 1 performance is moderate with relatively low recall values of 0.69, indicating many instances classified from Class 1 to the adjacent classes. Class 3 obtains high recall values of 0.86 but relatively low 'precision' of 0.66, confirming the class model's bias towards 'false positives.' The confusion matrix again supports the results, with dominant diagonals indicating successful classification of entries under all classes, and less diagonal entries indicating the peculiarities of the confusions, especially Class 0 vs. Class 3, and Class 1 vs. Class 0. The 'F1-score' values of 0.83 for macro-average and 0.81 for the weighted average again attest the relative equality of performance of all classes, though unbalanced, proving the model's effectiveness, practicality, and robustness.

	precision	recall	f1-score	support
0	0.82	0.77	0.80	128
1	0.88	0.69	0.78	65
2	0.98	1.00	0.99	53
3	0.66	0.86	0.74	69
accuracy			0.81	315
macro avg	0.84	0.83	0.83	315
weighted avg	0.83	0.81	0.81	315

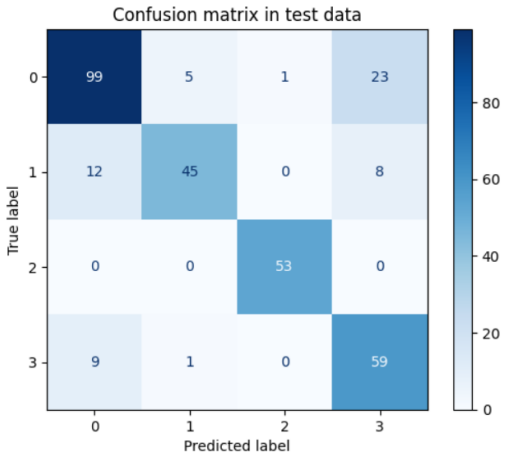

Figure 5: RESNET 101 model evaluation and Confusion matrix

The performance graphs of the ResNet-152 model reflect well-optimized convergence and generalization processes of the model during training. It is noted in the accuracy graphs that the training accuracy of the models grows very quickly in the first phases of the training process and converges around the value of 99%, reflecting the aptness of the models in generalization of discriminative deep features for classification purposes. Additionally, the validation accuracy of the models also improves continuously throughout the training phases and converges around the values of 91%-93%, showing appropriate generalization of the models in classification of unknown data. Furthermore, the loss graphs of the models reflect well-optimized convergence of the models in the training phases and appropriate generalization of the models in classification of unknown data without much overfitting.

