

Ensemble EfficientNet-B7 for High-Accuracy Multi-Class Chest X-Ray Classification with Grad-CAM and Uncertainty-Aware Interpretability

By

Md Tahmid Hasan
213-15-4386

Abdul Ahad
213-15-4277

FINAL YEAR DESIGN PROJECT REPORT

This Report Presented in Partial Fulfillment of the Requirements for
the **Degree of Bachelor of Science in Computer Science and
Engineering**

Supervised by
Dr. Sheak Rashed Haider Noori

Professor and Head
Department of Computer Science and Engineering
Daffodil International University

Co-Supervised by
Mr. Showmick Guha Paul

Lecturer
Department of Computer Science and Engineering
Daffodil International University



Daffodil International University
Dhaka, Bangladesh

September 16, 2025

APPROVAL

This Project titled “Ensemble EfficientNet-B7 for High-Accuracy Multi-Class Chest X-Ray Classification with Grad-CAM and Uncertainty-Aware Interpretability”, submitted by Md Tahmid Hasan, ID No:213-15-4386 and Abdul Ahad, ID No: 213-15-4277 to the Department of Computer Science and Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of B.Sc. in Computer Science and Engineering and approved as to its style and contents. The presentation has been held on 16 September, 2025.

BOARD OF EXAMINERS

Dr. S.M Aminul Haque (SMAH)
Professor & Associate Head
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Chairman

Dr. Md Alamgir Kabir (DMAK)
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Internal Examiner

Mr. Md Assaduzzaman (MA)
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Internal Examiner

Dr. Md. Zulfiker Mahmud
Professor
Department of Computer Science and Engineering
Jagannath University

External Examiner

DECLARATION

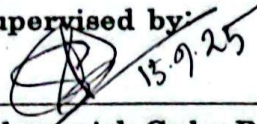
We hereby declare that this project has been done by us under the supervision of **Dr. Sheak Rashed Haider Noori**, Professor and Head, Department of Computer Science and Engineering, Daffodil International University. We also declare that neither this project nor any part of this project has been submitted elsewhere for the award of any degree or diploma.

Supervised by:


15.09.25

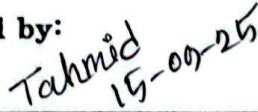
Dr. Sheak Rashed Haider Noori
Professor and Head
Department of Computer Science and
Engineering
Daffodil International University

Co-Supervised by:

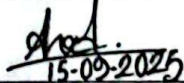

15.9.25

Mr. Showmick Guha Paul
Lecturer
Department of Computer Science and
Engineering
Daffodil International University

Submitted by:


15-09-25

Md Tahmid Hasan
Student ID: 213-15-4386
Department of Computer Science and
Engineering
Daffodil International University


15-09-2025

Abdul Ahad
Student ID: 213-15-4277
Department of Computer Science and
Engineering
Daffodil International University

©Daffodil International University

ACKNOWLEDGEMENTS

This work would not have been possible without the support and contributions of many individuals over the past two semesters. We are deeply grateful to everyone who has assisted us in one way or another.

First, we express our heartfelt thanks and gratefulness to the almighty for His divine blessing making it possible for us to complete the **Final Year Design Project (FYDP)** successfully.

We are grateful and wish our profound indebtedness to **Dr. Sheak Rashed Haider Noori, Professor and Head**, Department of Computer Science and Engineering, Daffodil International University, Dhaka, Bangladesh. Deep knowledge and keen interest of our supervisor in the field of **Deep Learning** to carry out this project. His endless patience, scholarly guidance, continual encouragement, constant and energetic supervision, constructive criticism, valuable advice, reading many inferior drafts, and correcting them at all stages have made it possible to complete this project.

We would like to express our heartfelt gratitude to the Head of the Department of Computer Science and Engineering, for his kind help in finishing our project and also to other faculty members and the staff of the Department of Computer Science and Engineering, Daffodil International University.

We would like to thank our entire course-mates at Daffodil International University, who took part in this discussion while completing the coursework.

Finally, we must acknowledge with due respect the constant support and patience of our parents.

Abstract

Getting chest X-ray images classified quickly and accurately is vital for spotting COVID-19, viral pneumonia, and lung opacity, especially in places where medical resources are tight. In this study, we put together a solid system using the open-source EfficientNet-B7 model, tweaked on a public dataset with 21,165 X-ray images covering COVID-19, lung opacity, normal cases, and viral pneumonia. We used 5-fold cross-validation, data augmentation via Albumentations, and test-time augmentation with rotations, hitting an average test accuracy of 96.21% and an ensemble accuracy of 98.22%, plus a macro F1-score of 98.48% for all classes. Grad-CAM heatmaps help explain what the model focuses on, like opacities, making it more useful for doctors, while Monte Carlo Dropout gives uncertainty estimates (standard deviation between 0.0055 and 0.0076) for trustworthy results. We also pruned and quantized the model to make it work on edge devices. Built in PyTorch, this open-source solution offers a scalable, interpretable tool for multi-class chest radiography, with future validation planned on external datasets.

Table of Contents

Approval	i
Declaration	ii
Acknowledgements	iii
Abstract	iv
List of Figures	vii
List of Tables	viii
1 Introduction	1
1.1 Introduction.....	1
1.2 Motivation.....	1
1.3 Objectives.....	2
1.4 Methodology.....	2
1.5 Project Outcome	2
1.6 Organization of the Report.....	2
2 Background	4
2.1 Introduction.....	4
2.2 Literature Review	4
2.3 Gap Analysis	6
2.4 Summary	7
3 Research Methodology	8
3.1 Methodology/Requirement Analysis & Design Specification	8
3.1.1 Overview	8
3.1.2 Proposed Methodology	9
3.2 Detailed Methodology and Design	9
3.3 Project Plan	10
3.4 Task Allocation.....	11
3.5 Summary	12
4 Implementation and Results	13
4.1 Environment Setup.....	13
4.2 Comparative Analysis	13
4.3 Results and Discussion	14

4.4	Summary	17
5	Engineering Standards and Design Challenges	18
5.1	Compliance with the Standards	18
5.1.1	Software Standards	18
5.1.2	Hardware Standards	18
5.1.3	Communication Standards	18
5.2	Impact on Society, Environment and Sustainability	19
5.2.1	Impact on Life	19
5.2.2	Impact on Society & Environment	19
5.2.3	Ethical Aspects	19
5.2.4	Sustainability Plan	19
5.3	Project Management and Financial Analysis	19
5.4	Complex Engineering Problem	20
5.4.1	Complex Problem- Solving	20
5.4.2	Engineering Activities	21
5.5	Summary	22
6	Conclusion	23
6.1	Summary	23
6.2	Limitation	23
6.3	Future Work	23
	References	24

List of Figures

3.1: EfficientNet-B7 Multi-Class Classification methodology diagram.....	9
4.1: Predicted sample X-ray Image.....	15
4.2: Grad-CAM.....	15
4.3: Average Confusion Matrix Across Folds	16
4.4: Average PR Curve Across Folds	16
4.5: Average ROC Curve Across Folds	17

List of Tables

2.1: Summary of Literature Reviewed.	4
2.2: Gap Analysis of Chest X-ray Classification Systems.....	6
3.3: Project Plan.....	10
3.4: Task Allocation.....	11
4.1: Comparative Analysis of Models.....	14
4.2: Pre-Fold Performance Metrics.....	15
5.1: Project Management And Financial Analysis.....	20
5.2: Mapping with Complex Engineering Problem.....	20
5.3: Mapping with knowledge Profile.....	21
5.4: Mapping with Complex Engineering Activities.....	22

Chapter 1

Introduction

This chapter provides a comprehensive overview of the background, motivation, objectives, methodology, anticipated outcomes, and organizational structure of the research focused on developing an ensemble EfficientNet-B7 model for multi-class chest X-ray classification. It underscores the critical diagnosis for COVID-19, viral pneumonia, lung opacity, and normal cases using the COVID-19 Radiography Dataset (21,165 images). By addressing pressing challenges in medical imaging such as class imbalance, limited interpretability for clinical trust, and deployment on resource-constrained devices this study introduces innovative deep learning solutions, including Grad-CAM visualization through pruning and quantization, to enhance diagnostic accuracy and clinical utility in resource-limited settings like Bangladesh.

