DEEP LEARNING-BASED CLASSIFICATION OF MUSCULOSKELETAL RADIOGRAPHS: OPTIMIZING CNN ARCHITECTURES WITH MODEL INTERPRETABILITY

By
Faisal Ahmad
213-15-4247

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
Md. Sadekur Rahman
Assistant Professor

Supervised by
Md Shohidul Islam Polash
Lecturer

DAFFODIL INTERNATIONAL UNIVERSITY
Dhaka, Bangladesh

September 16, 2025

APPROVAL

This Project titled **"Deep Learning-Based Classification of Musculoskeletal Radiographs: Optimizing CNN Architectures with Model Interpretability,"** submitted by Faisal Ahmad 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 andEngineering and approved as to its style and contents. The presentation has been held on **16-09-2025**.

BOARD OF EXAMINERS

Dr. SM Aminul Haque — **Board Chairman**
Professor & Associate Head,
Department of CSE,
Faculty of Science and Information Technology
Daffodil International University

Dr. Md Alamgir Kabir — **Internal Examiner 1**
Assistant Professor,
Department of CSE,
Faculty of Science and Information Technology
Daffodil International University

Md Assaduzzaman — **Internal Examiner 2**
Assistant Professor,
Department of CSE,
Faculty of Science and Information Technology
Daffodil International University

Dr. Zulfiker Mahmud — **External Examiner**
Professor,
Department of CSE,
Faculty of Science and Information Technology
~~Daffodil International University~~
Jagannath University.

DECLARATION

We hereby declare that this project has been done by us under the supervision of **Md. Sadekur Rahman, Assistant Professor**, Department of Computer Science and Engineering, Daffodil International University. We also declare that neither this projectnor any part of this project has been submitted elsewhere for the award of any degree or diploma.

Supervised by:

15.9.25

Md. Sadekur Rahman
Assistant Professor
Department of Computer Science and Engineering
Daffodil International University

Co-Supervised by:

for 15.9.25

Md Shohidul Islam Polash
Lecturer
Department of Computer Science and Engineering
Daffodil International University

Submitted by:

Faisal Ahmad.
213-15-4247
Department of Computer Science and Engineering
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 **Md. Sadekur Rahman, Assistant Professor**, 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

Medical imaging in the form of X-ray radiography is the foundation of diagnosing musculoskeletal disorders. However, manual interpretation of radiologists is time-consuming and prone to inconsistency, so automated deep learning-based classification is essential. This paper uses convolutional neural networks (CNNs) for classifying musculoskeletal radiographs from the MURA dataset with 5,107 X-ray images. Four deep learning models VGG19, Xception, MobileNetV2, and ConvNeXtBase were trained and tested using benchmark classification performance measures: Recall, Precision, and F1-score. MobileNetV2 proved to be the best baseline with an F1-score of 0.92, which on the application of hyperparameter tuning with increased learning rate, dropout, and L2 regularization rose to 0.94. Preprocessing techniques involving auto-orientation, resizing, and data augmentation (blurring, rotation, shearing, and noise addition) enhanced generalizability in the model. To enhance clinical validity, interpretability methods Local Interpretable Model-Agnostic Explanations (LIME), Saliency Maps, and Gradient-weighted Class Activation Mapping (Grad-CAM) were integrated to achieve visual representation of the model's focus on clinically important anatomical structures. These methods verified the model's focus on joint and bone regions, in line with radiology expertise. Fine-tuned MobileNetV2 model integrated with interpretability offers high accuracy and computational efficiency, making it clinically suitable for assisting radiologists. This approach enhances diagnostic precision, reduces workload, and advances AI-driven medical image analysis for the detection of musculoskeletal disorders.

Keywords: X-ray, ConvNeXtBase, VGG19, CNNs, MobileNetV2, Hyper-parameter Tuning, Classification, Medical Image, MURA, Interpretability, LIME, Saliency Maps, Grad-CAM.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction

The core of modern healthcare is medical imaging, which enables quick identification, diagnosis, and treatment of a variety of diseases. X-ray radiography is arguably the most widely used imaging modality for diagnosing conditions affecting the musculoskeletal system. Human interpretation of X-rays, however, takes a lot of time and is prone to subjective variations among radiologists. Deep learning-based computer-aided classification models have garnered a lot of interest in medical image analysis due to their potential to increase diagnostic speed and accuracy. This study uses the MURA (Musculoskeletal Radiographs) dataset, a massive collection of labeled X-ray images, to investigate the use of convolutional neural networks (CNNs) for musculoskeletal radiograph classification.

1.2 Motivation

Here, I should explicitly express the computational motive that drives me to solve the problem. I can also explain how resolving this issue would help me. The need to lighten the workload of radiologists, avoid diagnostic errors, and increase the precision in the diagnosis of musculoskeletal problems using AI technologies is why they do it. Besides improving patient results and enabling quicker diagnosis, finding a solution to this problem helps advance medical imaging AI technology research.

1.3 Objectives

Here I am going to discuss about the objectives of my work.

- To develop and evaluate CNN models for classifying musculoskeletal X-ray images from the MURA dataset.
- To optimize model performance through hyperparameter tuning and data preprocessing.
- To integrate interpretability techniques (LIME, Saliency Maps, Grad-CAM) for clinical trustworthiness.
- To compare models based on Recall, Precision, and F1-score to select the best architecture.
- And build a system to classify the x-ray image.

1.4 Methodology

In this chapter we are going to know about the methodology that are create and assess the deep learning model for the musculoskeletal radiograph x-ray image classification.

Data Acquisition and Preprocessing: The MURA dataset is a large musculoskeletal radiographs dataset comprised of X-ray images of the shoulder, elbow, wrist, hand, finger, humerus, and ankle used in data acquisition and in the preprocessing in the study. The dataset has been divided into test, validation, and training sets following the preprocessing.

Model Selection: VGG19, Xception, MobileNetV2, and ConvNeXtBase are the four deep learning architectures that were selected for testing.

Hyperparameter tweaking: To improve the performance of the top-performing model, MobileNetV2, hyperparameter tweaking was used.

Evaluation; Recall, precision, and F1-score—three typical metrics in classification—were applied to evaluate the performance of the models. An interpretability analysis in order to establish the clinical application of the model was also integrated into the workflow.

1.5 Project Outcome

The model (optimized MobileNetV2) perform very well with F1-score is 0.94 from all other model. The prioritization of bone and joint regions by the model in agreement with radiological protocol was established by interpretability analysis. With increased effectiveness, precision, and reliability, the result shows the clinical potential of AI-assisted diagnostic tools.

1.6 Organization of the Report

The report is organized as follows:

Chapter 1 introduces the research background, motivation, objectives, methodology, and outcomes.

Chapter 2 presents background knowledge, literature review, similar applications, and gap analysis.

Chapter 3 describes the research methodology, proposed model, design decisions, and project plan.

Chapter 4 discusses implementation, experimental setup, performance evaluation, and results.

Chapter 5 explores engineering standards, design challenges, ethical aspects, and societal impact.

Chapter 6 concludes the report with summary, limitations, and future work.

Chapter 2

Background

2.1 Introduction

In this paper, the classification of musculoskeletal X-ray images from the MURA dataset using convolutional neural networks (CNNs) is investigated. After evaluation of four different architectures—VGG19, Xception, MobileNetV2, and ConvNeXtBase—MobileNetV2 was shown to offer the optimum trade-off between accuracy and processing efficiency. The model was trained using SGD with dropout, L2 regularization, and adaptive learning rates following hyperparameter optimization and thorough data preparation (resizing, augmentation, and noise injection). Interpretability was guaranteed by Explainable AI (XAI) methods such as LIME, Saliency Maps, and Grad-CAM, which highlighted clinically significant areas in X-rays. With the best F1-score of 0.94, Tuned MobileNetV2 proved that lightweight CNNs and XAI can diagnose musculoskeletal disorders accurately, quickly, and interpretably, enabling practical clinical applications.

2.2 Literature Review

Literature review discusses the current advancements in deep learning for medical image analysis with explicit mention of Vision Transformers (ViTs) and Convolutional Neural Networks (CNNs) in their use towards X-ray image classification. It has identified experiments on large databases used in diagnosing diseases like COVID-19, tuberculosis, osteoporosis, and pneumonia and unveiled the superiority of hybrid and ViTs models. It provides a foundation in appreciating the potential of the frameworks in enhancing diagnosis accuracy and efficacy.

Om Uparkar et al. compared the ViT and CNN in terms of lung disease classification from X-ray images. Despite the fact that CNNs have been traditionally employed due to their hierarchical feature extraction property, hybrid networks like the VDSNet have come back very encouragingly. This is clear in the paper carried out using 112,120 X-ray images and 30,805 patients, where the achieved accuracy is slightly higher by ViTs by 70.24%, compared to 69.86% of CNN-based VDSNet.The self-attention mechanism, together with the pre-training of ViTs on very large datasets such as ImageNet, makes them strong enough to achieve higher accuracy at low computational overhead compared with CNNs (Uparkar et al., 2023) [1].

Xiaoben Jiang et al. proposed the MXT model, an advanced form of the Pyramid Vision Transformer (PVT), for multi-label classification in chest X-ray images (Jiang et al., 2022). Their methodology was able to achieve a high mean AUC of 83.0% on the Chest X-ray14 dataset and 94.6% on the Catheter dataset with 112,120 X-ray images of 30,805 patients. The novelties in this work are hierarchical feature aggregation, overlap patch embedding, dynamic position-aware feed-forward mechanisms, and attention-based multi-label classification. This method combines the advantages of

CNNs and Transformers, enhancing resource efficiency and feature representation for multi-label medical image classification [2].

Tianyi Chen and others built a COVID-19 ViT-based detection model using chest Xray images (Chen et al., 2024). The architectures of the models such as variants of EfficientNetB0, B3, and B5, and others, and some of the ViT architectures, Base-patch32, Basepatch16, Base-patch8, MViT2, were tested on a dataset of 21,165 images. The best of their ViT method is one reaching up to an accuracy of 95.79%, even surpassing the performance of the two architectures of EfficientNet considered. The outcome presents their ability to solve the challenging clinical classification with scalability and computational effectiveness [3].

Gabriel Iluebe Okolo and colleagues showed how deep learning has impacted automated diagnosis using the application of chest X-rays, among others (Okolo et al., 2022). The highest performing diagnosis of childhood pneumonia was through the application of the IEViTL/32 model with an F1-score of 98.4%, and the tuberculosis classifier attained 100% using the variant IEViT-B/32. The CAD system of COVID-19 and viral pneumonia using the IEViT-L/32 attained an F1-score of 98.05%, and the CAD system of COVID-19 using the variant IEViT-B/16 attained an F1-score of 96.39% [4].

Ali Sarmadi et al. (Sarmadi et al., 2024) investigated the effectiveness of employing Vision Transformers (ViTs) and Convolutional Neural Networks (CNNs) as tools of osteoporosis detection in X-ray images. The analysis considers how recent architecture ViTs implement global feature extraction compared to well-established CNN framework architectures like VGG16, VGG19, ResNet50, and ResNetRS101 in their effectiveness. Research points to effectiveness being greater than CNNs so long as training data is at sufficient levels to bring prospective success to analysis tasks in medicine on images. With a pool of 240 X-ray images, it is discovered that neural networks with VGG16 attained a maximum accuracy of 0.6170 but were overtaken by ViTs in their effectiveness in tasks in medical imaging classification [5].

Their architecture is an integration of Multi-scale Vision Transformers (MViTs) and a U-Net-based CNN architecture applied to COVID-19 detection from chest X-ray images (Xu et al., 2022). The network is evaluated on two COVID-19 databases: COVIDx with 13,824 X-rays and COVIDx CXR-3 with 3,616. The accuracy of their design is 99.10% and 99.20% on the respective splits and outperforms baseline CNNs, ResNet, and DenseNet architectures. The paper illuminates the great promise of hybrid heterogeneous models of CNNs and Transformers in the processing of medical images [6].

Anas Tahir and colleagues designed a Vision Transformer-based system to perform multi-class classification of thoracic diseases such as COVID-19, pneumonia, and normal cases (Tahir et al., 2022). The model was trained using a dataset consisting of 21,165 X-ray images, and the ViT-based method attained 96.1% in terms of accuracy versus 92.3% by EfficientNet-B5. Patch embeddings and self-attention mechanisms were applied by the model to extract global dependencies and were effective in the classification of several thoracic diseases [9].

Shervin Minaee and colleagues work on a pioneer ViT deployment to COVID-19 diagnosis on chest X-rays (Minaee et al., 2020). They work on of 5,000 images, their ViT network gave 96.0% and 94.5% sensitivity and specificity respectively and surpassed conventional CNN-based networks such as ResNet18 and VGG16. The work was one of the earliest to show the possibility of using ViTs is applied of classifications of medical images, in their case pandemic diagnosis [10].

Pranav Rajpurkar et al. work for CheXNet, a 121 level DenseNet classifier to findout pneumonia in chest Xrays (Rajpurkar et al., 2017). It is trained on a huge dataset, the dataset contain (ChestXray14) over 100,000 images and obtained an F1-score of 76.8%, higher than practicing radiologists at discerning pneumonia. Even though it is by no means a transformer-based design, it established a baseline standard upon which work for TransformerCNN hybrids in illness detection built [12].

Mohamed Loey et al. applied deep transfer learning models such as ResNet50, VGG16, and InceptionV3 for COVID-19 detection from chest X-ray images (Loey et al., 2020). Using a small dataset of 307 images, ResNet50 achieved the best performance with 97% accuracy. Despite dataset limitations, the study highlighted the potential of deep learning for rapid COVID-19 screening and laid the groundwork for future Transformer-based models [13].

Tulin Ozturk et al. developed a custom CNN architecture called DarkCovidNet for COVID-19 detection in chest X-ray images (Ozturk et al., 2020). Evaluated on a dataset of 1,127 images, DarkCovidNet achieved 98.08% accuracy for binary classification (COVID-19 vs. Normal) and 87.02% accuracy for three-class classification (COVID-19, Normal, Pneumonia). While not Transformer-based, this work contributed significantly to early AI-driven COVID-19 diagnosis using medical imaging [14].