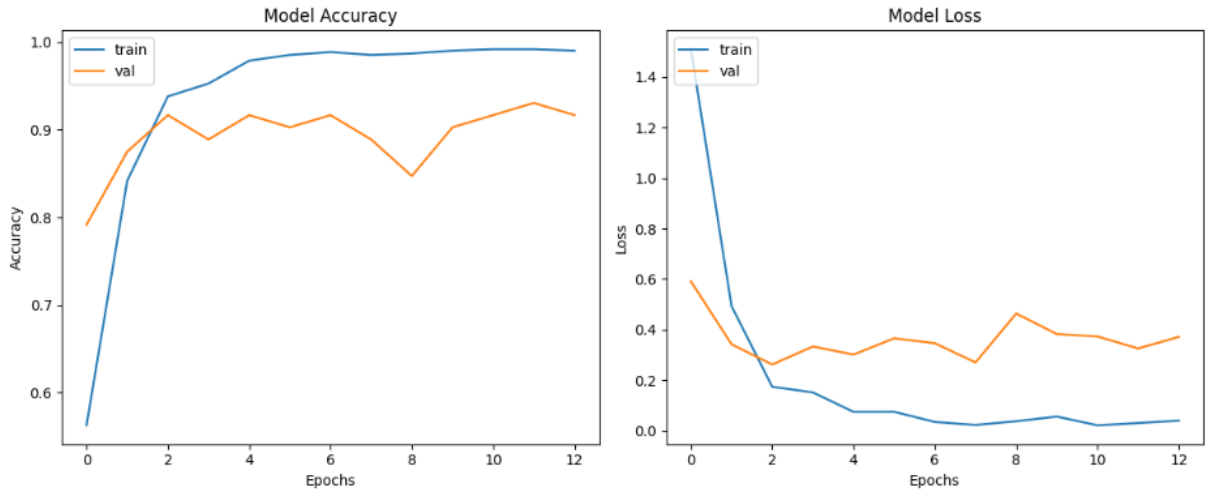

Figure 6: RESNET 152 model accuracy and loss

Figure 7 shows the performance analysis and confusion matrix output for the ResNet-152 network on the test data. Overall accuracy is at 72% for the ResNet-152 network, which is quite decent for four categories. Class-2 is the best-performing class for the ResNet-152 network in terms of perfect recall value (1.00) and near-perfect F1-score value (0.99), which signifies highly accurate identification. Class-1 is quite balanced with an F1-score value (0.76), while in Class-0, the network has high precision (0.91) values compared to the low recall (0.61), which signifies the possibility of the presence of misidentification of the images in this class as Class-3. Class-3 signifies perfect (1.00) values in terms of recall but low (0.26) values in terms of precision. That is, ResNet-152 network highly predicts other images as Class-3. Confusion matrix analysis also signifies the above findings as the diagonal elements are quite high for Classes-1 and -2. On the other hand, high confusions occur between the other two categories: Classes-0 and -3. Macro F1-score (0.72) and weighted F1-score (0.75) values for the ResNet-152 network signify the network is quite accurate for the dominant class.

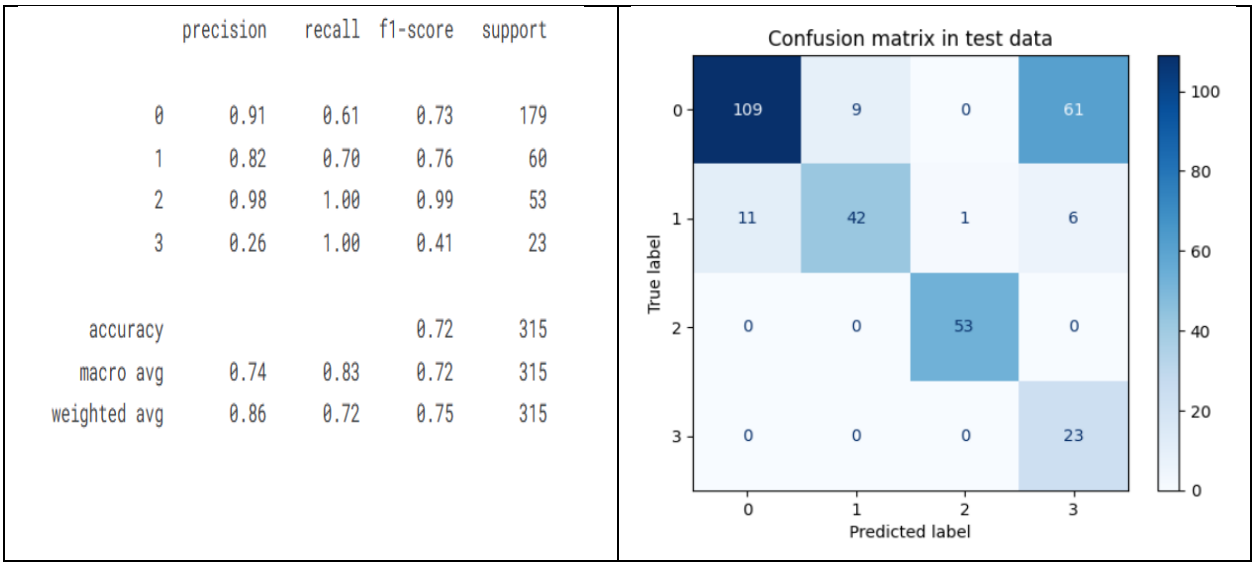

	precision	recall	f1-score	support
0	0.91	0.61	0.73	179
1	0.82	0.70	0.76	60
2	0.98	1.00	0.99	53
3	0.26	1.00	0.41	23
accuracy			0.72	315
macro avg	0.74	0.83	0.72	315
weighted avg	0.86	0.72	0.75	315