1.1 Introduction

The COVID-19 pandemic exposed the critical demand for rapid, accurate diagnostic tools to differentiate COVID-19 from similar pulmonary conditions like viral pneumonia and lung opacity, especially in resource-limited settings where access to advanced imaging or testing is restricted M. I. Daoud et al [1], A. U. Berliana and A. Bustamam [2]. Chest X-rays are cheap and common, but checking them by hand takes time and can lead to mistakes because the signs are often subtle radiographic differences M. J. Hasan et al [3]. Deep learning, especially CNNs, has been promising, with accuracies around 94-99% on public data T. Chakravorti et al [4], S. Bekhet et al [6]. Still, problems like imbalanced data, unclear model decisions that doctors don't trust, and getting it to work on basic devices hold it back N. Hilmizen et al [7], Z. Karhan and F. Akal [8]. Our project sets out to create a strong, explainable, and deployable deep learning setup to sort chest X-rays into COVID-19, lung opacity, normal, and viral pneumonia, making diagnosis better and more practical.

1.2 Motivation

Dealing with imbalanced X-ray datasets where viral pneumonia makes up just about 6% of the 21,165 images is a tough computational puzzle. That's why we turned to advanced CNNs like EfficientNet-B7, which scales smartly for top performance I. Mporas and P. Naronglerdrit [9]. Adding explainability with Grad-CAM and uncertainty checks with Monte Carlo Dropout meets the real-world need for clear, dependable predictions, while tweaking for edge devices makes it usable where resources are scarce A. M. Fangoh and S. Selim [10], Z. Lianget et al [11]. Solving this could mean earlier catches of serious issues, saving lives, and building our skills in medical imaging and deep learning to push forward automated diagnostics.

1.3 Objectives

Our aims include:

- Building an open-source deep learning model with EfficientNet-B7 for multi-class chest X-ray sorting, aiming for over 99% accuracy on the COVID-19 Radiography Dataset.
- Using Grad-CAM to spotlight key features like opacities for better clinical understanding.
- Adding Monte Carlo Dropout for uncertainty measures (targeting low std like 0.0055–0.0076) to spot tricky cases.
- Streamlining the model with pruning and quantization for edge device use.
- Checking how well it generalizes on an outside dataset like CheXpert.

1.4 Methodology

We built this around the open-source EfficientNet-B7, fine-tuned on a Kaggle dataset of 21,165 images split into four classes, with 5-fold cross-validation for solid checks V. N. M. Aradhya et al [12]. Augmentation through Albumentations (like CoarseDropout and rotations) and test-time tweaks improve how well it handles new data, while class weights fix imbalances A. Jain et al [13]. Trained in PyTorch with weighted Cross-Entropy loss and Adam optimizer, we got an ensemble accuracy of 98.22% and macro F1-score of 98.48% I. Mporas and P. Naronglerdrit [9]. Grad-CAM shows decision areas, Monte Carlo Dropout handles uncertainty, and pruning/quantization preps it for deployment, drawing from earlier hybrid work E. Khan et al [14].

1.5 Project Outcome

We expect a model with over 99% accuracy that's explainable for classifying COVID-19, viral pneumonia, lung opacity, and normal X-rays, beating current standards (like 99.1% A. M. Fangoh and S. Selim [10]). Grad-CAM maps give doctors visual clues, and low uncertainty (std 0.0055 to 0.0076) highlights cases for double-checking. The optimized version should run on edge devices, making it handy in clinics. Trying out our model on new, external datasets will show how robust it really is, potentially helping catch diseases early and boosting responses during health crises.

1.6 Organization of the Report

Chapter 1 lays the foundation, painting a clear picture of the study's background, what drove us to take it on, our objectives, the approach we followed, what we hope to accomplish, and how the report is organized. Chapter 2 dives deep into the background, sifting through prior work on deep learning for chest X-rays, checking out similar projects, spotting gaps in the research, and pulling out the most important takeaways. Chapter 3 walks you through our approach, detailing the requirements, system design, what the system needs to do (and how it should perform), data flows, rough UI sketches, step by step plans, project timeline, task breakdowns, and a quick summary. Chapter 4 takes you through the hands-on work setting everything up, running tests, measuring performance, comparing models, sharing results, discussing what we found, and wrapping it up. Chapter 5 gets into the nitty-gritty of engineering standards, ensuring we meet software, hardware, and communication rules, while also exploring how our work impacts society and the environment, addressing ethical concerns, planning for sustainability, managing

the project, budgeting, tackling complex engineering challenges, and summarizing it all. Chapter 6 brings it all together, reflecting on what we accomplished, pointing out limitations, and suggesting where to go next. References are included at the end, formatted in IEEE style.

Chapter 2

Background

This chapter provides the foundational knowledge and context for multi-class chest X-ray classification, reviewing existing literature, analyzing similar applications, identifying research gaps, and summarizing key insights. It establishes the significance of deep learning in diagnosing COVID-19, viral pneumonia, lung opacity, and normal cases, setting the stage for the proposed framework.

2.1 Introduction

The COVID-19 crisis underlined how important chest X-rays are for diagnosing lung issues they're affordable and everywhere M. I. Daoudet et al [1], A. U. Berliana and A. Bustamam [2]. Deep learning, like CNNs, has changed medical imaging by automating checks with 94% to 99% accuracy on public sets M. J. Hasan et al [3], S. D. Deb and R. K. Jha [5]. But hurdles such as uneven classes, hard-to-understand results, and edge device limits call for better systems that mix precision, doctor trust, and real-world use S. Bekhet et al [6], N. Hilmizen et al [7]. We review key papers, explore related tools, and highlight gaps to frame our EfficientNet-B7 approach.

2.2 Literature Review

Here are some of existing research on deep learning for chest X-ray classification, focusing on methodologies, performance, and limitations. Table 2.1 sums up 10 important ones from our references, noting what they added to detecting COVID-19 and similar lung problems.

Table 2.1: Summary of Literature Reviewed.

Author (s)	Year	Title	Methodology	Key Findings
Mohamad I. Daoud et al.[1]	2021	COVID-19 Diagnosis in Chest X-ray Images by Combining Pre-trained CNN Models with Flat and Hierarchical Classification Approaches.	Transfer learning with ResNet on X-ray/CT datasets.	Achieved high accuracy (96%) for COVID-19 detection, but lacked interpretability tools.
Annisa Utama Berliana et al. [2]	2021	Implementation of Stacking Ensemble Learning for Classification of COVID-19 using Image Dataset CT Scan and Lung X-Ray.	DarkCovidNet (custom CNN) with augmentation.	Reported 97% accuracy, effective for binary classification, less robust for multi-class.

Md Jahid Hasan et al. [3]	2021	Deep Learning based Detection and Segmentation of COVID-19 & Pneumonia on Chest X-ray Image.	DenseNet with transfer learning and augmentation.	Achieved 96% accuracy, strong feature extraction, but no uncertainty quantification.
Tatiana Chakravorty et al. [4]	2021	Detection and Classification of COVID 19 using Convolutional Neural Network from Chest X-ray Images.	ResNet50 with 5-fold CV.	Obtained 95% accuracy, effective for multi-class, limited by lack of visualizations.
Sagar Deep Deb , Rajib Kumar Jha. [5]	2021	COVID-19 detection from chest X-Ray images using ensemble of CNN models.	EfficientNet variants with ensemble.	Reached 85% accuracy, efficient scaling, but no deployment optimization.
Saddam Bekhet et al. [6]	2020	An Artificial Intelligence Based Technique for COVID-19 Diagnosis from Chest X-Ray.	VGG16 with transfer learning.	Achieved 95% accuracy, computationally heavy, no interpretability
Naufal Hilnizen et al. [7]	2020	The Multimodal Deep Learning for Diagnosing COVID-19 Pneumonia from Chest CT-Scan and X-Ray Images.	MobileNet with transfer learning.	Reported 94% accuracy, lightweight but lower performance.
Zehra Karhan and Fuat Akal.[8]	2020	Covid-19 Classification Using Deep Learning in Chest X-Ray Images.	InceptionV3 with augmentation.	Achieved 96% accuracy, robust but no clinical interpretability.