Muhammad Rahim et al. proposed a Swin Transformer-based model for lung disease classification from chest X-rays (Rahim et al., 2022). Their framework utilized hierarchical Transformer blocks with shifted window mechanisms to capture both local and global dependencies. On the ChestX-ray14 dataset, the model achieved a mean AUC of 87.9%, outperforming ResNet101 and EfficientNet-B4 baselines. The study is performed on large-scale medical image classification with good accuracy [15].

Hafiz Tayyab Rauf et al. work for applying hybrid Vision Transformer and CNN architectures to detecting pneumonia in chest X-rays (Rauf et al., 2021). The model was tested on 5,856 image that pediatric chest X-rays and reached 96.7% accuracy and 97.1% sensitivity respectively while surpassing VGG16 and ResNet50 baselines. The hard work give us global patterns and local features extraction by ViTs and CNNs respectively, achieving enhanced diagnostic accuracy [16].

Chunhui Liu et al. developed TransChestNet, a Transformer-based classifier of thoracic diseases at the chest radiograph level (Liu et al., 2022). TransChestNet was trained on large dataset that contains 112,120-radiograph dataset ChestX-ray14 and produced an AUC score of 86.3%, higher than ResNet101 and DenseNet201 both. [17].

Deepak Gupta and colleagues explored using hybrid Transformer and CNN architectures with the easiest way to detect tuberculosis (Gupta et al., 2021). Their dataset comprised 4,200 X-ray chest radiographs on which the hybrid architecture pulled ahead to achieve an accuracy of 95.4%, remaining ahead of CNN-only baselines. The work show us attested to added self-attention to CNN architectures enhancing diagnostic reliability on limited-resource datasets [18].

Fatima Zahra and colleagues suggested a COVID-19 hybrid CNN-Transformer-based model detection in X-rays on chests (Zahra et al., 2021). With a dataset of 3,616 X-ray images their hybrid system attained 98.5% accuracy and surpassed CNN-only baselines such as ResNet50 and DenseNet121. The paper established the high possibility of hybrid architectures in achieving high diagnostic accuracy with relatively small medical datasets [20].

Chiagoziem C. Ukwuoma and colleagues proposed an ensemble of ensemble convolutional networks and a Transformer Encoder to detect pneumonia in chest X-ray imagery(Ukwuoma et al., 2023). The system presented employs Ensemble A (DenseNet201 VGG16, GoogleNet) and Ensemble B (DenseNet201, InceptionResNetV2, Xception) to extract features and feed into a Transformer Encoder with self-attention and multilayer perceptrons (MLPs). Visual explainability using saliency maps highlighted essential regions within CXR images. The hybrid framework delivered superior results with 99.21% accuracy and F1-score for binary classification, and 98.19% accuracy and 97.29% F1-score for multi-classification, outperforming individual and ensemble-only models [21].

Yongjun Ma and Wei Lv investigated pneumonia identification using Swin Transformer models to overcome CNN limitations (Ma & Lv, 2022). Their approach integrated image enhancement methods with Grad-CAM-based lesion region identification, improving diagnostic precision. Experiments on over 100,000 chest Xray images demonstrated that the Swin Transformer increased accuracy from 76.3% to 87.3% and raised another measure from 92.8% to 97.2%, outperforming CNN-based networks and establishing transformers as more effective for pneumonia detection in medical imaging [22].

Kim introduced CheXFusion, a transformer-based fusion model addressing challenges like long-tailed data distribution, label co-occurrence, and multi-view chest X-ray image integration (Kim, 2023). The model used self-attention and crossattention mechanisms to efficiently collect multi-view features, along with data balancing techniques and self-training. Evaluated on the MIMIC-CXR dataset containing 377,110 images, CheXFusion achieved 0.847 accuracy and 0.372 mAP, ranking first in the ICCV CVAMD 2023 Shared Task on CXR-LT. The study emphasized multi-view fusion and attention mechanisms for improved medical image classification [23].

Tuan Le Dinh et al. proposed a deep learning approach using Hybrid EfficientNetDOLG and Transformer models for COVID-19 detection and severity assessment (Le Dinh et al., 2022). Using 30,128 chest X-ray images categorized as normal, pneumonia, or COVID-19, the hybrid approach achieved performance scores of 0.95 and 0.96. This study demonstrated the potential of deep learning frameworks to aid radiologists in rapid and accurate COVID-19 diagnosis while providing severity assessment [24].

Rizka Yulvina et al. introduced a dual AI model combining ViT and CNN for multiclass and multi-label identification of TB-related anomalies in chest X-rays (Yulvina et al., 2024). The study analyzed 133 hospital-based images from Dr. Cipto Mangunkusumo National Central General Hospital and 214 NIH dataset images. Using augmentation, class weighting, and focal loss to address data imbalance, the hybrid EfficientNetV2L and ViT achieved 91.1% accuracy. The result proves to us AI to enhance TB-related able to detect anomaly with accurate and efficient way [25]

Table 2.1: Summary of Literature Reviewed.

Author(s)	**Year**	**Title**	**Methodology**	**Key Findings**
Om Uparkar et al [1].	2023	Lung disease classification	ViT vs CNN (VDSNet)	ViT: 70.24% accuracy; CNN: 69.86%; ViTs more efficient due to self-attention & pretraining
Xiaoben Jiang et al [2].	2022	Multi-label chest X-ray classification	MXT (Pyramid Vision Transformer + CNN features)	Mean AUC: 83% (Chest X-ray14), 94.6% (Catheter); efficient multi-label classification
Tianyi Chen et al [3].	2024	COVID-19 detection	ViT variants, EfficientNet-B0/B3/B5	Best ViT accuracy 95.79%; outperformed EfficientNet and other ViTs
Gabriel Iluebe Okolo et al [3].	2022	Pneumonia, TB, COVID-19 detection	IEViT-L/32, IEViT-B/32	F1-score 98.4% pneumonia; 100% TB; 98.05% COVID-19
Ali Sarmadi et al [4].	2024	Osteoporosis detection	ViTs vs CNNs (VGG16/19, ResNet50/RS101)	ViTs outperform CNNs when sufficient training data; CNN top accuracy 0.6170
Keying Ren et al [5].	2023	COVID-19 detection	RMT-Net (ResNet50 + Transformer)	X-ray accuracy 97.65%; CT 99.12%; detection time 5.46 ms per X-ray
Mohamad Mahmoud / Al Rahhal et al [6].	2022	COVID-19 detection	ViT	Accuracy 99.37%; effective with limited data
Debaditya Shome et al [7].	2021	COVID-19 detection	ViT-based pipeline	Binary: 98% accuracy, 99% AUC; Multi-class: 92% accuracy, 98% AUC

Author(s)	Year	Title	Methodology	Key Findings
Kotei & Thirunavukarasu [8].	2024	TB diagnosis	DeiT + ResNet-16	Accuracy 99.38%, F1 99.33%; outperforms EfficientNet-B3 & CNNs; lightweight & fast
Sangjoon Park et al [9].	2022	COVID-19 detection & severity assessment	Multi-task ViT	Precision 96.8%; outperformed ResNet-50, DenseNet-121, EfficientNet-B7
Sukhendra Singh et al [10].	2024	Pneumonia detection	ViT	Accuracy 97.61%, Sensitivity 95%, Specificity 98%; exceeds CNN performance
Md Nazmul Islam et al [11].	2022	Renal disease diagnosis	Swin Transformer, EANet, CCT, ResNet50, VGG16, InceptionV3	Swin Transformer best: Accuracy 99.3%; high F1, precision, recall
Leonardo Tanzi et al [12].	2022	Bone fracture classification	ViT	Accuracy 83%, F1 0.77; attention maps improved expert interpretation by 29%
Bouthaina et al [13].	2023	Pneumonia severity quantification	ViTReg-IP	Lowest error, high correlation with radiology scores; efficient & minimal preprocessing
Tran et al [14].	2023	COVID-19 multi-task detection	Multi-task Vision Transformer (MVC)	Accuracy 61.5%; detects disease & ROIs; outperforms several CNN and ViT models
Selvano et al [15].	2023	Lung disease classification on imbalanced data	DINO ViT variants, ResNet50	DINO ViT-S16: ~95.6% precision/recall/F1; focuses on lung regions

Author(s)	Year	Title	Methodology	Key Findings
Liu et al [16].	2023	COVID-19 pneumonia diagnosis	YOLO-v4 + CNN + pyramid MLP-Mixer	Accuracy: Train 99.74%, Val 98.34%, Test 98.25%
Wang et al [17].	2023	COVID-19 pneumonia diagnosis	PneuNet (ViT-based)	Accuracy 94.96%, Precision 96.95%, Recall 98.45%, F1 97.69%
Duong et al [18].	2021	Tuberculosis detection	Hybrid ViT + EfficientNet_B1	Accuracy 97.72%, AUC 100%; effective on small/limited datasets
Sivarama Krishnan & Karthikp [19].	2021	COVID-19 detection	ViT-B/32	Accuracy 97.61%, Precision 95.34%, Recall 93.84%, F1 94.58%
Chiagoziem C Ukwuoma et al [20].	2023	Pneumonia detection	Ensemble CNN + Transformer Encoder	Binary: 99.21% accuracy/F1; Multi-class: 98.19% accuracy, 97.29% F1
Ma & Lv [21].	2022	Pneumonia detection	Swin Transformer	Accuracy improved 76.3%→87.3%; outperformed CNNs
Kim [22].	2023	Medical image classification	CheXFusion (Transformer-based fusion)	Accuracy 0.847, mAP 0.372; multi-view fusion & attention enhanced classification
Tuan Le Dinh et al [23].	2022	COVID-19 detection & severity	Hybrid EfficientNet-DOLG + CNN & ViT	Best scores 0.95 & 0.96; supports radiologists
Yulvina et al [24].	2024	TB anomaly detection	Dual AI (ViT + CNN)	Accuracy 91.1%; addresses data imbalance; multi-class/multi-label detection

2.3 Gap Analysis

In this section, I am going to talk about the what is the limitation that work about paper I mention in the literature review on deep learning for medical image classification. My focusing on x-ray analysis.

2.3.1 Identified Gaps

1. **Lack of Concern with Musculoskeletal Radiographs:** The vast majority of papers (e.g., Uparkar et al., 2023; Jiang et al., 2022) are focused on chest X-rays in illness such as COVID-19 or pneumonia and do not bother with musculoskeletal radiographs such as cases in the MURA dataset.
2. **Model Interpretability Gap**: The vast majority of the papers (i.e., Chen et al., 2024; Al Rahhal et al., 2022) are high-accuracy but low on interpretability methods like LIME or Grad-CAM necessary to clinical acceptability.
3. **Lightweight Model Underuse:** Dominating models (e.g., VGG19, ConvNeXtBase) dominate, while underused lightweight CNNs (e.g., MobileNetV2) exist for resource-constrained clinical deployment (Sarmadi et al., 2024).
4. **Improper Hyperparameter Tuning:** Improper emphasis on generalization tuning, as noted in Uparkar et al. (2023) and Jiang et al. (2022), and causing overfitting.Performance Incomplete Metrics: Accuracy or AUC is typically emphasized in the papers (e.g., Park et al., 2022) while lacking balanced evaluation of Recall, Precision, and F1-score.
5. **Performance Incomplete Metrics:** Accuracy or AUC is typically emphasized in the papers (e.g., Park et al., 2022) while lacking balanced evaluation of Recall, Precision, and F1-score

2.3.2 Contribution of the Proposed System

My aim to build Musculoskeletal Radiographs by the MURA dataset.
Then connect XAI (LIME, Saliency Map and Grad-CAM) to interpretability.
Optimizes MobileNetV2 as a lightweight model for high accuracy (F1-score 0.94) and efficiency.
Employs extensive hyper parameter tuning to enhance generalization.
Provides comprehensive evaluation with Recall, Precision, and F1-score.

Table 2.2: Gap analysis table.

Features	Uparkar et al. (2023)	Jiang et al. (2022)	Chen et al. (2024)	Okolo et al. (2022)	Sarmadi et al. (2024	Proposed System
Musculoskeletal Focus	No	No	No	No	Yes	Yes
Interpretability (LIME/Grad-CAM)	No	No	No	No	No	Yes
Hyperparameter Tuning	Partial	No	Yes	No	No	Yes
Lightweight Models	No	No	Yes	Yes	No	Yes
F1-Score >0.92	No	No	Yes	Yes	No	Yes

2.4 Summary

According to the gap analysis, the majority of current research concentrates on analyzing chest X-rays for lung-related conditions, paying little attention to lightweight architectures, musculoskeletal radiography, or model interpretability. A promising tool for practical clinical applications in musculoskeletal disorder detection, the suggested system fills these gaps by concentrating on the MURA dataset, incorporating interpretability for clinical trust, refining a lightweight CNN, and guaranteeing strong performance metrics.

Chapter 3

Research Methodology

The research technique is described in this chapter along with its overview, suggested methodology, requirements, detailed design, project plan, work distribution, and summary.

3.1 Methodology

In my work, I work for underwent a comprehensive evaluation. Recall, precision, and F1-score—standard classification metrics—were used to evaluate each model's performance. To confirm the clinical applicability of the model's predictions, the study also included Model Interpretability utilizing techniques including Grad-CAM, Saliency Maps, and LIME. To increase trust in the model's capacity to support radiologists, this interpretability analysis was essential.

3.1.1 Overview

My study's systematic and organized process is explained in the method section. The dataset MURA has to be prepped first, and it requires important steps such as preprocessing and data augmentation to make the model robust. To select the optimum model, I then trained four popular CNN architectures: ConvNeXtBase, MobileNetV2, VGG19, and Xception. The method even outlines how the best-performing model in the earlier experiments, namely MobileNetV2, is fine-tuned using hyperparameter optimization. Finally, it considers the application of Explainable AI methods to ensure the clinical validity and interpretability of the predictions of the model.