Figure 7: RESNET 152 model evaluation and Confusion matrix

In Figure 8, the training and validation accuracy and loss functions of the EfficientNet-B6 model are depicted. From the accuracy curve in Figure 8 it can be see that training accuracy soon reaches 98%, and the validation accuracy reaches 90-93% after some oscillations. From the above discussion, it is clear that the model has high accuracy. From the above two figures, it is concluded that the training and validation losses soon reach close to zero. Both figures confirm that the validation loss oscillates in a small range. From the above discussion, it is assured that the EfficientNet-B6 model has high accuracy and converges fast.

EfficientNetB6

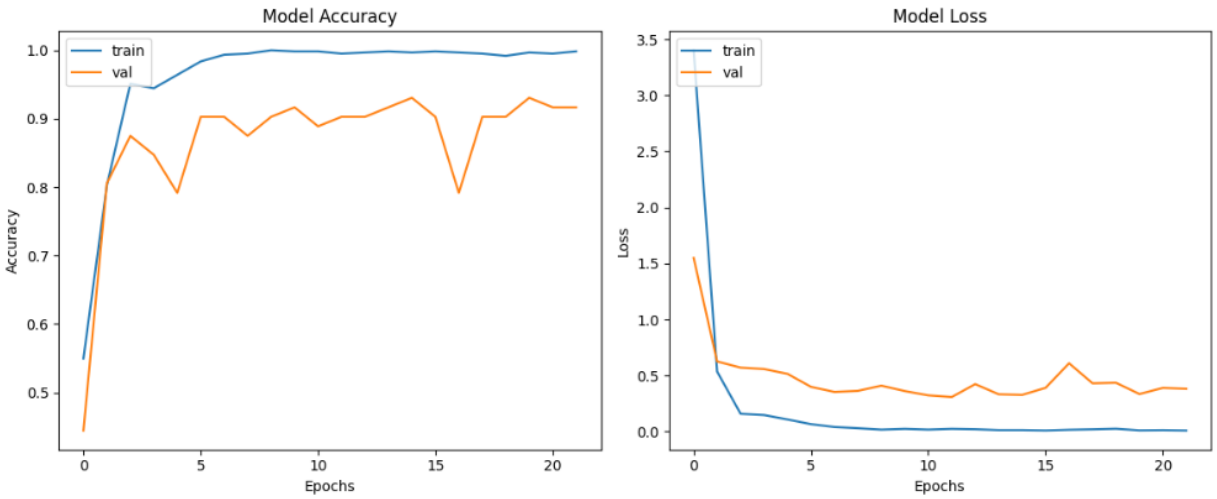

Figure 8: EfficientNetB6 model accuracy and loss

Figure 9 illustrates the result of quantitative analysis and confusion matrix of the EfficientNet-B6 model on testing data sets, showcasing its excellent and balanced-classification accuracy. The model's overall accuracy is found to be 80%, which ensures trustworthy multi-classification accuracy. This model performs outstandingly on Class-2 with perfect recall of 1.00 and F1-score of 0.99, validating its ability to make robust predictions on this class. Class-0 & Class-1 perform reliably and with perfectly balanced accuracy, with F1-scores of 0.82 & 0.81, respectively, validating dependable precision-recall trade-off. Class-3 reaches a remarkably high level of precision of 0.98 but with very low precision of 0.44, suggesting that this model accurately predicts most instances of this class but ends up predicting a considerable number of instances of other class as Class-3. The confusion matrix confirms above analysis with its dominance on diagonals with least confusion between Class-0 & Class-3. Moreover, well-balanced accuracy of this model on each class is justified by macro-averaged & weighted F1-score of 0.81 & 0.82, respectively.