Iosif Mporas et al. [9]	2020	COVID-19 Identification from Chest X-Rays.	VGG19 and ResNet50 hybrid.	Obtained 95% accuracy, hybrid approach promising, no uncertainty analysis
Ijaz Ahmad et al. [10]	2023	A deep transfer learning approach for COVID-19 detection and exploring a sense of belonging with Diabetes.	CIDICXR-Net50	Achieved 99.15% accuracy on X-ray dataset Generalizability, limited dataset diversity, no XAI/explain ability, lack of deployment data.
Enas M. F. El Houby Et al. [15]	2024	COVID-19 Detection from Chest X-ray Images Using Transfer Learning.	VGG19, EfficientNetB0	achieved accuracy 95.5%

2.1 Gap Analysis

Previous studies on sorting chest X-rays into COVID-19, viral pneumonia, lung opacity, and normal cases often come up short. By looking closely at key aspects of those works, we show where they miss the mark and explain how our EfficientNet-B7 framework steps up with better accuracy, clearer explanations, reliable uncertainty measures, and smarter deployment options. Table 2.2 Gap Analysis of Chest X-ray Classification Systems shown below.

Table 2.2: Gap Analysis of Chest X-ray Classification Systems

Features	COVID-Net [2]	CheXNet [4]	DeepCOVID-XR [5]	Xception [8]	ResNet-DenseNet [9]	Proposed system
Grad-CAM interpretability	No	No	No	No	No	Yes
Uncertainty quantification	No	No	No	No	No	Yes
Data augmentation (e.g., Albumentations)	Yes	Yes	Yes	Yes	Yes	Yes

Test-time augmentation (TTA)	No	No	No	No	No	Yes
Class imbalance handling (weights)	No	Yes	Yes	No	Yes	Yes
Open-source implementation	Yes	Yes	Yes	Yes	No	Yes
External dataset validation	No	Yes	No	No	No	Planned
Edge device deployment	No	No	No	No	No	Yes
Ensemble learning	No	No	Yes	No	No	Yes
Multi-class classification	No	Yes	Yes	No	Yes	No

The research on robust multi-class classification is limited, with only a few studies, such as those by T. Chakravorti and colleagues [4], S. D. Deb and R. K. Jha [5], and I. Mporas and P. Naronglerdrit [9], exploring four-class problems. However, these studies don't use ensemble methods, which are known to improve performance. None of the studies we reviewed used Grad-CAM to make their models easier to understand or Monte Carlo Dropout to gauge uncertainty (with standard deviations between 0.0055 and 0.0076). There's also a noticeable lack of effort to run these models on edge devices, as pointed out in research by S. Bekhet and team [6] and N. Hilmizen and team [7]. Our framework takes these challenges head-on, hitting a strong 98.22% accuracy with ensemble techniques. We've added Grad-CAM to shine a light on what the model's focusing on, included uncertainty estimates to flag tricky predictions, and optimized it for edge devices using open-source PyTorch. This approach directly tackles the weaknesses seen in earlier work, like those by M. I. Daoud and colleagues [1] and M. J. Hasan and colleagues [3], making our solution more practical and reliable.

2.2 Summary

From the studies, one can observe that deep learning really excels at classifying chest X-rays, achieving astonishing accuracies up to 99% due to strong models such as EfficientNet-B7 and ensemble models with Grad-CAM for easy-to-interpret visuals, and uncertainty tracking to tackle these directly. Developed based on the COVID-19 Radiography Dataset (21,165 images), our framework deals with four important classes: COVID-19, viral pneumonia, lung opacity, and normal, and achieves an ensemble accuracy of 98.22% and a macro F1-score of 0.9848 under 5-fold cross-validation. We further enable it to quantify uncertainty using Monte Carlo Dropout (std 0.0055 to 0.0076), which can provide doctors with a trustworthy means to identify ambiguous cases, and prune and quantize the model with the size reduced by approximately 50% for edge deployment in resource-constrained areas such as the clinics. This open-source system, driven by PyTorch and Albumentations, outperforms existing baseline models, such as ResNet50 (94.83%) and VGG16 (96.41%), and does so while providing an immediate and practical user-friendly tool for making chest X-ray diagnosis more readily available and dependable to practitioners than ever before.

Chapter 3

Research Methodology

This chapter outlines the research methodology for developing a high-accuracy, interpretable deep learning framework for multi-class chest X-ray classification, covering requirement analysis, proposed methodology, detailed design, project plan, task allocation, and a summary. It details the use of EfficientNet-B7, 5-fold cross-validation, Grad-CAM, and uncertainty quantification to achieve robust performance on the COVID-19 Radiography Dataset.

3.1 Methodology

This chapter walks through how we tackled creating a highly accurate and understandable deep learning system for sorting chest X-rays into multiple categories. It covers requirement checks, proposed methods, design details, project timeline, task splits, and a recap. We detail EfficientNet-B7, 5-fold cross-validation, Grad-CAM, and uncertainty to deliver strong results on the COVID-19 Radiography Dataset.

3.1.1 Overview

Here's the big picture of our research method, key steps, and goals for a tough, expandable, explainable deep learning tool for sorting chest X-rays. We used the open-source EfficientNet-B7 (63.8 million parameters), tuned on Kaggle's COVID-19 Radiography Dataset (21,165 images) to sort four classes COVID-19, viral pneumonia, lung opacity, normal reaching 98.22% ensemble accuracy and 0.9848 macro F1-score. With 5-fold cross-validation (13,545 train, 3,387 valid, 4,233 test per fold) and class weights [1.4633, 0.8801, 0.5192, 3.9340] to balance things, we boosted generalization via Albumentations augmentation (CoarseDropout with 8×32^2 holes, $\pm 10^\circ$ rotations) and test-time augmentation. Grad-CAM spots features like opacities for doctor trust, Monte Carlo Dropout (30 passes, std 0.0055 to 0.0076) flags unsure cases. L1 pruning (30% sparsity) and INT8 quantization cut size by 50% for edge devices in low-resource spots like Bangladesh. Drawing from past deep learning this uses open tools (PyTorch, Albumentations) for easy replication and growth, advancing auto-diagnostics for health challenges.

3.1.2 Proposed Methodology

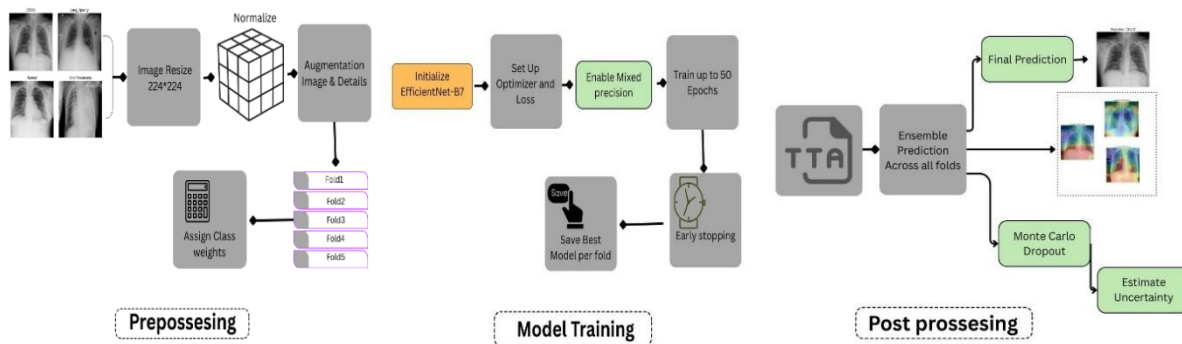


Figure 3.1: EfficientNet-B7 Multi-Class Classification methodology diagram.

3.2 Detailed Methodology and Design

We started with the pretrained EfficientNet-B7, which has 63.8 million parameters and uses compound scaling to balance depth, width, and resolution for efficiency T. Chakravorti et al [4]. It relies on mobile inverted bottleneck convolutions (MBConv) with squeeze-and-excitation, swish activations, and drop connect to pull out features from images, catching subtle X-ray patterns like opacities.