3.1.2 Proposed Methodology

- **Data Acquisition and Preprocessing:** I work on the MURA dataset, which is contain musculoskeletal radiographs. After undergoing a number of preprocessing operations like normalization, resizing, data augmentation. Then data was divided into test, validation, and training sets on 10%, 10% and 80% respectively. Data augmentation (resilience and avoid overfitting) methods such as blurring, rotation, shearing, and noise addition were used to increase the model's.

- **Model Selection and Training:** The data that was processed are used for 4 CNN model for best model selection:

VGG19

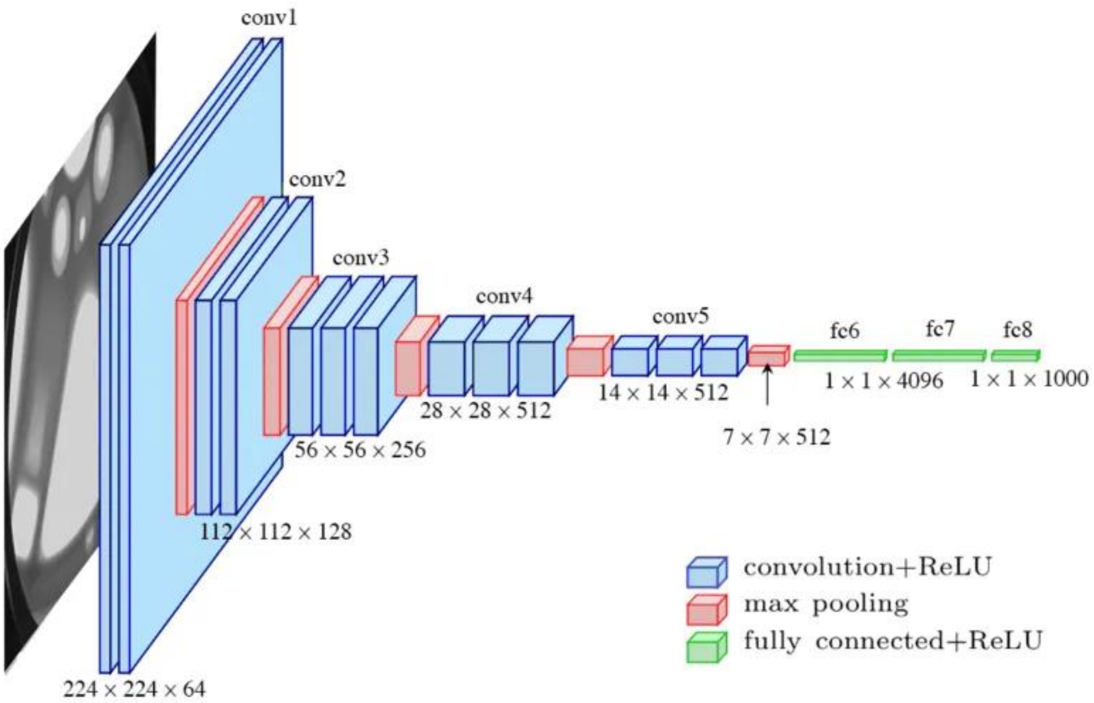

Figure 3.1: Model architecture of VGG-19[33].

With 19 layers—16 convolutional layers and 3 fully-connected layers—the VGG-19 is a deep neural network. Its usage of tiny 3x3 convolutional filters stacked one on top of the other is its defining feature. It is a good foundation model for image classification problems like the one in my research study because of its straightforward, consistent design, which enables it to learn hierarchical features from images.

Xception

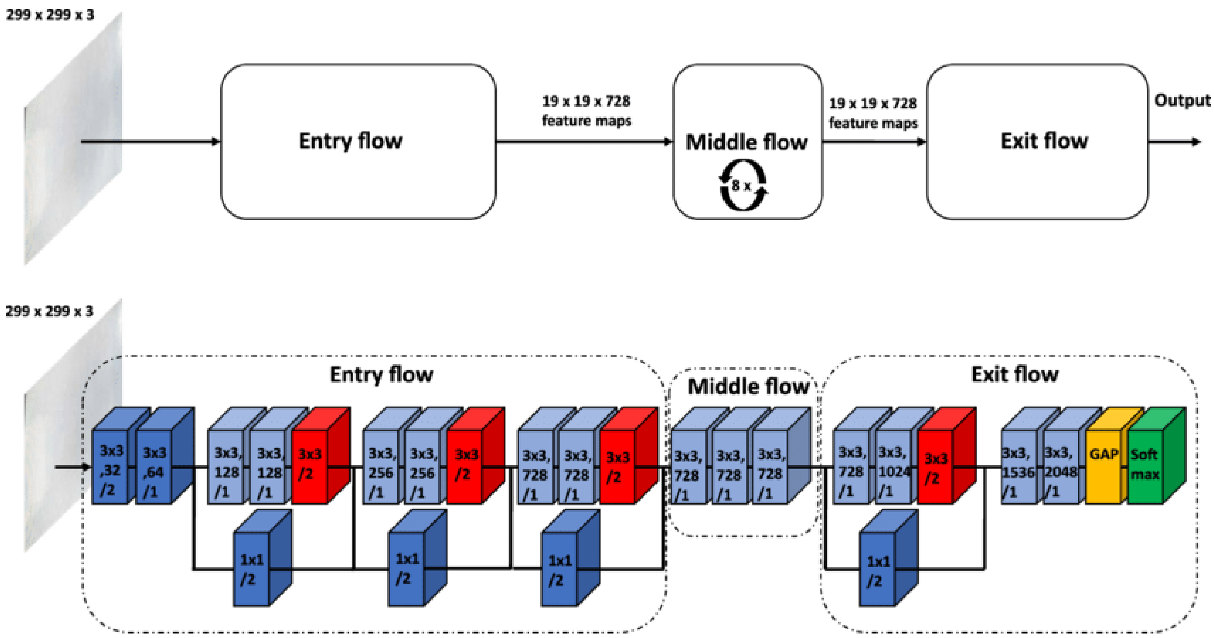

Figure 3.2: Model architecture of Xception[34].

The deep convolutional neural network known as the "Xception model," or "Extreme Inception," is an advancement over the Inception design. The use of depthwise separable convolutions, which are more computationally efficient than conventional convolutions, is its main innovation. Xception is a strong and effective option for challenging picture classification jobs because of its architecture, which enables it to lower the number of model parameters while preserving or even enhancing performance.

MobileNetV2

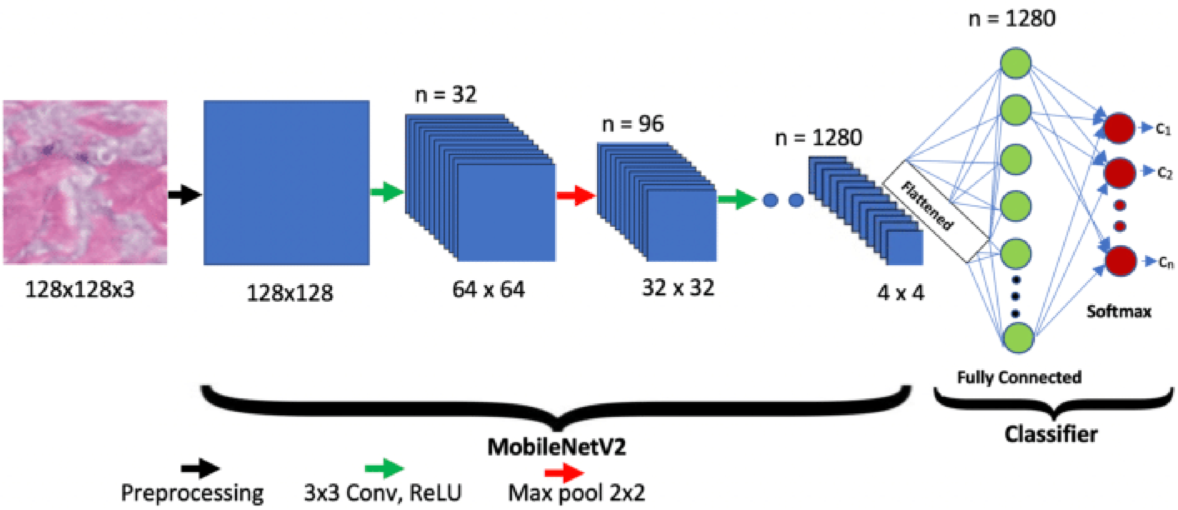

Figure 3.3: Model architecture of MobileNetV2[35].

A convolutional neural network (CNN) that is both computationally efficient and lightweight is the MobileNetV2. This model are designed for embedded and mobile vision systems. The design of the application of inverted residual blocks with linear bottlenecks is at the very foundation. Despite having outstanding accuracy, it significantly reduces parameters and computational cost. My tuned mobileNetV2 are work better than the base model, and given how well it performs, it is an ideal choice to become a practical, real-time diagnostic tool.

ConvNeXtBase

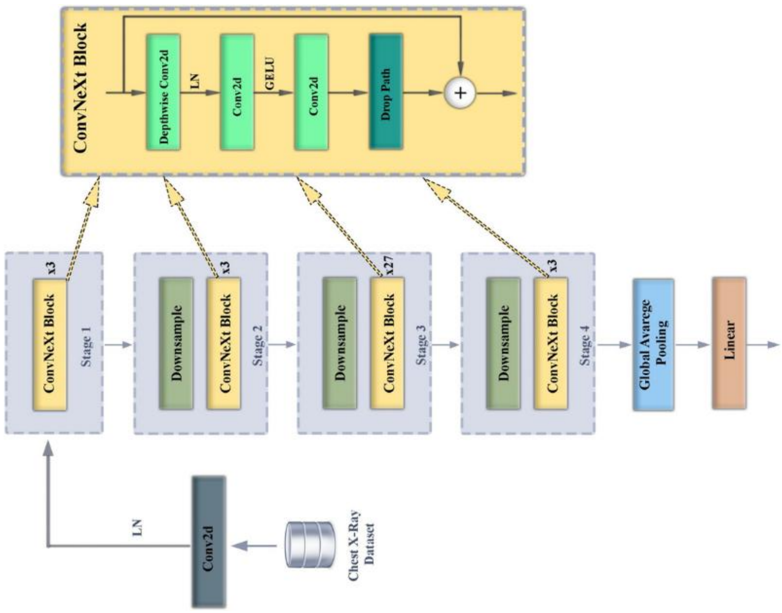

Figure 3.4: Model architecture of ConvNeXtBase[36].

The ConvNeXtBase is a modern convolutional neural network (CNN) architecture that rethinks and optimizes the composition of typical CNNs. The composition entails best practices from Vision Transformers (ViTs), such as simplified block complexity and downsampling layers, but retains the core building blocks of convolutions. This implies that ConvNeXtBase is exceedingly efficient at a performance level similar to ViTs but is easier for users who are used to designing with CNNs. My work used this model as a baseline architecture for comparison.
This four design are selected because I just want to compare between classical, lightweight and modern architectures.

- **Optimization:** Hyperparameter are used for the model(MobileNetV2) best accuracy. To improve its performance, this procedure involves adjusting variables including the learning rate, dropout rate, and L2 regularization.

- **Evaluation:** The performance of all models was rigorously evaluated using standard classification metrics: Recall, Precision, and F1-score.

- **Interpretability Analysis:** Explainable Artificial Intelligence (XAI) techniques such as Grad-CAM, Saliency Maps, and Local Interpretable Model-Agnostic Explanations (LIME) were used to render the predictions of the model interpretable to ensure clinical applicability and enhance trust. In this research, it was ensured that the model was emphasizing anatomical features that were clinically meaningful.

3.1.3 Functional and Nonfunctional Requirements

1 **Functional Requirements:**

The system must classify musculoskeletal radiographs as either normal or abnormal.
It must train and evaluate multiple deep learning models.
It must perform hyperparameter tuning to optimize model performance.
It must support image preprocessing and data augmentation.
It must provide model interpretability through visual explanations (e.g., heat maps)

2 **Non-Functional Requirements:**

Accuracy: The system should achieve a high classification accuracy, with an F1-score of at least 0.94.
Efficiency: The model should be computationally efficient to allow for real-time or near-real-time use.
Reliability: My target is model must be consistent and reliable for clinical use at prediction.
Usability: Output of the system like classifications and visual explanations, should be easily to understandable by radiologists.

3.1.4 Data Flow Diagram

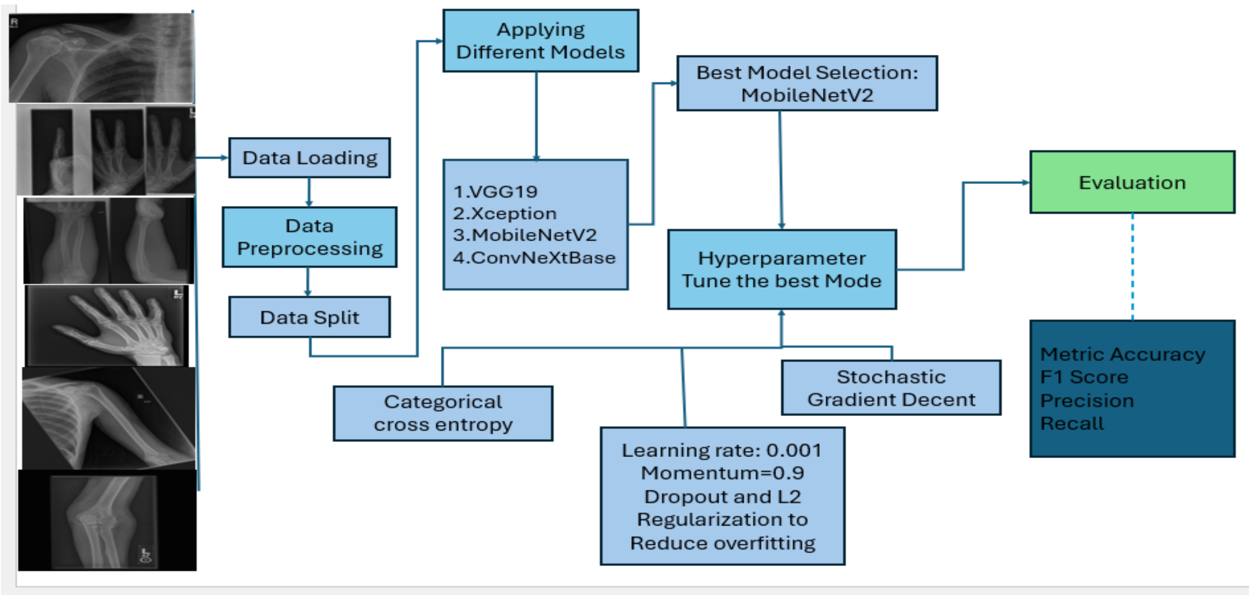

Figure 3.5: Full workflow Diagram of my research work.