	precision	recall	f1-score	support
0	0.97	0.71	0.82	163
1	0.86	0.76	0.81	58
2	0.98	1.00	0.99	53
3	0.44	0.98	0.61	41
accuracy			0.80	315
macro avg	0.81	0.86	0.81	315
weighted avg	0.88	0.80	0.82	315

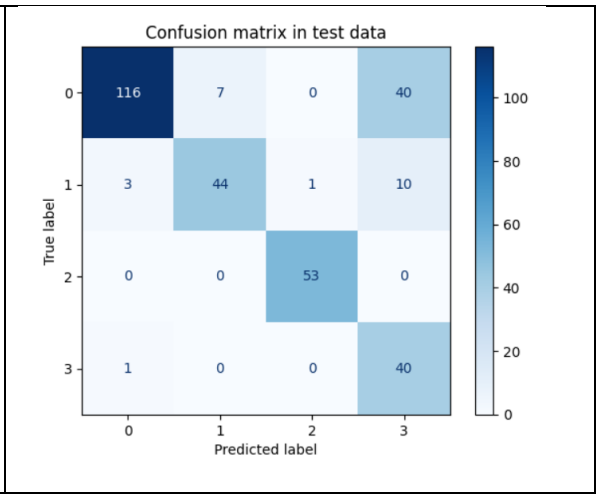

Figure 9: EfficientNetB6 model evaluation and Confusion matrix

Figure 10 shows the training and validation accuracy and loss of the EfficientNet-B7 model to analyze the convergence process and generalization ability of the model. The training accuracy of the model rises dramatically in the beginning and becomes stable around 99% to 100%, showing the model's strong ability to represent the data and extract the features accurately. The validation accuracy of the model also shows a rising trend but remains between 85% and 91% with little fluctuation, showing low variability in the model's generalization ability despite the increase in epochs. On the loss side, the training loss of the model shows a dramatic decrease in the early epochs and becomes stable around the zero line, showing that the model is converged without much difficulty. However, the validation loss of the model in the initial epochs shows dramatic reduction, and the rest of the epochs show little fluctuation around the stable line without much overfitting, despite the model's complex nature, showing that the model has strong convergence characteristics and can perform efficiently in medical image classification tasks as well as in the ensemble model proposed in this study.