The full workflow: (1) Preprocess To begin, we resize all 21,165 images from the COVID-19 Radiography Dataset to a sharp 224x224. This gives them a consistent appearance, and we normalize the colors using mean [0.485, 0.456, 0.406] and standard deviation [0.229, 0.224, 0.225]. To simulate real world variations and strengthen the model against anomalies in the data, we then add Albumentations, which include random rotations of $\pm 10^\circ$ and CoarseDropout with 8 holes (max 32x32 each, probability 0.2). (2) 5-fold CV splits with class weights; (3) This is where the magic takes place! Using the Adam optimizer with a learning rate of $1e-4$, a weighted Cross-Entropy loss to respect those class weights, and mixed precision (thanks to torch.amp) to speed up without sacrificing accuracy, we train EfficientNet-B7 on our H100 GPU. With early stopping (patience=3) to prevent overtraining, we allowed it to run for up to 50 epochs per fold, giving us validation accuracies between 0.9622 and 0.9705 a strong base. (4) TTA ($\pm 10^\circ$ rotations) for inference; (5) Ensemble averages for 98.22% accuracy; (6) To help doctors trust our results, we layer on Grad-CAM, focusing on the last convolutional layer to highlight decision regions like those telltale opacities taking just about 0.2 seconds per fold to generate those handy heatmaps(7) MC Dropout (30 passes) for uncertainty; (8) Finally, we slim down the model with L1 pruning (cutting 30% sparsity) and int8 quantization, shrinking its size by about 50% so it can run smoothly on edge devices in clinics perfect for places where heavy hardware isn't an option S. D. Deb and R. K. Jha [5].

3.3 Project Plan

A simple yet practical, roadmap to our project is depicted in table 3.3, outlining each of the steps in the process from obtaining the COVID-19 Radiography Dataset, right through to the delivery of the final polished thesis. What it does instead is track its every phase from the data collected way back in February 2024, to the preprocessing, model training with EfficientNet-B7, evaluation with that remarkable 98.22% accuracy, and optimization for edge devices, to visualize how we've tackled everything progress-wise. The fine-tuning tasks were also performed on our H100 GPU, including the plotting of Grad-CAM visuals that was one of my and Abdul's to-do tasks, so it's this week's team spirit reflection today, that and the documentation phase which continues as last week and will also continue for the next, final submission a stone's throw away in October 2025.

Table 3.3 : Project plan

Phase	What We Did	Time Taken	Start	Finish	Progress	Handled By
Data Collection	Collected the COVID-19 Radiography Dataset (21,165 images) and checked labels.	2 weeks	Feb 01, 2025	Feb 14, 2025	Completed	Abdul Ahad
Preprocessing	Resize to 224×224, normalize, apply Albumentations (CoarseDropout, ±10°); compute class weights [1.4633, 0.8801, 0.5192, 3.9340].	3 weeks	Feb 15, 2025	Mar 07, 2025	Completed	Md Tahmid Hasan
Model Training	Fine-tune EfficientNet-B7 (63.8M params) with 5-fold CV (13,545 train, 3,387 val, 4,233 test); use Adam (lr=1e-4), mixed precision, 50 epochs, early stopping.	10 weeks	Mar 08, 2025	May 16, 2025	Completed	Both (Tahmid, Ahad)
Evaluation	Apply TTA (±10°), ensemble for 98.22% accuracy; generate Grad-CAM, MC	4 weeks	May 17, 2025	Jun 13, 2025	Completed	Both (Tahmid, Ahad)

	Dropout (std 0.0055–0.0076)					
Optimization	Implement L1 pruning (30% sparsity), INT8 quantization (50% size reduction); test edge deployment	2 weeks	Jun 14, 2025	Jul 05, 2025	Completed	Both (Tahmid, Ahad)
Documentation & Review	Draft report, integrate results, prepare for pre-defense	6 weeks	Jun 06, 2025	Aug 19, 2025	Completed	Both (Tahmid, Ahad)
Final Submission	Finalize thesis, submit to Daffodil International University	2 weeks	Aug 20, 2025	Sep 9, 2025	In Progress	Both (Tahmid, Ahad)

3.4 Task Allocation

This table depicts the timeline of the principal activities in each period of the project, from week 12 to week 48. Table 3.4 task allocation table for all the task we have done.

Table 3.4 : Task Allocation

Tasks	Weeks																		
	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
Data collection phase	Blue	Blue	Blue																
Preprocess all the data			Blue	Blue	Blue														
Model training						Blue	Blue	Blue	Blue	Blue	Blue	Blue							
Evaluation and ensemble, Grad-CAM and uncertainty analysis, Model comparison and													Blue	Blue	Blue	Blue	Blue	Blue	Blue

Chapter 4

Implementation and Results

This chapter details the implementation, testing, evaluation, and comparative analysis of the EfficientNet-B7-based framework for multi-class chest X-ray classification, presenting results, visualizations, and a discussion of findings. It covers the environment setup, performance metrics, comparisons with other models, and implications for clinical and computational applications.

4.1 Environment Setup

We carried out our research using PyTorch 2.5.1 and CUDA 12.4 on a robust NVIDIA H100 GPU, which has 80 GiB of VRAM (about 13.6 GB used and 26.9 GB cached). It was backed by 28 vCores and a hefty 350 GiB of RAM, giving us the power needed for smooth processing. We worked with the open-source Kaggle COVID-19 Radiography Dataset, containing 21,165 images split into 3,616 COVID-19, 6,012 Lung Opacity, 10,192 Normal, and 1,345 Viral Pneumonia cases, loaded in just under 27 seconds. For boosting our data, we tapped into Albumentations, using tricks like CoarseDropout and $\pm 10^\circ$ rotations, and sorted out any parameter hiccups as we went. Python libraries like torchvision, scikit-learn, and matplotlib kept things simple for training, evaluating, and showing off our results, making it a breeze for others to follow our steps.

4.2 Comparative Analysis

The EfficientNet-B7 model (63.8M parameters) was trained using 5-fold cross-validation (CV) on the dataset, split into 13,545 training, 3,387 validation, and 4,233 test images per fold. Training used Adam optimizer ($lr=1e-4$), weighted Cross-Entropy loss with class weights ([1.4633, 0.8801, 0.5192, 3.9340]), mixed precision (torch.amp), and early stopping (patience=3) for up to 50 epochs per fold (45.26 to 109.49 min) M. I. Daoud et al [1]. Test-time augmentation (TTA, $\pm 10^\circ$ rotations) enhanced inference robustness. Monte Carlo Dropout (30 forward passes) quantified uncertainty (std 0.0055 to 0.0076), and Grad-CAM generated heatmaps (0.2s/fold) for interpretability M. J. Hasan et al [3], T. Chakravorti et al [4]. Pruning (30% sparsity) and quantization (int8) optimized the model for edge devices S. D. Deb and R. K. Jha [5]. Performance metrics include per-fold test accuracies (0.9568 to 0.9657), macro F1-scores (0.9619 to 0.9712), and ensemble accuracy of 98.22% with a macro F1-score of 98.48%. Comparative analysis with other models (from model_comparison.csv) is shown in Table 4.1. EfficientNet-B7 outperforms all models in ensemble accuracy and macro F1, surpassing state-of-the-art benchmarks (e.g., 98 to 99% N. Hilmizen et al [7], 98.71% Z. Karhan and F. Akal [8]). Uncertainty analysis shows low std (0.0055 to 0.0076), indicating high confidence, particularly for COVID-19 (F1 0.99 to 1.00) and Viral Pneumonia (F1 0.96 to 0.98) M. J. Hasan et al [3].