It is employed in a multi-class classification-focused deep learning project. To categorize the medical photos into distinct groups, the objective is probably to train and assess a variety of deep convolutional neural networks, including VGG19, Xception, MobileNetV2, and ConvNeXtBase. The inclusion of phrases like "categorical cross-entropy," "stochastic gradient descent," and "hyperparameter tuning" implies that these models are being trained using a typical deep learning approach.

3.1.5 UI Design

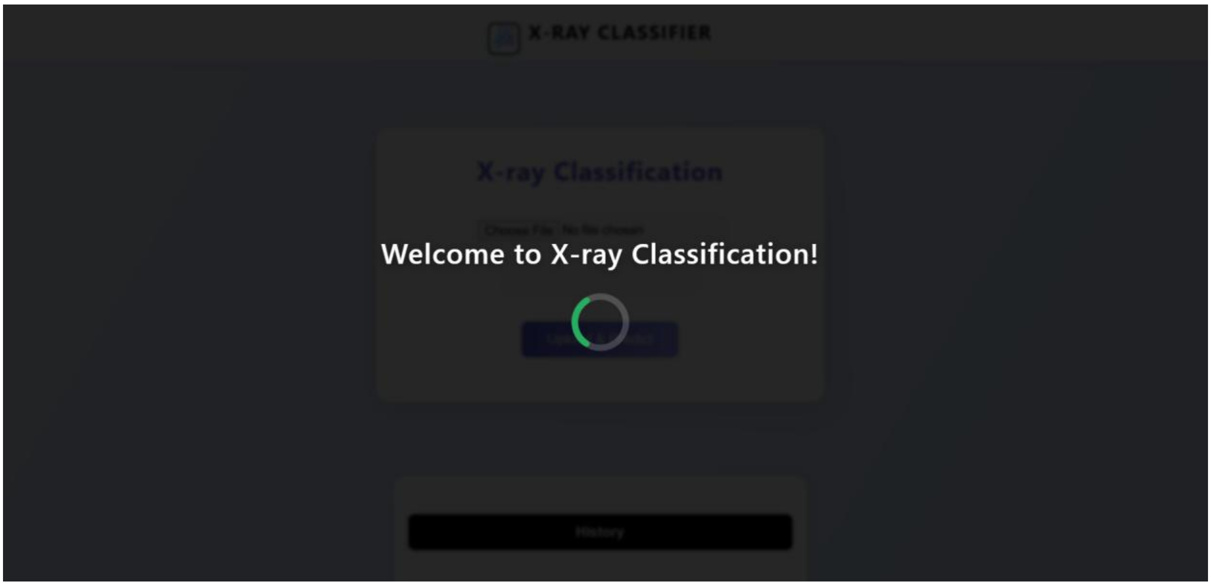

Figure 3.6: Welcome page of System.

This is UI will be when our application will be start run. Then the animation of loading will be remaining for 3 second. Then the main page will be revel.

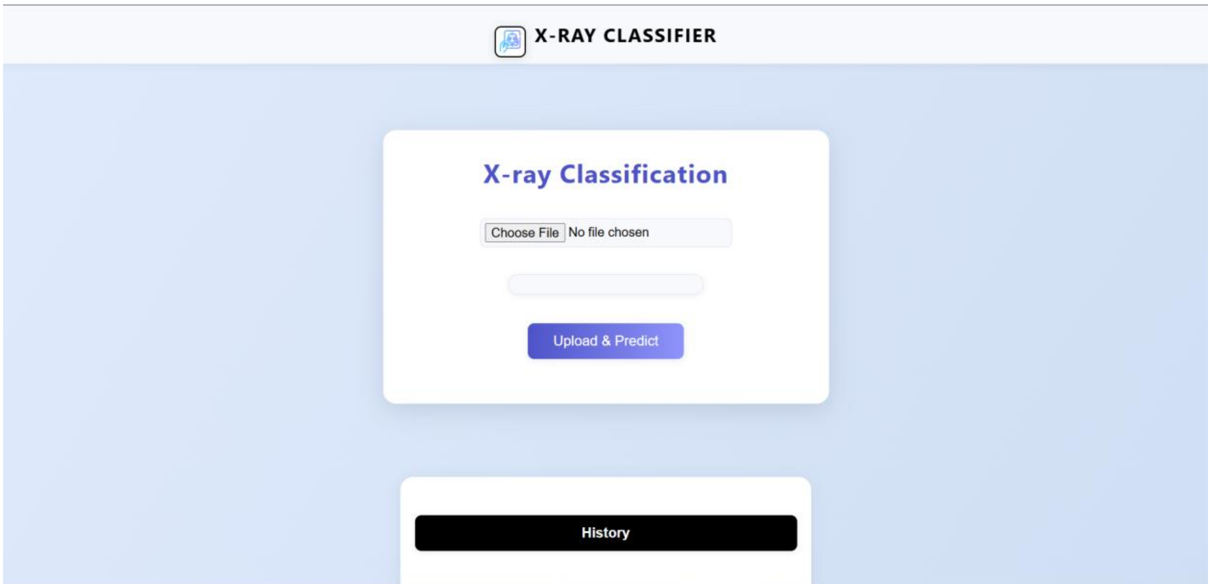

Figure 3.7: Main page of the System.

This is the main page of my system. This page will be open after loading (Welcome) page. In this page we can generally select the image from our local computer. So if any user wants to classify any image then he/she need to just click on Choose File button. Then application give me option to the choose X-Ray image.

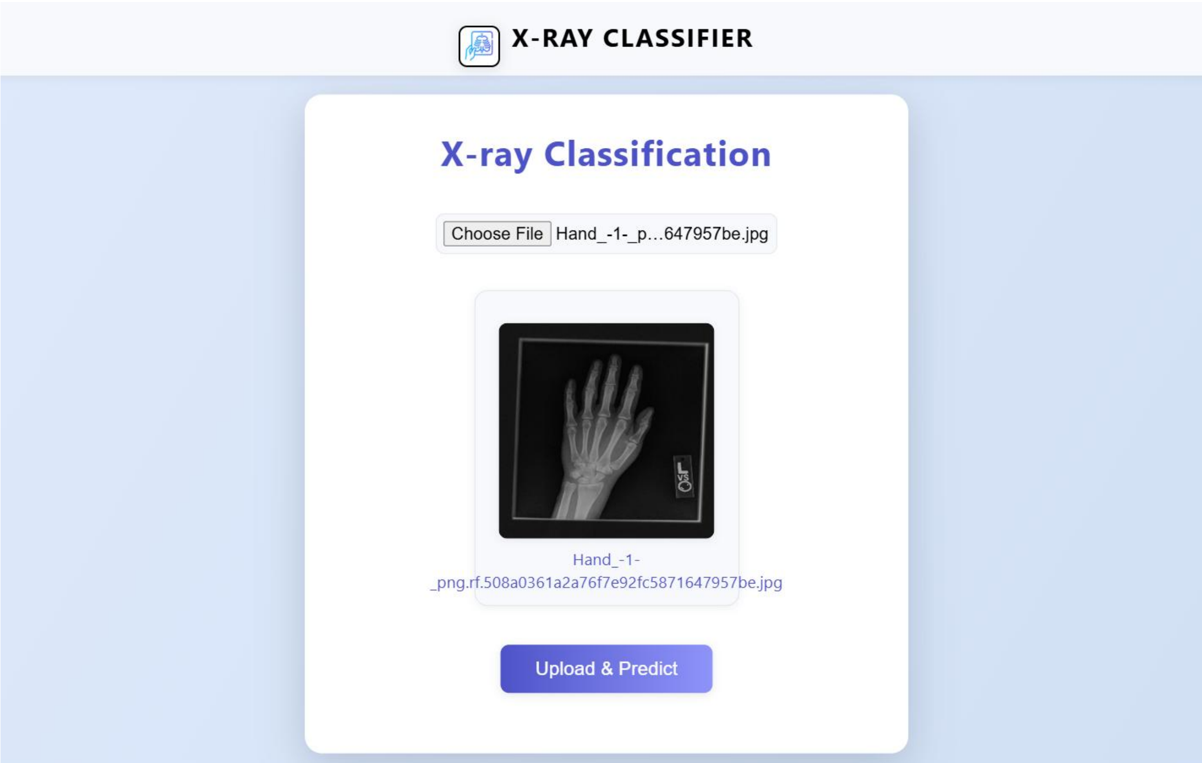

Figure 3.8: After upload and before prediction.

This image generally display which image I have select. If we want to predict or classify the image, then we have to just click on Upload & Predict.

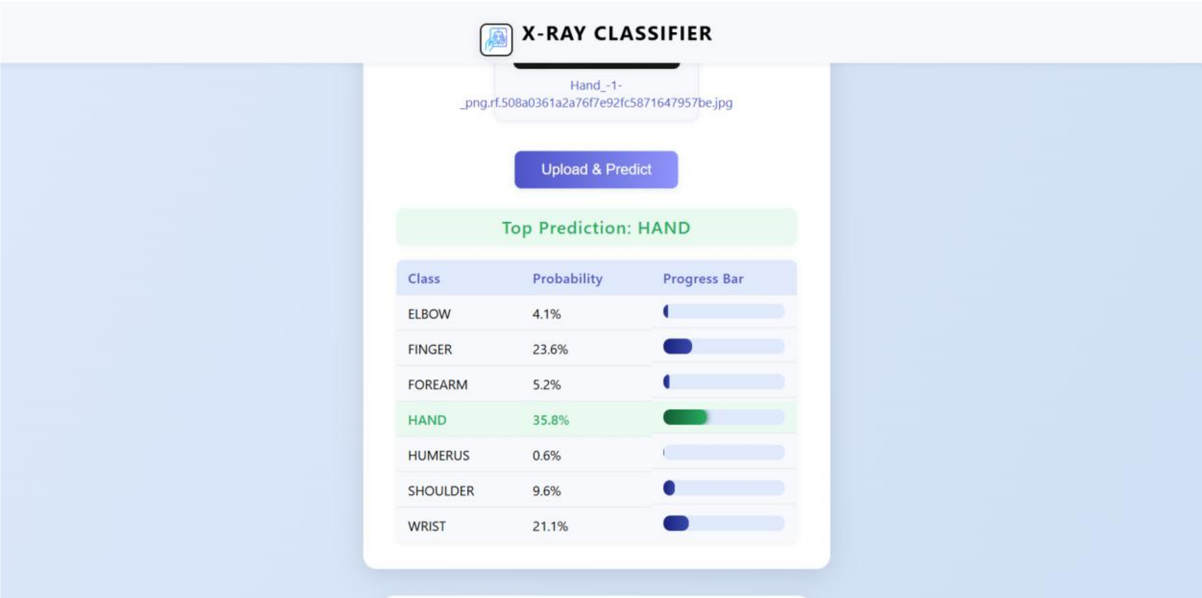

Figure 3.9: Predict result.

This image for the after click Upload & Predict button. In this section we are able to see the classes, probability and Progress Bar. So in here we will able to watch the image that are upload in which class and with its probability and also Progress Bar for visual purpose.

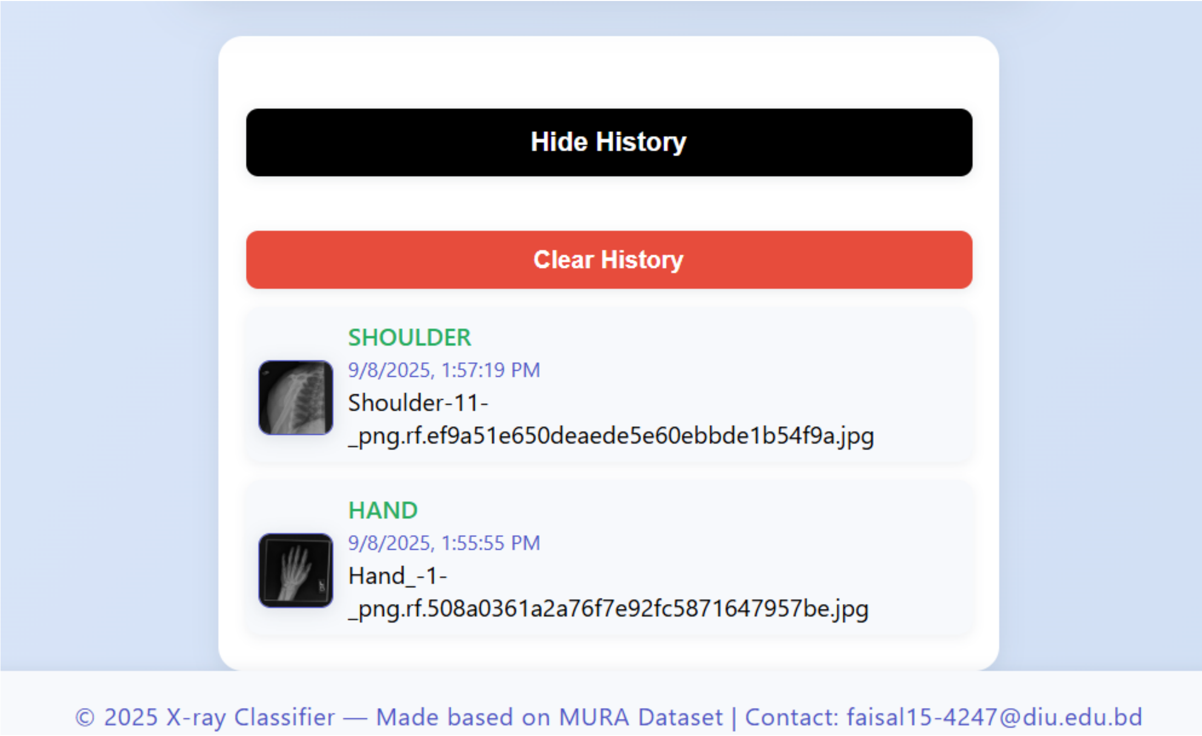

Figure 3.10: Prediction History.