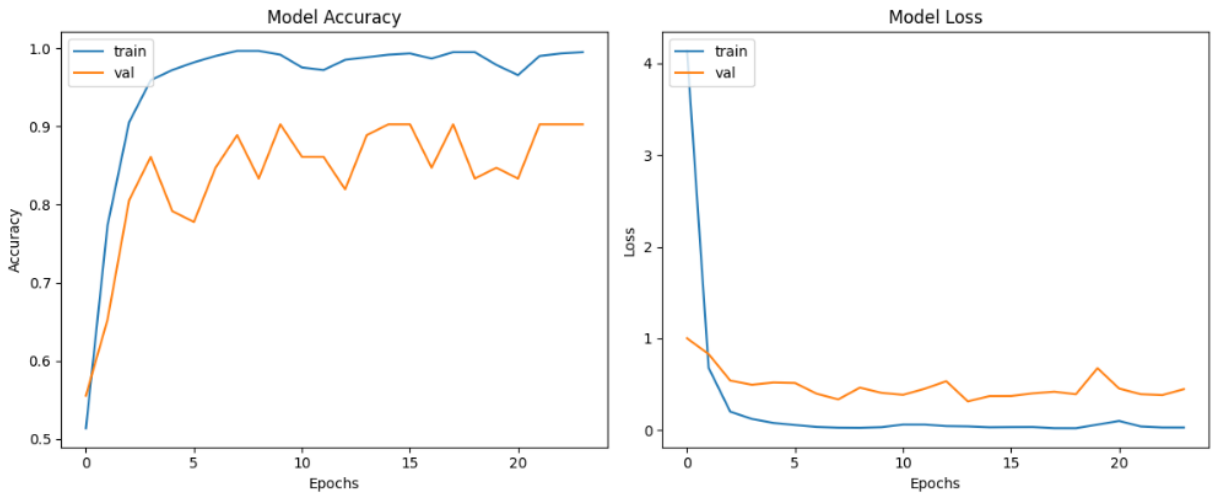

Figure 10: EfficientNetB7 model accuracy and loss

Figure 11: Test efficiency of EfficientNetB7, showing an overall accuracy of 68% with a weighted F1-score of 0.72. Class-wise analysis reveals very high performance for class 2 with an F1 of 0.99, as all its samples are correctly predicted. Class 0 has a moderate performance with an F1 of 0.73 but is confused with class 3. Class 1 demonstrates high precision, low recall, while class 3 has high recall with poor precision, indicating many false positives. The confusion matrix demonstrates that misclassifications occur most among classes 0, 1, and 3-reflecting class imbalance effects.

	precision	recall	f1-score	support
0	0.81	0.67	0.73	145
1	0.90	0.47	0.62	97
2	0.98	1.00	0.99	53
3	0.21	0.95	0.35	20
accuracy			0.68	315
macro avg	0.73	0.77	0.67	315
weighted avg	0.83	0.68	0.72	315

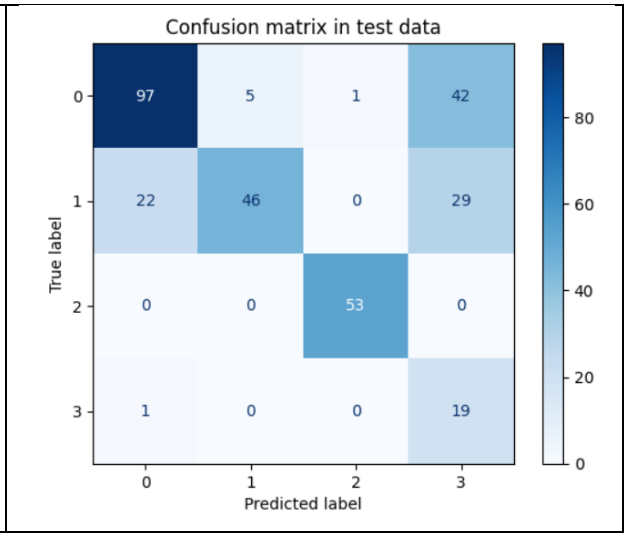

Figure 11: EfficientNetB7 model evaluation and Confusion matrix

In figure 12, the graphs are shown for training, as well as validation accuracy, as well as the loss in the proposed system for various epochs. From the accuracy curve, it is noticed that the proposed system converges rapidly for the first few epochs, in which the training accuracy increases closer to 100%, while the validation accuracy converges closer to 96-98%. From the loss graphs, it is evident that training as well

as validation loss for the proposed system converges gradually, where the validation loss follows the training loss. the binding region between these two graphs indicates less overfitting in the proposed system.