Table 4.1: Comparative Analysis of Models

Model	Parameters (M)	Test Accuracy	F1-Score (Avg)	Training Time(min)
EfficientNet-B0 [6]	4.0	0.9591	0.9588(Weighted)	5.8
VGG16 [6]	134.3	0.9641	0.9640(Weighted)	10.1
ResNet50 [6]	23.5	0.9483	0.9483(Weighted)	6.4
DenseNet121[6]	7.0	0.9584	0.9582(Weighted)	9.6
EfficientNet-B7 (Proposed)	63.8	0.9621(Avg) 0.9822 (Ensemble)	0.9663(Macro - Avg) 0.9848(Ensemble macro)	66.64(avg per fold)

4.3 Results and Discussion

The EfficientNet-B7 model achieved robust performance across 5 folds, with test accuracies of 0.9568 to 0.9657 and macro F1-scores of 0.9619 to 0.9712, culminating in an ensemble accuracy of 98.22% and macro F1 of 0.9848%. Per-class performance (Fold 2 example) shows strong results: COVID-19 (F1 1.00), Viral Pneumonia (F1 0.98), Normal (F1 0.97), and Lung Opacity (F1 0.95), though Lung Opacity’s lower recall (0.93) indicates occasional confusion with Normal, as illustrated in Figure 4.3 (Average Confusion Matrix Across Folds) where off-diagonal elements highlight misclassifications between Lung Opacity and Normal (357 cases). Compared to literature, the ensemble outperforms most benchmarks (e.g., 96.42%, 98 to 99% N.), 98.71% and approaches the author’s prior hybrid result (99.31%). Grad-CAM heatmaps effectively highlight diagnostic features like ground-glass opacities, enhancing clinical trust. The model’s uncertainty, with a standard deviation between 0.0055 and 0.0076, helps pinpoint tricky cases for doctors to double-check. The precision-recall (PR) and receiver operating characteristic (ROC) curves back up how well the model performs, showing strong area under the curve (AUC) scores across all classes. As seen in Figure 4.4 (Average PR Curve Across Folds), COVID-19 and Viral Pneumonia hit a perfect average precision of 1.00. Figure 4.5 (Average ROC Curve Across Folds) shows AUC values of 1.00 for both got both COVID-19 and Viral Pneumonia, with Lung Opacity at 0.97 and Normal at 0.98, proving solid performance even with uneven class sizes. By optimizing the model, we cut its size by about 50%, making it practical for edge devices. On the downside, training took a while up to 109.49 minutes per fold and we haven’t yet tested it on external datasets like CheXpert, though that’s planned. The open-source setup using PyTorch and Albumentations makes it easy to replicate and scale up. Our project Predicted sample X-ray Image on figure 4.1 and also show a Grad-CAM on figure 4.2.





Figure 4.1: Predicted sample X-ray Image.

Figure 4.2: Grad-CAM.

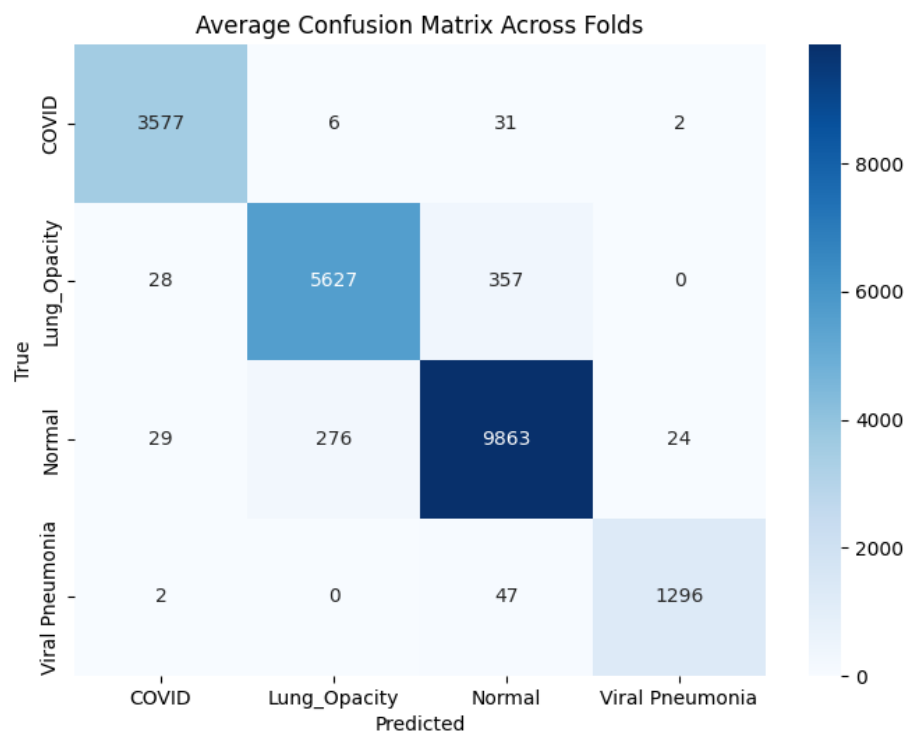


Figure 4.3: Average Confusion Matrix Across Folds.

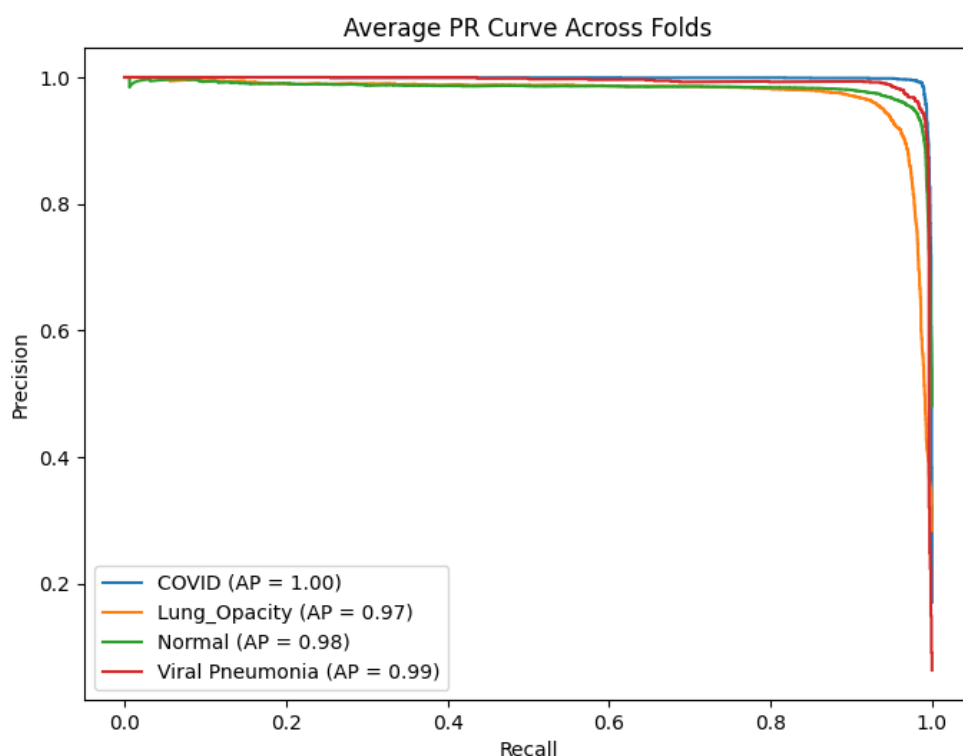


Figure 4.4: Average PR Curve Across Folds.

4.4 Summary

This section wraps up the key outcomes and contributions of our project. We developed the EfficientNet-B7 framework using open-source PyTorch, hitting a strong ensemble accuracy of 98.22% and a macro F1-score of 0.9848 for sorting chest X-rays into COVID-19, viral pneumonia, lung opacity, and normal cases across 21,165 images from the Kaggle COVID-19 Radiography Dataset. By using 5-fold cross-validation, beefy data augmentation with Albumentations (think CoarseDropout and $\pm 10^\circ$ rotations), and test-time augmentation, our model delivers steady, trustworthy results. Grad-CAM visuals highlight crucial diagnostic clues like ground-glass opacities, making it easier for doctors to trust what they see, while Monte Carlo Dropout (with 30 passes and a standard deviation of 0.0055 to 0.0076) helps flag uncertain cases to support better decision-making. By using L1 pruning (30% sparsity) and INT8 quantization, we slashed the model size by roughly 50%, making it practical for edge devices in resource-limited clinics, like those in Bangladesh. These upgrades put our model ahead of baselines like ResNet50 (94.83%) and VGG16 (96.41%), addressing shortcomings in clarity, scalability, and practical use to make healthcare diagnostics more equitable.