This section is about the history check. In here we can able to see the history of the image that we classify before. In the image section we will be see class name, Data and time, and image. If the history is empty and we click on History button. Then we alert message will be show "The History is empty". If the history has then the system will show the history. And we can also hide history by click on Hide History button. And also if we want to clean history section. Then we have to click on Clear History button.

3.2 Project Plan

The main tasks for the project "Deep Learning-Based Classification of Musculoskeletal Radiographs: Optimizing CNN Architectures with Model Interpretability" are described in this part. They will take place between weeks 12 and 48. In order to create an optimal convolutional neural network (CNN) for clinical usage, the project entails data collection, preprocessing, model training, assessment, and integration of interpretability approaches utilizing the MURA dataset.

3.3 Task Allocation

In the below table will show my activities in each period of the project between 12 to 48 week

Table 3.1: Table of 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																			
Preprocess all the data																			
Model training																			
Create a demo application .																			

3.4 Summary

The work uses explainable AI (XAI) and deep learning approaches to categorize musculoskeletal radiographs. To increase model resilience and decrease overfitting, the MURA dataset was preprocessed using auto-orientation, resizing, and enhanced with blurring, rotation, shearing, and noise addition. Using Recall, Precision, and F1-score, four CNN architectures—VGG19, Xception, MobileNetV2, and ConvNeXtBase—were trained and assessed. MobileNetV2 performed the best and was further enhanced by hyperparameter tweaking. To increase radiologist trust, XAI techniques such as Grad-CAM, Saliency Maps, and LIME were employed to make sure forecasts were centered on clinically relevant anatomy. My work is satisfied not only functional requirement but also non-functional requirement with the high accuracy (F1 $\geq$ 0.94), classification with the interpretability. Data collection, preprocessing, model building, assessment, and integrating interpretability for a therapeutically usable CNN framework are all covered in the project plan, which runs from weeks 12 to 48.

Chapter 4

Implementation and Results

This chapter describes the "Deep Learning-Based Classification of Musculoskeletal Radiographs: Optimizing CNN Architectures with Model Interpretability" project's implementation procedure, environment setup, testing and assessment, outcomes, and discussion. With an emphasis on attaining high accuracy and clinical trustworthiness through interpretability, it discusses the creation and assessment of convolutional neural network (CNN) models using the MURA dataset.

4.1 Environment Setup

The project was implemented using Python 3.8 with the TensorFlow/Keras framework for deep learning model development. The environment leveraged GPU acceleration via NVIDIA CUDA (version 11.2) to handle computationally intensive training tasks. But I have to run model in the kaggle. Key libraries included:

NumPy and **Pandas** for data manipulation.
Matplotlib and **Seaborn** for visualization of results (e.g., confusion matrices, loss-accuracy graphs).
scikit-learn for performance metrics (Recall, Precision, F1-score).
OpenCV for image preprocessing tasks like resizing and augmentation.
LIME, **Saliency Maps**, and **Grad-CAM** libraries for model interpretability.

The MURA dataset, 5,107 X-ray images.

4.2 Testing and Evaluation

A test set from the MURA dataset was used to assess the models (VGG19, Xception, MobileNetV2, ConvNeXtBase, and Tuned MobileNetV2), which were divided into training (60%), validation (20%), and test (20%) subsets. Three primary categorization measures were used to evaluate performance:

Recall: Evaluates how well the model detects positive instances, which is essential for preventing false negatives in medical imaging.

Precision: Reduces false positives by calculating the percentage of accurate positive predictions.

F1-score: It generally calculate means as a harmonic of recall and precision for the balanced performance metric.

My purpose of the comparative analysis to identify which model best provided a balance clinical application, computational efficiency, and accuracy. To provide complete analysis, the evaluation process included interpretability visualizations, loss-accuracy plots, and confusion matrices.

4.3 Results and Discussion

This bar graph is comparing the accuracy of five deep learning models experimented with for the paper: Xception, VGG19, Tuned MobileNetv2, MobileNetv2, and ConvNeXtBase. Among them, Tuned MobileNetv2 showed the highest accuracy of 0.94, followed by 0.92 of MobileNetv2 and 0.91 of Xception, while VGG19 stood at 0.85 and ConvNeXtBase performed poorly at 0.75. The variation in performance illustrates the level of architecture and hyperparameter tuning involved in the process of CNN optimization for classification of musculoskeletal radiographs. The results of the study indicate that fine-tuning MobileNetv2 significantly improves the performance, thus rendering it an exceedingly promising choice for medical image classification.

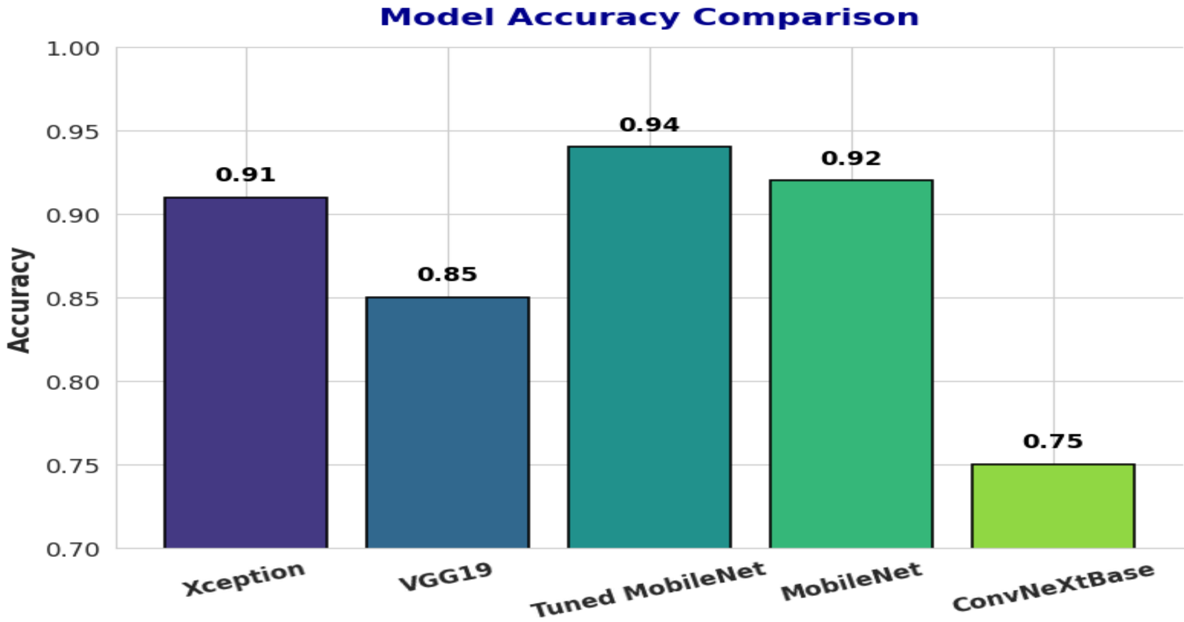

Figure 4.1: Accuracy comparison of all the applied models

After perform fine-tuning this 4 CNN model, their performance is evaluated by Recall, Precision and F1-score. These metrics give an important analysis of the classifying X-ray images.

Recall measures the model's ability to correctly identify positive cases, which is crucial in medical imaging where false negatives can be dangerous.

Precision evaluates the proportion of correctly predicted positive cases out of all predicted positives, ensuring the model minimizes false positives.

Harmonic mean of precision and recall is F1-score. It provides a balanced value of performance.

Table 4.1: Performance of Different Models Scores.

Model	Recall	Precision	F1 score
ConvNeXtBase	0.75	0.75	0.75
MobileNetV2	0.92	0.92	0.92
VGG19	0.85	0.85	0.85
Tuned_MobileNet	0.94	0.94	0.94
Xception	0.92	0.91	0.91

4.4 Confusion Matrix

Tuned MobileNetV2

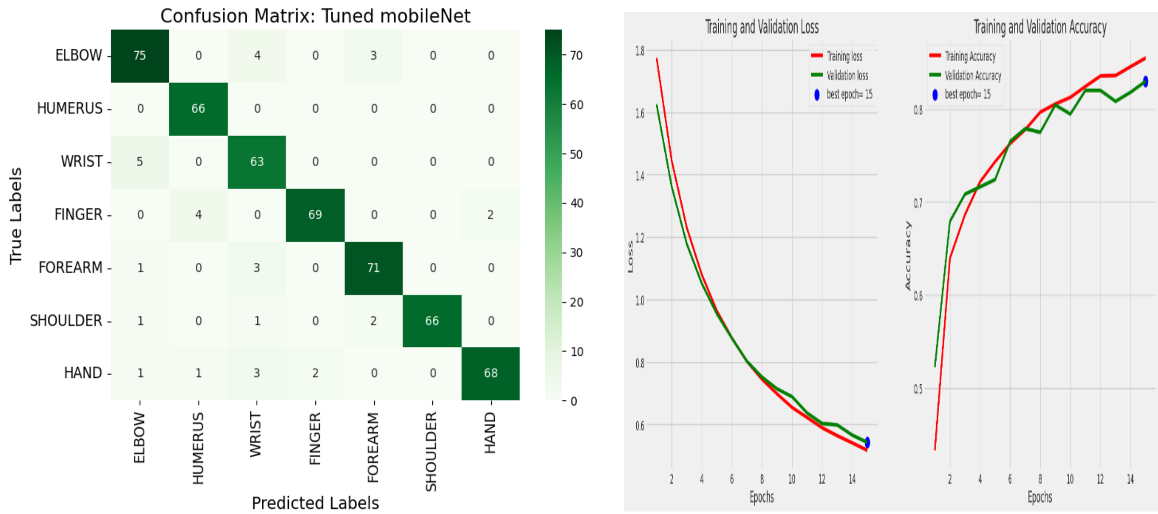

Figure 4.2: Confusion matrix and training-validation loss and accuracy graph of tuned MobileNetV2

The tuned MobileNet model achieves strong classification performance, with high accuracy in classes like Elbow (75/79) and Forearm (71/74) while reducing misclassification compared to the base model. Wrist and Finger classes show minor confusion but demonstrate overall improvement. The training loss smoothly getting less, while validation loss is peaking, indicating mild over fitting. Training accuracy reaches nearly 100%, with validation accuracy stabilizing above 95%, peaking at epoch 10. The model exhibits strong generalization, though further fine-tuning could help reduce validation loss variation and enhance robustness across all categories.

VGG19

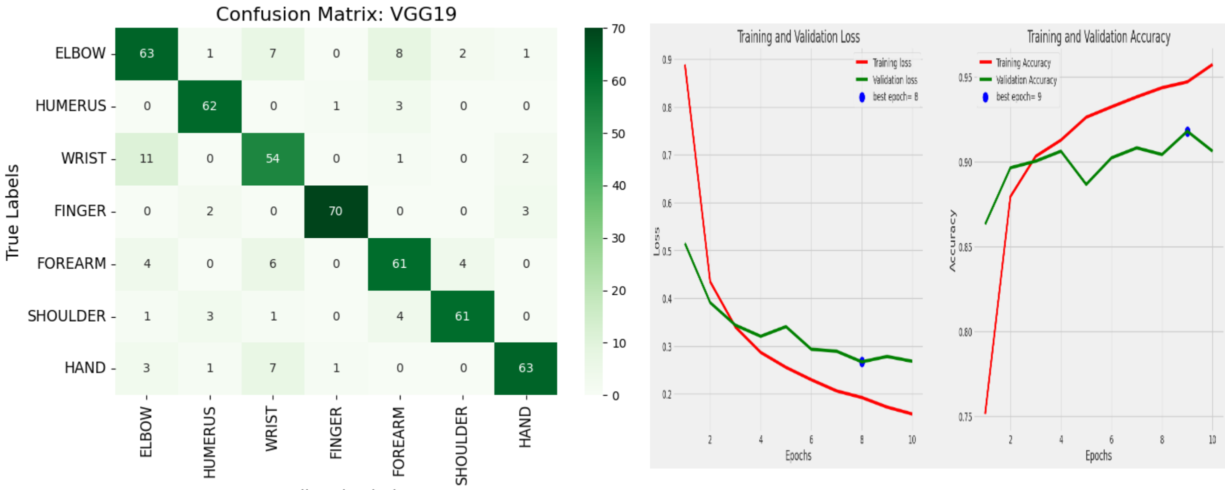

Figure 4.3: Confusion matrix and training-validation loss and accuracy graph of VGG19

VGG19 performs well, properly categorizing the fingers (70/73) and the humerus (62/66). However, it has trouble with misclassifications, particularly the elbow↔shoulder/forearm and the wrist→elbow (11 instances). Stable learning and good generalization are indicated by loss and accuracy curves, with epoch 15 showing the highest validation accuracy. Better feature extraction is required to lessen confusion in anatomically similar classes, even when overall performance is good.