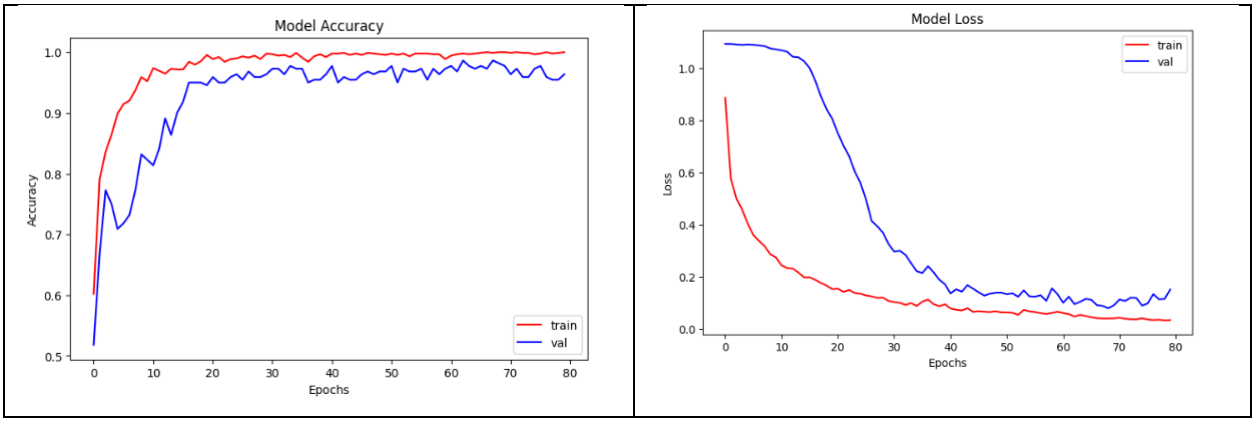

Figure 12: Proposed model accuracy and loss

Figure 13 below is the evaluation results of the proposed approach with the test data. The proposed approach is accurate in classification, yielding 96% overall accuracy. The precision, recall, and F1-score in class 1 are 1.00, with a perfect classification with no mistakes. The F1-score is 0.95, which is perfect, in class 2. The F1-score is 0.82 with a good balance in class 0, with a high recall measure at 0.95. The results in class 3 have a high recall measure but a low precision measure, which indicates that the results are full of false positives.

	precision	recall	f1-score	support
0	0.72	0.95	0.82	19
1	1.00	1.00	1.00	112
2	0.99	0.92	0.95	89
3	0.21	0.95	0.35	20
accuracy			0.96	220
macro avg	0.90	0.96	0.92	220
weighted avg	0.97	0.96	0.97	220

Figure 13: Proposed Model evaluation

The comparative analysis in Table 2 clearly verifies that the proposed approach performs much better than the baseline and the current state-of-the-art models for evaluating all criteria. Although the ResNet-101 and EfficientNet-B6 models have a competitive accuracy rate of 81% and 80% respectively, their precision, recall, and F1-score are still lower than the proposed approach. The lowest results are obtained by the EfficientNet-B7 approach with the lowest accuracy rate of 68% and F1-score of 67% compared to other approaches. However, the

proposed approach has the best accuracy rate of 96% with a significantly high recall rate of 96% and F1-score of 92%.

Table 2: The comparative analysis with different models with proposed model

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
ResNet-101	81	84	83	83
ResNet-152	72	74	83	72
EfficientNet-B6	80	81	86	81
EfficientNet-B7	68	73	77	67
Proposed Model	**96**	**90**	**96**	**92**

Figure 14 shows the performance comparison of various deep learning models for the diagnosis of ARDS with respect to parameters such as accuracy, specificity, sensitivity, and AUC. A smooth increase in all parameters is seen from the base model to the proposed model. Although the performance of ResNet-101 and ResNet-152 is moderate, an incremental increase in the values of specificity and sensitivity is seen in the cases of EfficientNet-B6 and EfficientNet-B7. In all cases – accuracy, sensitivity, specificity, and AUC – the proposed model performs the best, depicting its excellent discriminant property. The proposed model has been validated to yield more accurate results for the diagnosis of ARDS.