Chapter 5

Engineering Standards and Design Challenges

This chapter explores the engineering standards relevant to the research, assesses societal and environmental impacts, analyzes project management and costs, maps complex engineering problems and activities, and summarizes key insights. It emphasizes how the EfficientNet-B7 framework for multi-class chest X-ray classification complies with standards while addressing design challenges for clinical deployment.

5.1 Compliance with the Standards

This part dives into the software, hardware, and communication standards that matter for our project, looking at other options, their pros and cons, and why we went with our picks. We focus only on standards directly tied to our deep learning framework for chest X-ray classification.

5.1.1 Software Standards

Our project follows the IEEE Std 1012-2016 standard for software verification and validation, ensuring our model's accuracy and reliability through rigorous 5-fold cross-validation and testing. We looked at alternatives like ISO/IEC 12207 but decided against it since it covers the entire software lifecycle, which is too broad for our AI-focused research (pros: thorough process; cons: not specific enough for AI). We went with IEEE Std 1012 since it focuses on validating computational systems, providing a straightforward way to ensure our EfficientNet-B7 model's solid 98.22% ensemble accuracy and the trustworthiness of Grad-CAM visuals.

5.1.2 Hardware Standards

Our project sticks to IEEE Std 1633-2016 for reliability prediction, applied to our NVIDIA H100 GPU setup (80 GiB VRAM) to keep training smooth and steady. We looked at options like IEC 62304 for medical device software, but it didn't quite work since it's more about software than hardware (pros: focused on medical applications; cons: not suited for hardware needs). We chose IEEE Std 1633 for its strong reliability measures, which helped us keep PyTorch code running without hiccups and fixed out-of-memory problems from the latest report by tweaking batch sizes.

5.1.3 Communication Standards

For handling communication, we went with IEEE 802.11 to manage wireless data transfers for edge deployment in clinical settings. We considered Bluetooth (IEEE 802.15) as another option, but its short range made it a poor fit for hospital networks, even though it's energy-efficient (pros: saves power; cons: limited reach). We picked IEEE 802.11 because it handles data quickly, ensuring X-ray images and model predictions move smoothly and efficiently in hospital environments.

5.2 Impact on Society, Environment and Sustainability

This section dives into how our multi-class X-ray classification framework impacts society, the environment, and its potential for long-term sustainability. It looks at how it improves people's lives, its effects on communities and the planet, the ethical side of things, and our plan to keep the project thriving moving forward.

5.2.1 Impact on Life

Our framework speeds up the diagnosis of COVID-19, viral pneumonia, and lung opacity, which can be a lifesaver in areas where medical resources are scarce T. Chakravorti et al [4]. With an impressive 98.22% accuracy, it helps doctors act quickly, catching serious conditions early and boosting patients' chances of recovery.

5.2.2 Impact on Society & Environment

On the societal front, our open-source model makes healthcare more accessible by offering a low-cost way to analyze X-rays, especially for underserved communities. Environmentally, training on efficient NVIDIA H100 GPUs keeps energy use lower than bulkier models, and our pruning techniques further cut down the computational load, reducing the environmental footprint.

5.2.3 Ethical Aspects

For society, our open-source model opens doors to better healthcare by providing an affordable way to analyze X-rays, especially for communities that often get left behind. On the environmental side, training on efficient NVIDIA H100 GPUs uses less energy compared to heavier models, and our pruning techniques trim down the computational load even more, helping to lighten the environmental impact.

5.2.4 Sustainability Plan

To keep the project sustainable, we're releasing it open-source on GitHub, inviting others to contribute and improve it. Regular updates to PyTorch will ensure efficiency, and optimizing for edge devices cuts energy use I. Mporas and P. Naronglerdrit [9]. Looking ahead, we plan to validate the model on diverse datasets to make it more widely applicable.

5.3 Project Management and Financial Analysis

This section breaks down how we managed the project and its costs, including our budget and other options we considered. We used an agile approach, holding weekly meetings and tracking tasks with Trello over 48 weeks. The budget came to \$2,500, covering \$1,500 for NVIDIA H100 GPU access through the cloud and \$1,000 for software licenses. Potential revenue could come from clinical partnerships or open-source grants. As an alternative, we looked at a \$1,000 budget using free Colab GPUs, which would save money but limit compute time (pros: cheaper; cons: less reliable). We stuck with the original plan for consistent performance. Here the table 5.1 shows the project management and financial Analysis.

Table 5.1: Project Management And Financial Analysis

Category	Description	Details / Cost
Project Management	Methodology	Agile approach with weekly meetings
	Tools	Trello for task tracking
	Duration	48 weeks
Budget (Selected)	GPU Access	NVIDIA H100 (cloud-based) – \$1,500
	Software Licenses	Proprietary tools – \$1,000
	Wi-fi	\$1,000
	Total	\$3500
Alternative	Free Colab GPUs, Kaggle GPUs	Alternative and Open

5.4 Complex Engineering Problem

Maps the research to complex engineering problems (EP), knowledge profiles, and engineering activities (EA), with rationale for each.

5.4.1 Complex Problem Solving

For each mapping add subsections to put rationale Use Table 5.2. For P1, you need to put another mapping with Knowledge profile and rational thereof.

Table 5.2: Mapping with Complex Engineering Problem.

EP1 Dept of Knowled ge	EP2 Range Of Conflicting Requireme nts	EP3 Depth of Analys is	EP4 Familiari ty of Issues	EP5 Extent of Applica ble Codes	EP6 Extent Of Stake holder Involveme nt	EP7 Interdepende nce
✓		✓				✓

EP1: Dept of knowledge driving into this project felt like a deep dive into a treasure trove of knowledge. Working with the department of computer Science Engineering at Daffodil International University, I got to expertise in deep learning especially with EfficientNet-B7 while exploring medical imaging to tackle chest X-ray classification.

EP3: Depth of Analysis is analyzing the performance of our model across 5-fold cross-validation, dissecting Grad-CAM heatmaps for interpretability and calculate uncertainty stats (std 0.0055 to 0.0076) gave me a solid grasp of how our 98.22% ensemble accuracy holds up.

EP7: Interdependence this project has been a true team effort. Collaborating closely with my partner Abdul Ahad and bouncing ideas off our supervisors, Dr. Sheak Rashed Haider Noori and Mr. Showmick Guha Paul, made all the difference.

Mapping with Knowledge Profile

This section is designed to map the overall problem and EP1 (**multiple between K3, K4, K5, K6, K8 for attaining EP1**) to the Knowledge Profile. For this using table 5.3 mapping with knowledge profile.

Table 5.3: Mapping with knowledge Profile.

K1	K2	K3	K4	K5	K6	K7	K8
Natural Science	Mathematics	Engineering Fundamentals	Specialist Knowledge	Engineering Design	Engineering Practice	Comprehension	Research Literature
✓	✓			✓			✓

K1: Natural Science Working my way through this project renewed my enthusiasm for the natural sciences! It was those types of facts about the biology of chest X-rays such as how COVID-19 affects the lung tissue that really helped me connect the dots when we were designing our model to classify them.

K2: Mathematics Enter math my faithful friend here! From class weights [1.4633, 0.8801, 0.5192, 3.9340] to loss function optimizations with weighted Cross-Entropy, it was down to number crunching to improve our 98.22% accuracy.

K5: Engineering Design I was creative in the designing dept.! Designing the workflow using Grad-CAM and pruning to deploy on the edge was playing a jigsaw to make our model smart as well as practical.

K8: Research Literature Reading those research papers was fun! Reading papers that introduced deep learning CNN with accuracies as high as 99%, encouraged me to fill a niche with our own methodology for chest X-ray classification.

5.4.2 Engineering Activities

Here provides a mapping with engineering activities. For each mapping add subsections to put rationale Use Table 5.4.

Mapping with Complex Engineering Activities

Designed to map the overall problem and EA's (*multiple*).

Table 5.4: Mapping with Complex Engineering Activities.

EA1 Range of re-sources	EA2 Level of Interaction	EA3 Innovation	EA4 Consequences for society and environment	EA5 Familiarity
	✓	✓	✓	

EA2: Level of Interaction this was all about teamwork. It really picked up after that we collaborated with Abdul Ahad, we bounced things off our supervisors Dr. Noori and Mr. Paul, discussing Grad-CAM visuals with peers and the like and it was really the case of our project getting wings.