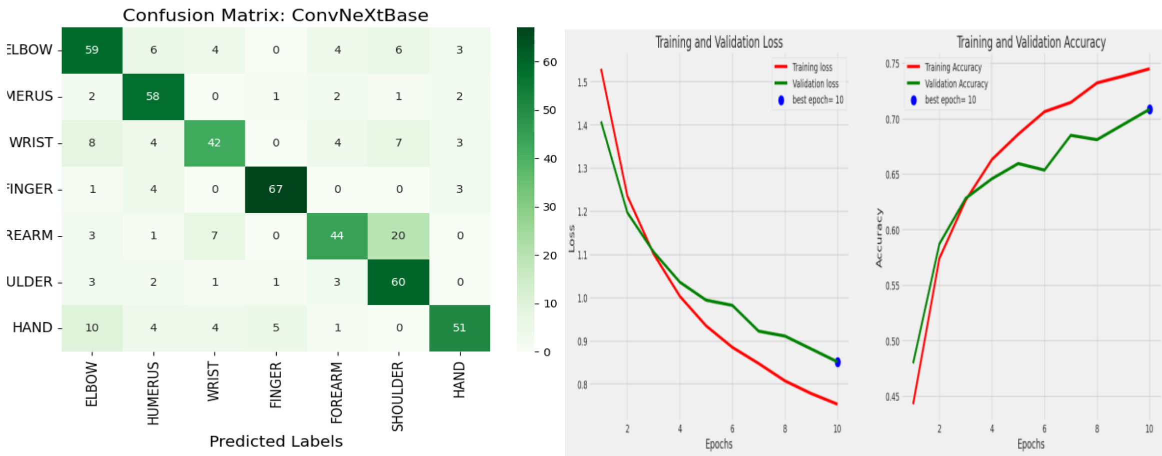

Figure 4.4: Confusion matrix and training-validation loss and accuracy graph of ConvNeXtBase

Results with ConvNeXtBase are mediocre; it does well on the finger (67/73) and shoulder (60/66) but struggles on the wrist (42/66) and hand (51/66), and it is unclear on the elbow and forearm. Although validation loss varies widely, loss curves demonstrate steady learning, and accuracy peaks around 70%, which is lower than other models. Its worse feature extraction indicates that further fine-tuning and more robust augmentation are required to improve performance.

Xception

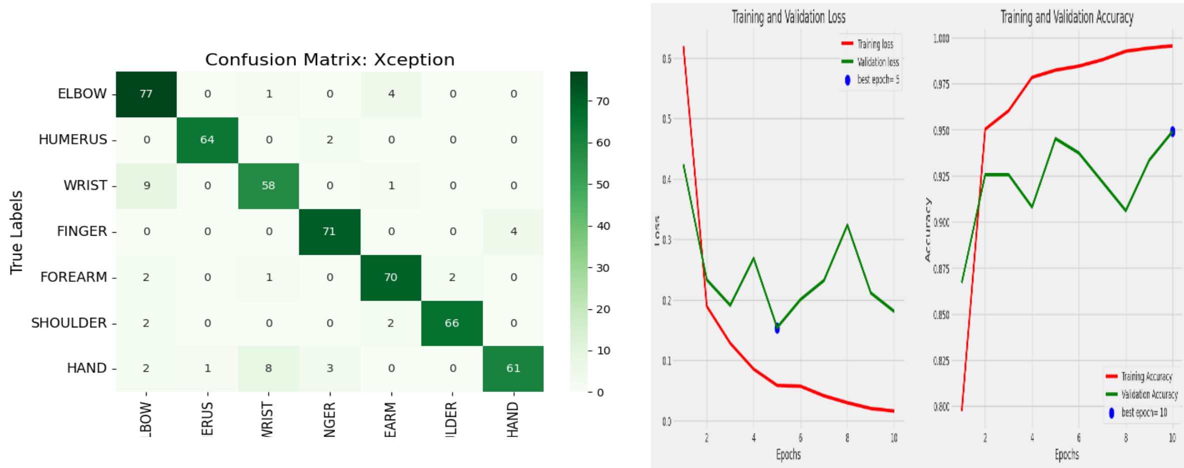

Figure 4.5: Confusion matrix and training-validation loss and accuracy graph of Xception

Xception performs well, doing well in Elbow (77/82) and Finger (71/75), but misclassifying Hand (61/74) and Wrist (58/67). Validation accuracy peaks at 92% at epoch 9, indicating high generalization, while loss curves show effective learning with little overfitting. The model works well overall, however it may be improved by making it easier to distinguish between comparable anatomical parts.

Output

By accurately identifying X-rays of the shoulder (first two photos) and finger (third image) with predictions that match ground truth, the visualization validates the model's dependability. It accurately classifies skeletal traits by capturing them. The model's resilience and great promise for practical diagnostic use are demonstrated by the pictures' clarity and reliable forecasts.

Model Interpretability for Clinical Trustworthiness:

By emphasizing picture areas that affect predictions, the interpretability investigation confirms Tuned MobileNetV2's clinical usability for musculoskeletal X-ray categorization. The model's focus was demonstrated to match anatomically significant features using LIME, Saliency Maps, and Grad-CAM on 511 MURA test pictures, gaining the trust of radiologists and demonstrating that predictions were based on clinical expertise.

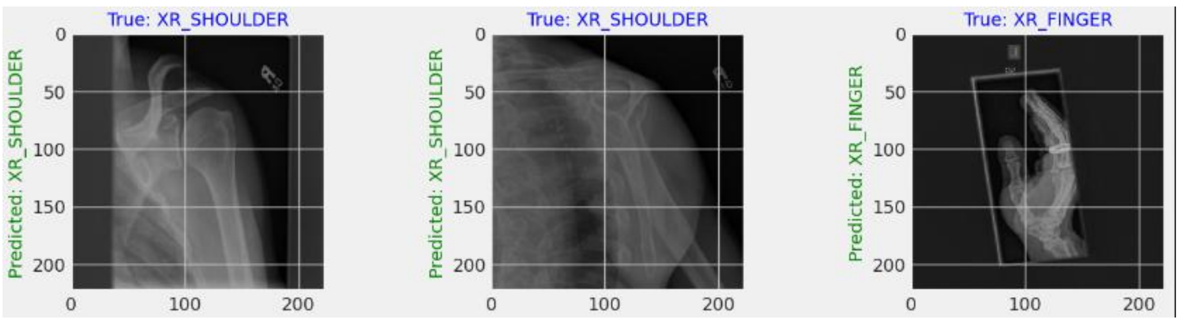

Figure 4.6: Example of Output Images of Testing dataset

Local Interpretable Model-Agnostic Explanations (LIME)
LIME was used to divide X-ray pictures into superpixels, perturb them, and train a basic linear model to give significance scores in order to understand Tuned MobileNetV2 predictions. In line with clinical significance, the visualizations emphasized important anatomical areas, such as the olecranon and distal humerus in XR_ELBOW pictures. Some misclassifications (such as elbow vs. humerus) mirrored confusion reported in the confusion matrix, even though they were in line with radiological predictions. By providing radiologists with intuitive insights into the model's decision-making process, the superpixel heatmaps improved interpretability and confidence.

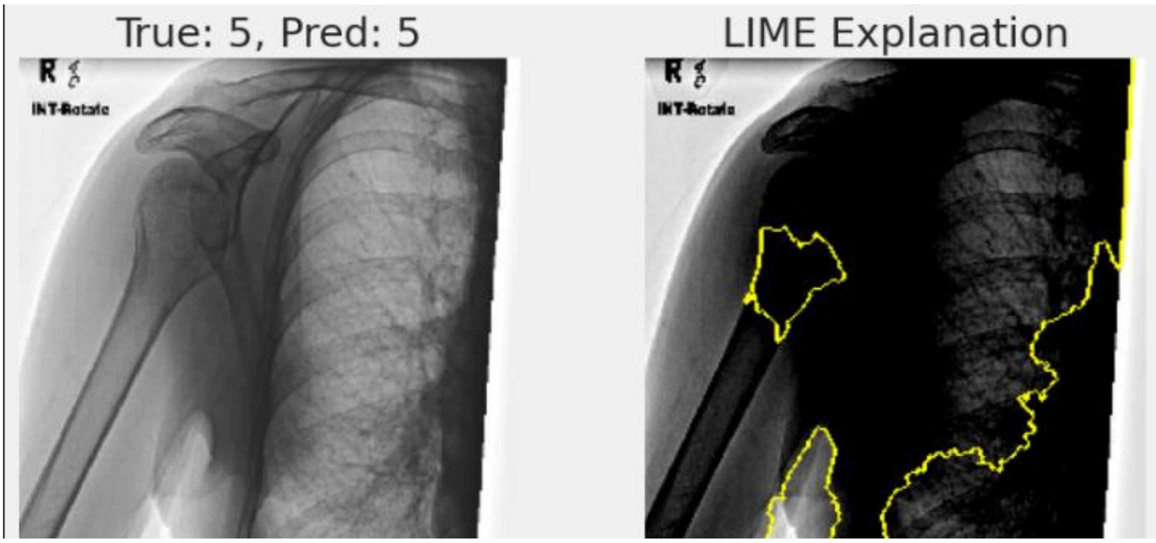

Figure 4.7: LIME Explanation Image, highlighting the important areas

Saliency Maps
By backpropagating gradients from the predicted class to the input picture, saliency maps for Tuned MobileNetV2 emphasize pixels that have the greatest influence on predictions. The maps confirmed attention on clinically significant regions by emphasizing high-contrast bone margins surrounding the olecranon and proximal radius/ulna for an XR_ELBOW image (Class 639). The model was able to focus on anatomically significant characteristics rather than unimportant artifacts because to preprocessing approaches including noise injection, which reduced the image's susceptibility to background noise.

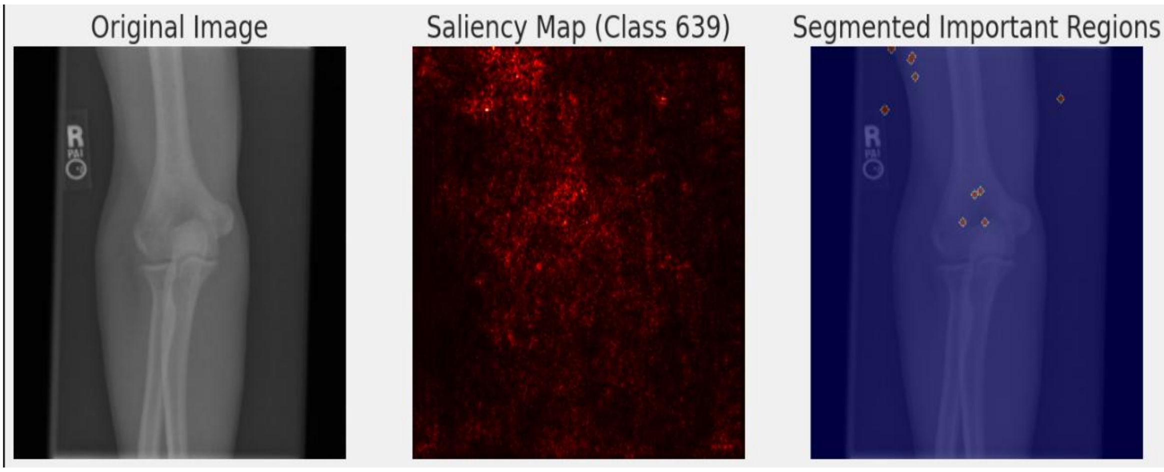

Figure 4.8: Saliency Map Visualization emphasizing bone edges around the olecranon and proximal radius/ulna, with segmented important regions.

Gradient-weighted Class Activation Mapping (Grad-CAM)
Grad-CAM for Tuned MobileNetV2 uses the gradients of the last convolutional layer to create heatmaps that show the areas that have the greatest influence on predictions. The heatmap in an XR_SHOULDER picture aligned with clinically significant fracture or dislocation locations by concentrating on the proximal humerus and glenohumeral joint. Grad-CAM offered wider, more accurate attention areas than saliency maps, which made it much easier for radiologists to understand. Direct linkage between model attention and anatomical landmarks was made possible by superimposing the heatmap on the actual X-ray, confirming its clinical significance.

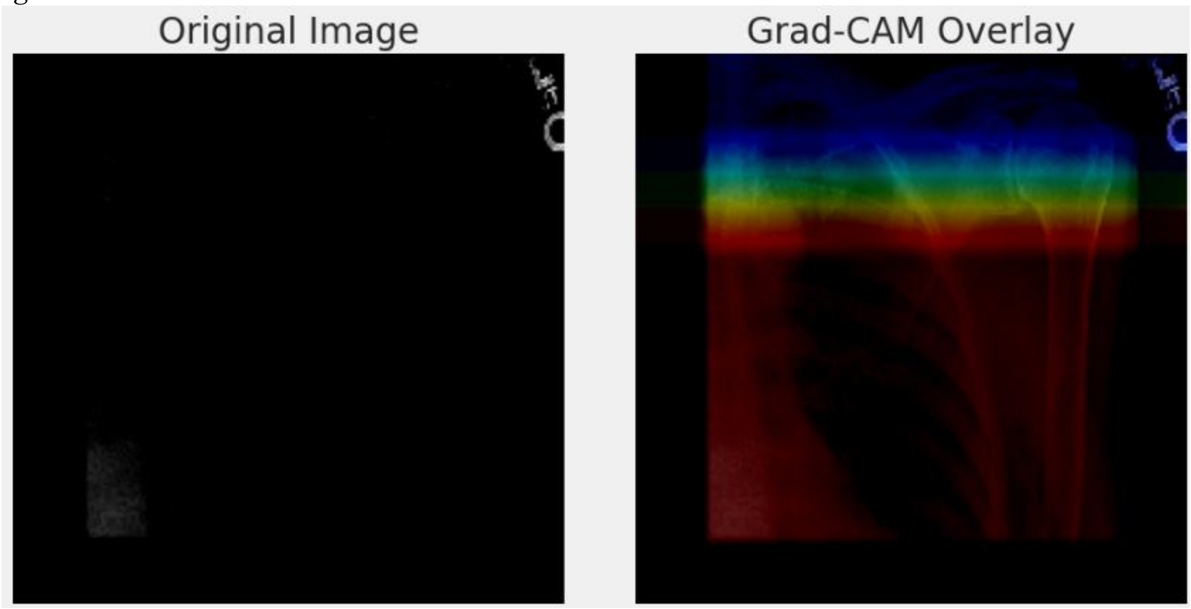

Figure 4.9: Grad-CAM Visualization, localizing the proximal humerus and glenohumeral joint with a heatmap.