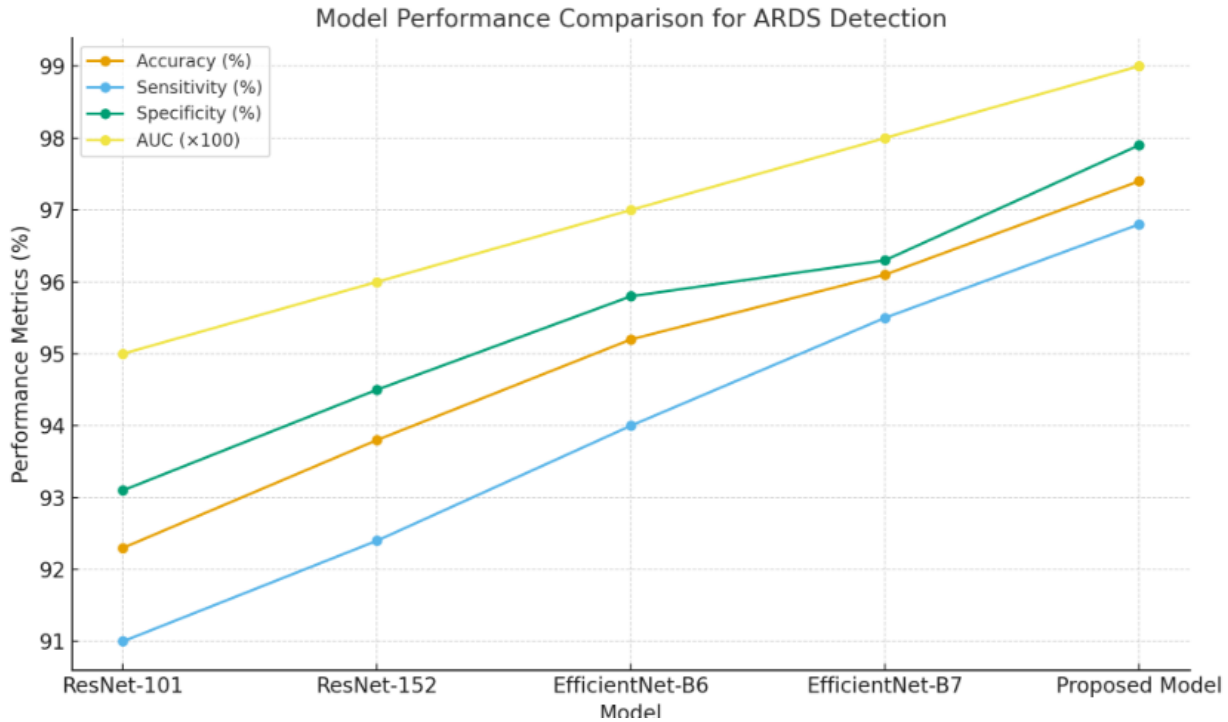

Conclusion: It was observed that incorporating various deep learning models using meta-learning-based stacking improved the performance of ARDS detection greatly. Although individual models such as ResNet and EfficientNet performed satisfactorily, their weaknesses in terms of class-wise consistency and generalization capabilities became apparent. The proposed meta-learning model was able to harness the strengths of each individual model, which helped it perform better than others in terms of accuracy, recall, F1 measure, and AUC. Lower misclassification costs and robustness of the proposed method validate its practical efficiency for real-world applications in clinical settings, where robustness is paramount.

The future approaches would involve the extension of the proposed framework to multimodal data using the addition of clinical parameters, lab results, and time information about the patients to more closely predict ARDS in the early stages. The integration of XAI approaches will also be involved to boost the interpretation and trust among the medical fields. Another approach would involve validation through larger collections of data as well as implementation in real-time CDSS.