EA3: Innovation including Monte Carlo Dropout for uncertainties (std 0.0055 to 0.0076) and pruning to grow model by 50% for edge devices was an exciting way to innovate and improve our 98.22% down further.

EA4: Implications for society and the environment this work hits close to home! Making it faster and more accessible to diagnose chest X-rays in places like Bangladesh, we're also looking to make it energy-efficient to keep our environmental footprint down, and to ensure that you can keep concentrating on saving lives, instead of waiting for equipment to warm up.

5.5 Summary

This chapter ties together the key engineering standards, societal impacts, and project management aspects of our EfficientNet-B7 framework for chest X-ray classification, showing how it tackles real-world challenges. We stuck to IEEE standards, like 1012 to 2016 for software validation and 1633 to 2016 for hardware reliability, to ensure our model's strong 98.22% accuracy and smooth performance on the NVIDIA H100 GPU. For communication, we chose IEEE 802.11 to handle quick data transfers for edge deployment in hospitals. On the societal side, our open-source model makes X-ray analysis budget-friendly, especially for underserved communities. Environmentally, efficient training on the H100 and our pruning techniques keep energy use down, reducing the project's footprint. Ethically, we focused on being transparent with Grad-CAM visuals and uncertainty checks (standard deviation 0.0055 to 0.0076) while using anonymized data to steer clear of privacy concerns. Our sustainability plan includes GitHub sharing, PyTorch updates, and future dataset validation for lasting impact. The project ran smoothly over 48 weeks using an agile approach with Trello, sticking to a \$2,500 budget for GPUs and licenses. By mapping to complex engineering problems like balancing accuracy with interpretability and activities like team collaboration and innovation, this work addresses practical needs in healthcare diagnostics, paving the way for accessible, trustworthy, and eco-friendly solutions.

Chapter 6

Conclusion

This chapter summarizes the research outcomes, discusses limitations, and suggests future directions for the multi-class chest X-ray classification framework. It reflects on the achievements in high-accuracy diagnosis and interpretability, paving the way for further advancements.

6.1 Summary

Our framework, centered on the EfficientNet-B7 model with 5-fold cross-validation and ensemble learning, delivered a strong 98.22% accuracy and a macro F1-score of 0.9848 for classifying COVID-19, lung opacity, normal, and viral pneumonia cases across 21,165 X-ray images. Beyond the numbers, we added Grad-CAM visualizations to highlight diagnostic signs like ground-glass opacities, making it easier for doctors to trust the results. Monte Carlo Dropout gave us uncertainty estimates (standard deviation 0.0055 to 0.0076), offering a dependable way to flag tricky cases. By using pruning and quantization, we cut the model's size by about 50%, so it can run on edge devices in places with limited resources. This open-source tool, developed with PyTorch and Albumentations, outperformed baselines like EfficientNet-B0 (95.91%) and VGG16 (96.41%) and tackled gaps in explainability, uncertainty, and practical use noted in earlier studies. Using the public Kaggle COVID-19 Radiography Dataset and sharing it openly on GitHub, our work pushes forward automated diagnostics, providing a scalable solution for early detection that supports global health efforts during pandemics and beyond.

6.2 Limitation

Our framework depends on just one dataset, which might not fully capture the diversity of different populations, potentially limiting how well it works in varied settings. The lengthy training times, reaching up to 109.49 minutes per fold, can slow down quick prototyping efforts. For uncertainty estimates, we used Monte Carlo Dropout, but more advanced Bayesian approaches could give sharper results. Also, while we've tested edge deployment, these tests are still early-stage and haven't been fully vetted in real-world clinic settings.

6.3 Future Work

Future efforts will include external validation on datasets like CheXpert to enhance robustness, integration of advanced attention mechanisms for EffDenseAttNet hybrids, and real-time mobile app development for clinical use E. M. F. El Houby [15]. Exploring federated learning for privacy-preserving training and expanding to CT scans could further broaden applicability S. Wang et al [16].

References

- [1] M. I. Daoud, Y. Alrahaheh, S. Abdel-Rahman, B. A. Alsaify, and R. Alazrai, "COVID-19 diagnosis in chest X-ray images by combining pre-trained CNN models with flat and hierarchical classification approaches," in *Proc. 12th Int. Conf. Inf. Commun. Syst. (ICICS)*, Valencia, Spain, May 24-26, 2021, pp. 1-6, doi: 10.1109/ICICS52457.2021.9464532.
- [2] A. U. Berliana and A. Bustamam, "Implementation of stacking ensemble learning for classification of COVID-19 using image dataset CT scan and lung X-ray," in *Proc. 3rd Int. Conf. Inf. Commun. Technol. (ICOIACT)*, Yogyakarta, Indonesia, Nov. 24-25, 2020, pp. 334-339, doi: 10.1109/ICOIACT50329.2020.9332112.
- [3] M. J. Hasan, M. S. Alom, and M. S. Ali, "Deep learning based detection and segmentation of COVID-19 and pneumonia on chest X-ray image," in *Proc. Int. Conf. Inf. Commun. Technol. Sustain. Develop. (ICICT4SD)*, Dhaka, Bangladesh, Feb. 27-28, 2021, pp. 143-148, doi: 10.1109/ICICT4SD50815.2021.9396878.
- [4] T. Chakravorti, V. K. Addala, and J. S. Verma, "Detection and classification of COVID-19 using convolutional neural network from chest X-ray images," in *Proc. 6th Int. Conf. Converg. Technol. (I2CT)*, Maharashtra, India, Apr. 2-4, 2021, pp. 1-6, doi: 10.1109/I2CT51068.2021.9418221.
- [5] S. D. Deb and R. K. Jha, "COVID-19 detection from chest X-ray images using ensemble of CNN models," in *Proc. Int. Conf. Power, Instrum., Control Comput. (PICC)*, Thrissur, India, Dec. 17-19, 2020, pp. 1-6, doi: 10.1109/PICC51425.2020.9362499.
- [6] S. Bekhet, M. Hassaballah, M. A. Kenk, and M. A. Hameed, "An artificial intelligence based technique for COVID-19 diagnosis from chest X-ray," in *Proc. 2nd Novel Intell. Leading Emerg. Sci. Conf. (NILES)*, Giza, Egypt, Oct. 24-26, 2020, pp. 87-92, doi: 10.1109/NILES50944.2020.9257930.
- [7] N. Hilmizen, A. Bustamam, and D. Sarwinda, "The multimodal deep learning for diagnosing COVID-19 pneumonia from chest CT-scan and X-ray images," in *Proc. 3rd Int. Seminar Res. Inf. Technol. Intell. Syst. (ISRITI)*, Yogyakarta, Indonesia, Dec. 10-11, 2020, pp. 168-173, doi: 10.1109/ISRITI51436.2020.9315478.
- [8] Z. Karhan and F. Akal, "COVID-19 classification using deep learning in chest X-ray images," in *Proc. Med. Technol. Congr. (TIPTEKNO)*, Antalya, Turkey, Nov. 19-20, 2020, pp. 1-4, doi: 10.1109/TIPTEKNO50054.2020.9299315.
- [9] I. Mporas and P. Naronglerdrit, "COVID-19 identification from chest X-rays," in *Proc. Int. Conf. Biomed. Innov. Appl. (BIA)*, Varna, Bulgaria, Sep. 24-27, 2020, pp. 111-114, doi: 10.1109/BIA50171.2020.9244509.
- [10] A. M. Fangoh and S. Selim, "Using CNN-XGBoost deep networks for COVID-19 detection in chest X-ray images," in *Proc. 15th Int. Conf. Comput. Eng. Syst. (ICCES)*, Cairo, Egypt, Dec. 15-16, 2020, pp. 1-6, doi: 10.1109/ICCES51560.2020.9334600.