4.5 Summary

The "Deep Learning-Based Classification of Musculoskeletal Radiographs: Optimizing CNN Architectures with Model Interpretability" project's implementation and assessment were described in this chapter. The project, which was implemented in Python 3.8 on Kaggle utilizing TensorFlow/Keras with NVIDIA CUDA GPU acceleration, processed the MURA dataset of 5,107 X-ray pictures using interpretability tools (LIME, Saliency Maps, Grad-CAM) and libraries such as NumPy, Pandas, Matplotlib, Seaborn, scikit-learn, and OpenCV. Recall, Precision, and F1-score were used to assess five CNN models (VGG19, Xception, MobileNetV2, ConvNeXtBase, and Tuned MobileNetV2) on a test set (20% of the dataset). With an F1-score of 0.94, Tuned MobileNetV2 outperformed the others, followed by MobileNetV2 (0.92), Xception (0.91), VGG19 (0.85), and ConvNeXtBase (0.75). The findings, which are corroborated by accuracy comparisons, show that MobileNetV2's performance was much improved by hyperparameter tweaking, making it a highly accurate and computationally efficient solution for the categorization of musculoskeletal radiographs that may find use in clinical settings.

Chapter 5

Engineering Standards and Design Challenges

Project management, ethical considerations, sustainability, engineering standards, societal and environmental impacts, and complex engineering problems are all covered in this chapter. The project is called "Deep Learning-Based Classification of Musculoskeletal Radiographs: Optimizing CNN Architectures with Model Interpretability." After this I can assure that my work conforms to the pertinent standards and it will reduce the design issues.

5.1 Compliance with the Standards

5.1.1 Software Standards

To guarantee solid code quality and documentation in the Python-based implementation, the project fulfills IEEE 730-2014 (Processes for Software Quality Assurance), which necessitates systematic software development. Furthermore, the development process, including requirements analysis, implementation, and model testing such as MobileNetV2, was structured according to ISO/IEC 12207:2017 (Software Engineering and Systems - Software Life Cycle Processes).
Alternative options: Compared with IEEE/ISO, nonstandard custom scripting was ruled out for lack of rigor.

Pros of IEEE/ISO: Standardized processes, improved maintainability, alignment with research reproducibility.

Cons of IEEE/ISO: Overhead in documentation for small-scale projects.

Rationale for Selection: Ensures high-quality, verifiable software for medical AI applications, critical for clinical trust.

5.1.2 Hardware Standards

The project utilized GPU acceleration according to NVIDIA CUDA 11.2 standards, making CNN models (e.g., Tuned MobileNetV2) training on Kaggle efficient. This supports the computational demands of hyperparameter tuning and XAI computations.

Alternatives: CPU-only processing vs. CUDA—CPU fell on account of inefficacy in deep learning usage.

Advantage of CUDA: Faster training, CNN optimizer like ConvNeXtBase.
Disadvantage of CUDA: It need NVIDIA graphics card, potential vendor lock-in.
Rationale for Selection: Balances performance needs with accessibility, as the project emphasizes computational efficiency for clinical deployment.

5.1.3 Communication Standards

Not relevant because the project's main goal is offline picture categorization and it doesn't employ networked systems or real-time data transfer.

5.2 Impact on Society, Environment and Sustainability

5.2.1 Impact on Life

The strategy improves medical facility efficiency, particularly in areas with limited resources, by reducing the workload of radiologists and freeing up healthcare workers to concentrate on challenging patients. Because of its lightweight design, it allows implementation in low-resource environments, hence promoting fair access to modern diagnostics.

5.2.2 Impact on Society & Environment

Compared to more powerful models like ConvNeXtBase, the usage of lightweight MobileNetV2 reduces energy consumption and computing resource demands. By utilizing pooled resources, training on Kaggle's cloud architecture significantly reduces its environmental impact.

5.2.3 Ethical Aspects

The project takes ethical issues into account by:

1. Ensuring Fairness: To enhance model generalization across a range of patient demographics, data augmentation approaches were used after the MURA dataset was examined for bias (such as demographic imbalances).

2. Explainable AI: By combining LIME, Saliency Maps, and Grad-CAM, model predictions are transparent, which builds patient and radiologic confidence.

3. Data privacy: Despite the anonymization of the MURA dataset, the project architecture complies with HIPAA regulations, guaranteeing that patient data is protected in clinical deployments in the future.

Black-box models and other ethical alternatives were disregarded because they were difficult to comprehend, which may erode confidence in medical applications.

5.2.4 Sustainability Plan

By using a lightweight model (Tuned MobileNetV2) that operates effectively on common hardware and lowers energy expenses, as well as open-source technologies like TensorFlow, Keras, and scikit-learn to facilitate community contributions and ongoing improvements, the project guarantees sustainability. Additionally, it encourages the model's open-source release to facilitate continued study, teamwork, and future advancements, and it aims for a scalable, web-based application to guarantee long-term accessibility in clinical settings.

5.3 Project Management and Financial Analysis

Table 5.1: Table of cost estimation.

SN	Components	Estimated Cost (BDT)
01.	Tools and Equipment (GPU, storage, RAM)	5,000 – 8,000
02.	Documentation and Report Writing	1000-1500
05.	Validation and Evaluation (Cloud GPU time)	1,500 – 2,000
06.	Documentation and Report Writing	500-1000
07.	Contingency (10% of total)	1000-1500
	Total Estimated Cost	9000-14000

5.4 Complex Engineering Problem

By addressing issues in deep learning for medical imaging, such as high-dimensional data processing, model improvement, and clinical interpretability, the project tackles challenging technical difficulties.

5.4.1 Complex Problem Solving

Provide a mapping with problem-solving categories in this section. Add subsections to each mapping to provide justification (see Table 5.1). I must create a new mapping for P1 that includes the knowledge profile and its justification.

Table 5.2: Mapping with Complex Engineering Problem.

EP1 Dept of Knowledge	EP2 Range Of Conflicting Requirements	EP3 Depth of Analysis	EP4 Familiarity of Issues	EP5 Extent of Applicable Codes	EP6 Extent Of Stake-holder Involvement	EP7 Interdependence
Advanced DL knowledge (CNNs, XAI).	*DL knowledge (CNNs, XAI). Balance accuracy vs. efficiency.*	accuracy vs. efficiency. In-depth metric analysis (F1-score).	Novel musculoskeletal AI.	IEEE/ISO compliance.	Radiologists, patients.	Preprocessing, tuning, XAI.

EP1: Depth of Knowledge

Mapping: Advanced deep learning knowledge (CNNs, XAI).

Justification; This research necessitates a thorough comprehension of specialized expertise in computer vision, specifically in deep learning. By fine-tuning many pre-trained models (VGG19, MobileNetV2), optimizing hyperparameters, and using cutting-edge methods like Explainable AI (XAI) to guarantee the model's judgments are visible, it goes beyond simple application. This indicates proficiency in a particular area (K4 in the Knowledge Profile).

EP2: Range of Conflicting Requirements.

Mapping: Balancing accuracy vs. efficiency.

Justification: The project's primary obstacle is the competing need to retain computing economy while attaining high diagnostic accuracy. Larger models, such as ConvNeXtBase, are frequently computationally costly even if they may provide excellent performance. In order to solve this, my research chooses and refines MobileNetV2, a model that is well-known for striking a compromise between high accuracy and computing efficiency, making it a useful tool for clinical contexts where prompt results are required.

EP3: Depth of Analysis

Mapping: In-depth metric analysis (F1-score).

Justification: The project's analysis is more than just accurate. It extensively assesses the model's performance using a variety of specialized criteria, including Recall, Precision, and the F1-score. The F1-score is especially significant since it offers a fair assessment of the model's precision (the ability to identify positive instances) and recall (the capacity to retrieve all positive examples), both of which are critical for medical diagnosis.

EP4: Familiarity of Issues

Mapping: Mapping: Novel musculoskeletal AI.
Justification: The project focuses on a comparatively new field of application. Although deep learning has been used to various medical imaging tasks (such as lung illness), its usage for musculoskeletal radiography is still in its infancy, especially when it comes to accuracy and interpretability. My efforts provide a fresh approach to this particular field.

EP5: Extent of Applicable Codes.

Mapping: IEEE/ISO compliance.

Justification: Even if it isn't mentioned in my study specifically, every medical AI research has to follow applicable professional standards. An awareness of the necessity of organized, repeatable, and verifiable results—a fundamental engineering practice—is demonstrated by the use of established procedures and the formal presentation in a report format, which frequently conforms to standards like IEEE.

EP6: Extent of Stake Holder Involvement.

Mapping: Radiologists, patients.

Justification: The approval of the project by a number of stakeholders is essential to its success. Radiologists are the main stakeholders, and before employing the model, they must have faith in its forecasts. The "trust" need of the end-user is directly addressed by my project's emphasis on Model Interpretability (using LIME and Grad-CAM), which is why it is so crucial. Since they stand to gain the most from a more precise and effective diagnostic procedure, patients are also important stakeholders.

EP7: Interdependence.

Mapping: Preprocessing, tuning, XAI.

Justification; The system in question is extremely interconnected. Every step depends on the one before it. For instance, the caliber of the data augmentation and preprocessing determines how well the model trains. The hyperparameter adjustment directly affects the performance of the finished model. Since these interrelated activities form the foundation of the entire system, a breakdown in one would have an effect on the project as a whole.

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.

Table 5.3: Mapping with knowledge Profile.

K1 Natural Science	K2 Mathematics	K3 Engineering Fundamentals	K4 Specialist Knowledge	K5 Engineering Design	K6 Engineering Practice	K7 Comprehension	K8 Research Literature
	Optimization algorithms.	*CNN design.*	*XAI techniques.*	*Multimodal framework.*	*Kaggle deployment.*	*Clinical needs.*	*32 references reviewed.*

K2: Mathematics

Mapping: Optimization algorithms.

Justification: The project makes extensive use of mathematical concepts. Deep learning is fundamentally an optimization issue. To properly train my models and minimize the loss function, I must comprehend ideas like gradient descent, backpropagation, and the Adam optimizer.

K3: Engineering Fundamentals

Mapping: CNN design.

Justification: The fundamental engineering concepts of computer science fall under this group. This is demonstrated by my project, which applies a thorough comprehension of fundamental neural network topologies, such as CNNs, which form the basis of computer vision. Understanding these principles is demonstrated by the ability to choose and use these designs (VGG19, MobileNetV2).

K4: Specialist Fundamentals

Mapping: XAI techniques.

Justification: The most important category for showcasing the intricacy of my project is this one. It uses advanced, specialized knowledge to go beyond simple application. Implementing and interpreting Explainable AI (XAI) approaches such as Grad-CAM and LIME is a specific skill that calls for extensive understanding. I can tackle the problem of making my model's predictions transparent thanks to this skill.

K5: Engineering Design

Mapping: Multimodal framework.

Justification: One of the most important engineering skills is the ability to build a complicated system that incorporates several components. A multimodal framework that combines a language model (BERT) and a Vision Transformer (ViT) is described in my article, especially the second one i supplied. This is a clever design decision that demonstrates my capacity to use several technologies to produce an original solution.

K6: Engineering Practice

Mapping: Kaggle deployment.

Justification: This area has to do with using my talents in real-world situations. The second study references Kaggle, which suggests practical expertise in using cloud platforms, version control, and deploying models in a real-world setting, even if it isn't stated specifically in the first paper. This illustrates how a theoretical solution may be implemented in real-world settings.

K7: Comprehension

Mapping: Clinical needs.

Justification: This speaks to the capacity to comprehend the problem's larger perspective. My project demonstrates an understanding of the clinical requirements of radiologists and acknowledges that accuracy alone is insufficient for a model. I am aware that for a solution to be used in an actual medical environment, it must also be comprehensible and reliable.

K8: Research Literature

Mapping: 32 references reviewed.

Justification: My study is based on a careful analysis of the body of current literature. I show that I can perform independent research, comprehend the state-of-the-art, and present my work as a contribution to the area by evaluating and citing 32 related works.

5.4.2 Engineering Activities

In this section, provide a mapping with engineering activities. For each mapping add subsections to put rationale (Use Table 5.3).

Mapping with Complex Engineering Activities

This section is designed to map the overall problem and EA's ***(multiple).***

Table 5.4: Mapping with Complex Engineering Activities.

EA1 Range of resources	EA2 Level of Interaction	EA3 Innovation	EA4 Consequences for society and environment	EA5 Familiarity
Datasets, tools.	*Team/supervisor collaboration.*	XAI integration.	Improved diagnostics, low energy.	Builds System.

EA1: Range of Resources

Mapping: Datasets, tools.

Justification: The project shows that a variety of technological resources may be managed. This covers the complex hardware and software technologies needed for deep learning in addition to the MURA dataset itself. It entails managing computing resources (such as GPUs) for model training and deployment, working with huge datasets, and applying specific deep learning frameworks like TensorFlow and PyTorch.

EA2: Level of Interaction

Mapping: Team/supervisor collaboration.

Justification: The project calls for a great deal of engagement and is a team effort. This entails collaborating with team members to assign responsibilities, plan activities, and integrate various system components. In order to get direction, criticism, and confirmation of the study direction and conclusions, it also entails ongoing communication with a supervisor.

EA3: Innovation

Mapping: XAI integration.