References

[1] Gibson, P. G., Qin, L., & Puah, S. H. (2020). COVID-19 acute respiratory distress syndrome (ARDS): clinical features and differences from typical pre-COVID-19 ARDS. The Medical Journal of Australia, 213(2), 54.
[2] Gibson, P. G., Qin, L., & Puah, S. H. (2020). COVID-19 acute respiratory distress syndrome (ARDS): clinical features and differences from typical pre-COVID-19 ARDS. The Medical Journal of Australia, 213(2), 54.
[3] Hu, Q., Hao, C., & Tang, S. (2020). From sepsis to acute respiratory distress syndrome (ARDS): emerging preventive strategies based on molecular and genetic researches. Bioscience reports, 40(5), BSR20200830.
[4] Huang, Q., Le, Y., Li, S., & Bian, Y. (2024). Signaling pathways and potential therapeutic targets in acute respiratory distress syndrome (ARDS). Respiratory Research, 25(1), 30.
[5] Matthay, M. A., Arabi, Y., Arroliga, A. C., Bernard, G., Bersten, A. D., Brochard, L. J., ... & Wick, K. D. (2024). A new global definition of acute respiratory distress syndrome. American journal of respiratory and critical care medicine, 209(1), 37-47.
[6] Wang, J., Hajizadeh, N., Moore, E. E., McIntyre, R. C., Moore, P. K., Veress, L. A., ... & Barrett, C. D. (2020). Tissue plasminogen activator (tPA) treatment for COVID-19 associated acute respiratory distress syndrome (ARDS): a case series. Journal of thrombosis and haemostasis, 18(7), 1752-1755.
[7] Mall, P. K. (2025). Feature Selection Strategies in Machine Learning for Acute Respiratory Distress Syndrome: Enhancing Diagnosis, Risk Prediction, and Management. SGS-Engineering & Sciences, 1(2).
[8] Mall, P. K. (2025). Integrating Machine Learning and Deep Learning for Early Detection of Acute Respiratory Distress Syndrome Using Clinical and Radiological Data. SGS-Engineering & Sciences, 1(4).
[9] Mall, P. K. (2025). Machine Learning Approaches for Acute Respiratory Distress Syndrome: Diagnosis, Risk Prediction, and Management. SGS-Engineering & Sciences, 1(1).
[10] Henry, B. M., & Lippi, G. (2020). Poor survival with extracorporeal membrane oxygenation in acute respiratory distress syndrome (ARDS) due to coronavirus disease 2019 (COVID-19): pooled analysis of early reports. Journal of critical care, 58, 27.
[11] Rashid, M., Ramakrishnan, M., Chandran, V. P., Nandish, S., Nair, S., Shanbhag, V., & Thunga, G. (2022). Artificial intelligence in acute respiratory distress syndrome: a systematic review. Artificial intelligence in medicine, 131, 102361.
[12] Rubulotta, F., Bahrami, S., Marshall, D. C., & Komorowski, M. (2024). Machine learning tools for acute respiratory distress syndrome detection and prediction. Critical care medicine, 52(11), 1768-1780.
[13]Rubulotta, F., Bahrami, S., Marshall, D. C., & Komorowski, M. (2024). Machine learning tools for acute respiratory distress syndrome detection and prediction. Critical care medicine, 52(11), 1768-1780.

[14] Villar, J., González-Martín, J. M., Hernández-González, J., Armengol, M. A., Fernández, C., Martín-Rodríguez, C., ... & Slutsky, A. S. (2023). Predicting ICU mortality in acute respiratory distress syndrome patients using machine learning: the predicting outcome and STratifiCation of severity in ARDS (POSTCARDS) study. Critical care medicine, 51(12), 1638-1649.
[15] Pai, K. C., Chao, W. C., Huang, Y. L., Sheu, R. K., Chen, L. C., Wang, M. S., ... & Chan, M. C. (2022). Artificial intelligence–aided diagnosis model for acute respiratory distress syndrome combining clinical data and chest radiographs. Digital Health, 8, 20552076221120317.
[16] Chiumello, D., Coppola, S., Catozzi, G., Danzo, F., Santus, P., & Radovanovic, D. (2024). Lung imaging and artificial intelligence in ARDS. Journal of Clinical Medicine, 13(2), 305.
[17] Lazzarini, N., Filippoupolitis, A., Manzione, P., & Eleftherohorinou, H. (2022). A machine learning model on Real World Data for predicting progression to Acute Respiratory Distress Syndrome (ARDS) among COVID-19 patients. PLoS One, 17(7), e0271227.
[18] Wei, S., Zhang, Y., Dong, H., Chen, Y., Wang, X., Zhu, X., ... & Guo, S. (2023). Machine learning-based prediction model of acute kidney injury in patients with acute respiratory distress syndrome. BMC Pulmonary Medicine, 23(1), 370.
[19] Mu, S., Yan, D., Tang, J., & Zheng, Z. (2025). Predicting mortality in Sepsis-Associated acute respiratory distress syndrome: A machine learning approach using the MIMIC-III database. Journal of Intensive Care Medicine, 40(3), 294-302.
[20] Ding, N., Nath, T., Damarla, M., Gao, L., & Hassoun, P. M. (2024). Early predictive values of clinical assessments for ARDS mortality: a machine-learning approach. Scientific reports, 14(1), 17853.