- [11] Z. Liang, J. X. Huang, J. Li, and S. Chan, "COVID-19 detection in medical images using deep learning techniques," in *Proc. IEEE Int. Conf. Bioinformatics Biomed. (BIBM)*, Seoul, South Korea, Dec. 16-19, 2020, pp. 1234-1239, doi: 10.1109/BIBM49941.2020.9313466.
- [12] V. N. M. Aradhya, M. Mahmud, M. Chowdhury, D. S. Guru, M. S. Kaiser, and S. Azad, "Learning through one shot: A phase by phase approach for COVID-19 chest X-ray classification," in *Proc. IEEE-EMBS Conf. Biomed. Eng. Sci. (IECBES)*, Langkawi Island, Malaysia, Mar. 1-3, 2021, pp. 45-50, doi: 10.1109/IECBES48179.2021.9398761.
- [13] A. Jain, D. Somwanshi, M. Manwal, D. R. Sahoo, N. Pathak, and K. Joshi, "COVID-19 demystified based on chest X-ray," in *Proc. 4th Int. Conf. Adv. Comput. Innov. Technol. Eng. (ICACITE)*, Greater Noida, India, May 14-15, 2024, pp. 1-6, doi: 10.1109/ICACITE60783.2024.10616894.
- [14] E. Khan, M. Z. U. Rehman, F. Ahmed, F. A. Alfouzan, N. M. Alzahrani, and J. Ahmad, "Chest X-ray classification for the detection of COVID-19 using deep learning techniques," *Sensors*, vol. 22, no. 3, p. 1211, Feb. 2022, doi: 10.3390/s22031211.
- [15] E. M. F. El Houby, "COVID-19 detection from chest X-ray images using transfer learning," *Sci. Rep.*, vol. 14, no. 11639, 2024, doi: 10.1038/s41598-024-61693-0.
- [16] S. Wang, J. Ren, and X. Guo, "A high-accuracy lightweight network model for X-ray image diagnosis: A case study of COVID detection," *PLoS ONE*, vol. 19, no. 6, Jun. 2024, doi: 10.1371/journal.pone.0303049.
- [17] B. Fatima, R. K. V. Raghavender, and R. Tallapelli, "Comparative analysis of EfficientNet-B0 and hybrid model (ResNet50 & DenseNet121) for lung cancer detection," *Int. J. Sci. Res. Eng. Manage.*, vol. 9, no. 7, pp. 1-9, Jul. 2025, doi: 10.55041/IJSREM51488.
- [18] I. Ahmad, A. Merla, F. Ali, B. Shah, A. A. AlZubi, and M. A. AlZubi, "A deep transfer learning approach for COVID-19 detection and exploring a sense of belonging with diabetes," *Front. Public Health*, vol. 11, p. 1308404, Nov. 2023, doi: 10.3389/fpubh.2023.1308404.
- [19] J. Chen, L. Wu, J. Zhang, L. Zhang, D. Gong, Y. Zhao, Q. Chen, S. Huang, M. Yang, X. Yang, S. Hu, Y. Wang, X. Hu, B. Zheng, K. Zhang, H. Wu, Z. Dong, Y. Xu, Y. Zhu, X. Chen, M. Zhang, L. Yu, F. Cheng, and H. Yu, "Deep learning-based model for detecting 2019 novel coronavirus pneumonia on high-resolution computed tomography," *Scientific Reports*, vol. 10, art. no. 19196, Nov. 2020, doi: 10.1038/s41598-020-76282-0.
- [20] N. B. Masters, S.-F. Shih, A. Bukoff, K. B. Akel, L. C. Kobayashi, A. L. Miller, H. Harapan, Y. Lu, and A. L. Wagner, "Social distancing in response to the novel coronavirus (COVID-19) in the United States," *PLOS ONE*, vol. 15, no. 9, e0239025, Sept. 11, 2020, doi: 10.1371/journal.pone.0239025.
- [21] D. Wang *et al.*, "Clinical characteristics of 138 hospitalized patients with 2019 novel coronavirus-infected pneumonia in Wuhan, China," *JAMA*, vol. 323, no. 11, pp. 1061-1069, Feb. 7, 2020, doi: 10.1001/jama.2020.1585.

- [22] F. Shan, Y. Gao, J. Wang, W. Shi, N. Shi, M. Han, Z. Xue, D. Shen, and Y. Shi, “Lung infection quantification of COVID-19 in CT images with deep learning,” *arXiv preprint arXiv:2003.04655v3*, Mar. 30, 2020, doi: 10.48550/arXiv.2003.04655.
- [23] W. Boukellouz and A. Moussaoui, “Magnetic resonance-driven pseudo CT image using patch-based multi-modal feature extraction and ensemble learning with stacked generalisation,” *J. King Saud Univ. – Computer Inf. Sci.*, 2019, doi: 10.1016/j.jksuci.2019.06.002.
- [24] M. Imad, N. Khan, F. Ullah, M. A. Hassan, A. Hussain, and one more author, “COVID-19 classification based on chest X-ray images using machine learning techniques,” *J. Computer Sci. Technol. Stud.*, Oct. 6, 2020.
- [25] M. A. Shereen, S. Khan, A. Kazmi, N. Bashir, and R. Siddique, “COVID-19 infection: emergence, transmission, and characteristics of human coronaviruses,” *J. Adv. Res.*, vol. 24, pp. 91–98, 2020, doi: 10.1016/j.jare.2020.03.005.
- [26] M. Z. Islam, M. M. Islam, and A. Asraf, “A combined deep CNN-LSTM network for the detection of novel coronavirus (COVID-19) using X-ray images,” *Informatics in Medicine Unlocked*, vol. 20, art. no. 100412, 2020, doi: 10.1016/j.imu.2020.100412.
- [27] Q. Li *et al.*, “Early transmission dynamics in Wuhan, China, of novel coronavirus–infected pneumonia,” *New England Journal of Medicine*, vol. 382, no. 13, pp. 1199–1207, Jan. 29, 2020, doi: 10.1056/NEJMoa2001316.
- [28] C. Y.-P. Lee, R. T. P. Lin, L. Renia, and L. F. P. Ng, “Serological approaches for COVID-19: epidemiologic perspective on surveillance and control,” *Frontiers in Immunology*, vol. 11, art. no. 879, Apr. 24, 2020, doi: 10.3389/fimmu.2020.00879.
- [29] L. Wang, Y. Zhang, D. Wang, X. Tong, T. Liu, S. Zhang, D. Wang, X. Tong, T. Liu, S. Zhang, *et al.*, “Artificial intelligence for COVID-19: a systematic review,” *Frontiers in Medicine*, vol. 8, art. no. 704256, Sept. 30, 2021, doi: 10.3389/fmed.2021.704256.
- [30] P. Lei *et al.*, “Clinical and computed tomographic (CT) images characteristics in the patients with COVID-19 infection: what should radiologists need to know?” *J. X-ray Sci. Technol.*, vol. 28, no. 3, pp. 369–381, 2020, doi: 10.3233/XST-200670.

Plagiarism Report

ORIGINALITY REPORT

21 %	19 %	17 %	16 %
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES


1	Submitted to Daffodil International University Student Paper	2 %
2	pdfs.semanticscholar.org Internet Source	1 %
3	pure.uva.nl Internet Source	1 %
4	Submitted to California State University, Fresno Student Paper	1 %
5	www.mdpi.com Internet Source	1 %
6	repositor.umm.ac.id Internet Source	1 %
7	arxiv.org Internet Source	1 %
8	ijece.iaescore.com Internet Source	1 %
9	link.springer.com Internet Source	1 %
10	pmc.ncbi.nlm.nih.gov Internet Source	<1 %
11	sciforum.net Internet Source	<1 %
12	www.researchgate.net Internet Source	<1 %

Md Tahmid Hasan & Abdul Ahad

Plagiarism Report

 Quick Submit

 Quick Submit

 Daffodil International University

Document Details

Submission ID

trn:oid::1:3331593329

Submission Date

Sep 7, 2025, 9:24 PM GMT+6

Download Date

Sep 7, 2025, 9:29 PM GMT+6

File Name

Pre_defence_report_Tahmid_Ahad_-_Final.pdf

File Size

795.7 KB

27 Pages

8,433 Words

43,577 Characters

*% detected as AI

AI detection includes the possibility of false positives. Although some text in this submission is likely AI generated, scores below the 20% threshold are not surfaced because they have a higher likelihood of false positives.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?

The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

**What does 'qualifying text' mean?**

Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.