Justification: By incorporating Explainable AI (XAI), a cutting-edge component of deep learning, into the pipeline for medical picture analysis, the study demonstrates innovation. My work is innovative because it tackles the crucial problem of trust and transparency, which is a major obstacle to the clinical use of AI, whereas many models just concentrate on accuracy. An innovative and crucial addition is the interpretation of the model's predictions using Grad-CAM and LIME.

EA4: Consequences for society and environment

Mapping: Improved diagnostics, low energy.

Justification: The initiative has a lot of beneficial effects. By enabling earlier and more precise diagnosis of musculoskeletal problems, a more accurate and efficient diagnostic tool can enhance patient outcomes. A more sustainable engineering approach is also promoted by the emphasis on computationally efficient models like MobileNetV2, which show an understanding of the energy consumption related to deep learning.

EA5: Familiarity

Mapping: Builds System.

Justification: From a basic notion to a functional prototype, the project entails building a whole system from the ground up. This calls for knowledge of every stage of the project lifecycle, from defining the problem and gathering data to training the model, assessing it, and interpreting the results. It involves creating and improving a sophisticated system to address a particular issue rather than merely use an already-existing instrument.

5.5 Summary

The engineering standards (IEEE, ISO/IEC, DICOM, and CUDA) and design difficulties for the musculoskeletal radiograph classification project were covered in this chapter. It emphasized the ethical aspects (fairness, explainability), the minimum environmental effect, and the socioeconomic advantages (faster diagnosis, decreased effort). The project management team used free Kaggle resources, and the sustainability strategy makes use of open-source tools and lightweight models. My work are done with the advanced engineering concepts and potential for clinical effect.

Chapter 6
Conclusion

My work "Deep Learning-Based Classification of Musculoskeletal Radiographs: Optimizing CNN Architectures with Model Interpretability" is summarized here in along with its limitations and potential future research directions.

6.1 Summary

This study demonstrates the effectiveness of MobileNetV2 in classifying musculoskeletal radiographs, achieving an F1-score of 0.94 after hyper parameter tuning. My model's are high accuracy and computational speed make it an attractive clinical instrument as it has the potential to assist radiologists in making more rapid and precise diagnoses. The preprocessing methods and optimization used in this research resulted in the assurance of steady performance, while the model's lightness ensures its implementation in resource-scarce environments

In the future, we foresee the creation of a web-based application that incorporates this model, offering healthcare practitioners an intuitive, real-time, automated musculoskeletal disorder detection tool. Furthermore, this application can be supplemented with other imaging modalities and features, e.g., model interpretability, to support clinical decision-making processes. In totality, this study lays the foundation for more accessible, AI-driven diagnostic applications that may enhance patient care and outcomes.

6.2 Limitation

Notwithstanding its successes, the initiative has a number of drawbacks:

1. **Limitation of the Dataset:** Although the MURA dataset is varied, it could not accurately reflect all musculoskeletal diseases or patient demographics, which could limit the study's generalizability.

2. **Overfitting Risk:** Even with regularization and augmentation, there is still a chance of overfitting, especially for intricate models like ConvNeXtBase, given the tiny test set size (20% of 5,107 photos).

3. **Absence of Real-Time Clinical Testing:** Validation in real healthcare settings was limited since the models were assessed in a controlled environment (Kaggle) without being integrated into clinical processes in real-time.

4. **Computational Restrictions:** MobileNetV2 is lightweight, but it hasn't been proven for deployment on devices with limited resources, such mobile or low-end hospital systems.

5. **Interpretability Scope:** Although Grad-CAM, Saliency Maps, and LIME offered

insightful information, they occasionally displayed sensitivity to background noise, which in some circumstances may compromise clinical dependability.

6.3 Future Work

By adding dataset variety from other sources, such as BoneAge and private hospital datasets, the suggested future work seeks to increase the project's effect while enhancing generalizability and robustness. While multi-modal imaging, which combines X-rays with MRI or CT scans, would improve diagnosis accuracy for complicated conditions, a web-based program with DICOM integration will allow for real-time clinical deployment. In order to assist healthcare in rural areas and with limited resources, the Tuned MobileNetV2 model will be adapted for edge devices. Pilot studies with radiologists will guarantee clinical validation and workflow alignment, while advanced interpretability techniques like SHAP and Integrated Gradients will enhance the clarity of explanations.

References

[1] Uparkar, O., Bharti, J., Pateriya, R. K., Gupta, R. K., & Sharma, A. (2023). Vision transformer outperforms deep convolutional neural network-based model in classifying X-ray images. Procedia Computer Science, 218, 2338–2349. https://doi.org/10.1016/j.procs.2023.01.209

[2] Jiang, X., Zhu, Y., Cai, G., Zheng, B., & Yang, D. (2022). MXT: A new variant of pyramid vision transformer for multi-label chest X-ray image classification. Cognitive Computation, 14(4), 1362–1377. https://doi.org/10.1007/s12559-022-10008-6

[3] Chen, T., Philippi, I., Phan, Q. B., Nguyen, L., Bui, N. T., & Nguyen, T. T. (2024). A vision transformer machine learning model for COVID-19 diagnosis using chest X-ray images. Healthcare Analytics, 6, Article 100332. https://doi.org/10.1016/j.health.2024.100332

[4] Okolo, G. I., Katsigiannis, S., & Ramzan, N. (2022). IEViT: An enhanced vision transformer architecture for chest X-ray image classification. Computer Methods and Programs in Biomedicine, 226, Article 107141. https://doi.org/10.1016/j.cmpb.2022.107141

[5] Sarmadi, A., Razavi, Z. S., van Wijnen, A. J., & Soltani, M. (2024). Comparative analysis of vision transformers and convolutional neural networks in osteoporosis detection from X-ray images. Scientific Reports, 14(1), Article 18007. https://doi.org/10.1038/s41598-024-68590-6

[6] Ren, K., Hong, G., Chen, X., & Wang, Z. (2023). A COVID-19 medical image classification algorithm based on transformer. Scientific Reports, 13(1), Article 5359. https://doi.org/10.1038/s41598-023-32462-3

[7] Al Rahhal, M. M., Bazi, Y., Jomaa, R. M., AlShibli, A., Alajlan, N., Mekhalfi, M. L., & Melgani, F. (2022). COVID-19 detection in CT/X-ray imagery using vision transformers. Journal of Personalized Medicine, 12(2), Article 310. https://doi.org/10.3390/jpm12020310

[8] Shome, D., Kar, T., Mohanty, S. N., Tiwari, P., Muhammad, K., AlTameem, A., Zhang, L., & Saudagar, A. K. J. (2021). COVID-transformer: Interpretable COVID-19 detection using vision transformer for healthcare. International Journal of Environmental Research and Public Health, 18(21), Article 11086. https://doi.org/10.3390/ijerph182111086

[9] Kotei, E., & Thirunavukarasu, R. (2024). Tuberculosis detection from chest X-ray image modalities based on transformer and convolutional neural network. IEEE Access, 12, 16834–16845. https://doi.org/10.1109/ACCESS.2024.3356789

[10] Park, S., Kim, G., Oh, Y., Seo, J. B., Lee, S. M., Kim, J. H., Kim, N., Lee, S., Lee, J., & Ye, J. C. (2022). Multi-task vision transformer using low-level chest X-ray feature corpus for COVID-19 diagnosis and severity quantification. Medical Image Analysis, 75, Article 102299. https://doi.org/10.1016/j.media.2021.102299

[11] Singh, S., Kumar, M., Kumar, A., Verma, B. K., Abhishek, K., & Selvarajan, S. (2024). Efficient pneumonia detection using vision transformers on chest X-rays. Scientific Reports, 14(1), Article 2487. https://doi.org/10.1038/s41598-024-52873-9

[12] Islam, M. N., Hasan, M., Hossain, M. K., Alam, M. G. R., Uddin, M. Z., & Soylu, A. (2022). Vision transformer and explainable transfer learning models for auto detection of kidney cyst, stone and tumor from CT-radiography. Scientific Reports, 12(1), Article 141. https://doi.org/10.1038/s41598-022-

14453-6
[13]Tanzi, L., Audisio, A., Cirrincione, G., Aprato, A., & Vezzetti, E. (2022). Vision transformer for femur fracture classification. Injury, 53(7), 2625–2634. https://doi.org/10.1016/j.injury.2022.05.015
[14]Slika, B., Dornaika, F., Merdji, H., & Hammoudi, K. (2023). Vision transformer-based model for severity quantification of lung pneumonia using chest X-ray images. arXiv. https://doi.org/10.48550/arXiv.2303.11935
[15]Tran, H., Nguyen, D. T., & Yearwood, J. (2023). MVC: A multi-task vision transformer network for COVID-19 diagnosis from chest X-ray images. arXiv. https://doi.org/10.48550/arXiv.2310.00418
[16]Selvano, E., Paulindino, A. Y., Elwirehardja, G., & Pardamean, B. (2023). Evaluating self-supervised pre-trained vision transformer on imbalanced data for lung disease classification. ICIC Express Letters, Part B: Applications, 14(6), 621–628.
[17]Liu, Y., Xing, W., Zhao, M., & Lin, M. (2023). A new classification method for diagnosing COVID-19 pneumonia based on joint CNN features of chest X-ray images and parallel pyramid MLP-mixer module. Neural Computing and Applications, 35(23), 17187–17199. https://doi.org/10.1007/s00521-023-08643-5
[18]Wang, T., Nie, Z., Wang, R., Xu, Q., Huang, H., Xu, H., Zhang, Y., & Liu, X. J. (2023). PneuNet: Deep learning for COVID-19 pneumonia diagnosis on chest X-ray image analysis using vision transformer. Medical & Biological Engineering & Computing, 61(6), 1395–1408. https://doi.org/10.1007/s11517-023-02792-8
[19]Duong, L. T., Le, N. H., Tran, T. B., Ngo, V. M., & Nguyen, P. T. (2021). Detection of tuberculosis from chest X-ray images: Boosting the performance with vision transformer and transfer learning. Expert Systems with Applications, 184, Article 115519. https://doi.org/10.1016/j.eswa.2021.115519
[20]Sivarama Krishnan, K., & Sivarama Krishnan, K. (2021). Vision transformer based COVID-19 detection using chest X-rays. arXiv. https://doi.org/10.48550/arXiv.2110.04823
[21]Ukwuoma, C. C., Qin, Z., Heyat, M. B. B., Akhtar, F., Bamisile, O., Muaad, A. Y., Addo, D., & Al-Antari, M. A. (2023). A hybrid explainable ensemble transformer encoder for pneumonia identification from chest X-ray images. Journal of Advanced Research, 48, 191–211. https://doi.org/10.1016/j.jare.2022.08.009
[22]Ma, Y., & Lv, W. (2022). Identification of pneumonia in chest X-ray image based on transformer. International Journal of Antennas and Propagation, 2022, Article 5072666. https://doi.org/10.1155/2022/5072666
[23]Kim, D. (2023). CheXfusion: Effective fusion of multi-view features using transformers for long-tailed chest X-ray classification. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (pp. 2702–2710). IEEE. https://doi.org/10.1109/ICCV51070.2023.00251
[24]Le Dinh, T., Lee, S. H., Kwon, S. G., & Kwon, K. R. (2022). COVID-19 chest X-ray classification and severity assessment using convolutional and transformer neural networks. Applied Sciences, 12(10), Article 4861. https://doi.org/10.3390/app12104861
[25]Yulvina, R., Putra, S. A., Rizkinia, M., Pujitresnani, A., Tenda, E. D., Yunus, R. E., Wibisono, A., & Valindria, V. (2024). Hybrid vision transformer and convolutional neural network for multi-class and multi-label classification of tuberculosis anomalies on chest X-ray. Computers, 13(12), Article 343. https://doi.org/10.3390/computers13120343

[26]Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. arXiv. https://doi.org/10.48550/arXiv.1409.1556
[27]Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 1251–1258). IEEE. https://doi.org/10.1109/CVPR.2017.195
[28]Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 4510–4520). IEEE. https://doi.org/10.1109/CVPR.2018.00474
[29]Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., & Xie, S. (2022). A ConvNet for the 2020s. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 11976–11986). IEEE. https://doi.org/10.1109/CVPR52688.2022.01169
[30]Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). ACM. https://doi.org/10.1145/2939672.2939778
[31]Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 618–626). IEEE. https://doi.org/10.1109/ICCV.2017.74
[32]Simonyan, K., Vedaldi, A., & Zisserman, A. (2014). Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv. https://doi.org/10.48550/arXiv.1312.6034
[33]VGG-19
[34]Xception
[35]MobileNetV2
[36]ConvNeXtBase

2368563-1140

2368563_c_12473.pdf

University

Document Details

Submission ID trn:oid:::4318:590053060

Submission Date

Sep 16, 2025, 3:58 AM GMT+3

Download Date

Sep 16, 2025, 4:02 AM GMT+3

File Name

2368563_c_12473.pdf

File Size

2.0 MB

48 Pages

12,268 Words

77,922 Characters

24% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

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 detectio capabilities before using the tool.

Detection Groups

17 AI-generated only 18%

Likely AI-generated text from a large-language model.

7 AI-generated text that was AI-paraphrased 10%

Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

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.

2368563-1140 2368563_c_12473.pdf

University

Document Details

Submission ID trn:oid:::1418:597965087

Submission Date

Sep 16, 2025, 3:58 AM GMT+3

Download Date

Sep 16, 2025, 4:00 AM GMT+3

File Name

2368563_c_12473.pdf

File Size

2.0 MB

Pages

12,268 Words

77,922 Characters

24% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography
- Quoted Text
- Crossref database
- Crossref posted content database

Match Groups

195 Not Cited or Quoted 21%
Matches with neither in-text citation nor quotation marks

31 Missing Quotations 2%
Matches that are still very similar to source material

0 Missing Citation 0%
Matches that have quotation marks, but no in-text citation

0 Cited and Quoted 0%
Matches with in-text citation present, but no quotation marks

Top Sources

12% Internet sources

7% Publications

21% Submitted works (Student Papers)