MelanomaNet: RL-Guided Hyperparameter Optimization in Deep Meta-Ensembles for Melanoma Image Classification

BY

Jakia Khanom
ID: 241-25-023

This Report Presented in Partial Fulfillment of the Requirements for
The Degree of Master of Science in Computer Science and Engineering

Supervised By

Dr. Md Zahid Hasan
Associate Professor
Department of CSE
Daffodil International University

Co-Supervised By

Mr. Abdus Sattar
Associate Professor
Department of CSE
Daffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY

DHAKA, BANGLADESH

MAY 2025

APPROVAL

This Project titled "**MelanomaNet: RL-Guided Hyperparameter Optimization in Deep Meta-Ensembles for Melanoma Image Classification**", submitted by **Jakia Khanom**, ID No: **241-25-023** 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.

BOARD OF EXAMINERS

Dr. S.M Aminul Haque — **Chairman**
Professor and Associate Head
Department of Computer Science and Engineering
Faculty of Science & Information Technology

Dr. Fizar Ahmed — **Internal Examiner**
Associate Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Dr. Md Alamgir Kabir — **Internal Examiner**
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Nazibur Rahman — **External Examiner**
Technical Lead, Database Administrator
Wipro, Telenor - Grameen Phone Account
Dhaka, Banglades

DECLARATION

I hereby declare that this research has been done by me under the supervision of **Dr. Md Zahid Hasan, Associate Professor, Department of CSE,** Daffodil International University. I 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:

Dr Md Zahid Hasan
Associate Professor
Department of CSE
Daffodil International University

Co-Supervised by:

Mr. Abdus Sattar
Associate Professor
Department of CSE
Daffodil International University

Submitted by:

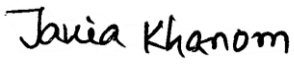

Jakia Khanom
ID: 241-25-023
Department of CSE
Daffodil International University

ACKNOWLEDGEMENT

First, I express my heartiest thanks and gratefulness to Almighty Allah for His divine blessing which makes it possible to complete the final year project/internship successfully.

I am grateful and wish my profound indebtedness to **Dr Md Zahid Hasan, Associate Professor**, Department of CSE, Daffodil International University, Dhaka, deep knowledge & keen interest of my 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.

I would like to express my heartiest gratitude to **Dr. Sheak Rashed Haider Noori**, **Head,** Department of CSE, for his kind help to finish our project and to other faculty members and the staff of the CSE department of Daffodil International University.

Finally, I must acknowledge with due respect the constant support and patience of my parents.

ABSTRACT

Skin cancers, particularly melanoma, constitute a public health issue as they are highly aggressive and deadly if not caught early. This research study proposes an energetically inspired, conscientious deep learning framework for dermoscopic image classification of skin lesions (benign/malignant). Our method relies on an ensemble of state-of-the-art convolutional neural networks combined with expressive morphological and color-based shape descriptors, which are further ensemble learning using ensemble learning, test-time augmentation (TTA), and reinforcement learning (RL) based hyperparameter search. Hybrid deep architecture of EfficientNetV2, NFNet, and ResNet were utilized here, with both focal loss and ArcFace margin product fused at training to balance class imbalance and improve feature discriminability. Shape-aware dual-branch model structure was utilized for fusing visual features and dermatologically important handcrafted features. Validation of the model was carried out with 5-fold cross-validation and large test set validation. The system performed 95.35% classification accuracy, AUC of 0.992, and PR-AUC of 0.992 on the test set with superior diagnostic performance and generalizability. Our findings illustrate the potential of state-of-the-art ensemble AI systems to aid dermatological diagnostics with high reliability and explainability.

.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1
INTRODUCTION

1.1 Introduction

Melanoma is the most dangerous type of skin cancer and is noted for its high metastasis and mortality. Its treatment outcome is significantly improved with early detection, but conventional diagnostic practices—essentially dermatologists' visual examination—are subjectively based, very often leading to delayed or incorrect diagnoses. Dermoscopy has improved the visual resolution of skin lesions yet remains largely clinically experiential-based.

Computer-aided diagnosis (CAD) system development has a revolutionary potential for clinical dermatology. Nonetheless, most current deep learning models experience overfitting, poor generalization, and insufficient integration of clinical domain expertise. Moreover, medical datasets are typically plagued with class imbalance since malignant cases are severely underrepresented relative to benign cases. All these issues call for an integrative framework with active image understanding, domain-invariant features, and adaptive training paradigms.

This work presents a new hybrid method for melanoma classification with a two-branch deep learning system that combines synergistically CNN-extracted visual features and hand-designed shape and color features motivated by dermatology principles. The system is also strengthened by an ensemble learning method, test-time augmentation, and a reinforcement learning-based optimizer that can dynamically alter hyperparameters for best model performance. By systematic experiments and performance evaluation, we present a very accurate and interpretable melanoma detection system, in close connection with clinical requirements and operational feasibility.

1.2 Motivation

The early and precise detection of life-threatening diseases continues to be one of the most pressing issues in the rapidly changing field of modern healthcare, where technological integration is becoming more and more essential. One type of skin cancer where early

diagnosis can significantly change patient outcomes is melanoma, a very aggressive and deadly form of the disease. The healthcare industry is still having difficulty creating diagnostic systems that are not only accurate but also flexible, interpretable, and deployable in a variety of real-world contexts, even with advances in imaging techniques and machine learning.

The urgent need to close this gap by creating a system that mimics the flexibility and learning effectiveness of an experienced clinician is what inspired this dissertation. In tasks like skin lesion classification, traditional machine learning models—especially those driven by deep learning—have demonstrated impressive promise. To control how the model learns from data, these models frequently rely significantly on manually adjusted hyperparameters. Typically, adjusting these parameters entails trial-and-error procedures that call for substantial time, computational resources, and domain expertise. Due in large part to their static nature and poor generalization, these models often perform poorly when applied to new datasets or different clinical environments, even when they perform well in lab settings.

Medical data itself presents the second significant obstacle. Medical image datasets are frequently small and extremely unbalanced, in contrast to data-rich fields like natural image recognition. Compared to benign cases, malignant ones are uncommon, and gathering data is costly, time-consuming, and morally delicate. These drawbacks can result in overfitting and subpar performance in practical situations, which is especially risky in clinical settings where precise diagnosis can mean the difference between life and death.

By enabling models to utilize knowledge from extensive datasets across various domains, transfer learning has surfaced as a partial solution to the data scarcity issue. However, there are disadvantages to transfer learning as well. Most transfer learning approaches still require retraining or fine-tuning when applied to new tasks, making them less suitable for dynamic and continually evolving healthcare environments.

The idea behind this research is that we can revolutionize the way we develop AI systems for medical image analysis by fusing the flexibility of Reinforcement Learning (RL) with the effectiveness of Meta-Learning. By learning from its successes and failures in real time, Reinforcement Learning enables a model to make well-informed decisions about its own

learning process, including which hyperparameters to use. Often called "learning to learn," meta-learning enables the system to quickly adjust to new tasks with little data, using past performance to speed up and enhance performance. The goal of this dissertation is to develop a clever, self-adjusting framework for diagnosing melanoma that is reliable, scalable, and appropriate for use in a variety of clinical contexts. Reducing the need for expert tuning, cutting down on training time, and improving generalization across various medical image types and clinical workflows are the objectives. This work is ultimately driven by the goal of empowering healthcare professionals to use AI as a trustworthy partner rather than merely a tool. A partner who, like a human doctor, grows, learns, and adapts with time. With the goal of producing useful, moral, and significant solutions that significantly improve international healthcare systems and, most importantly, save lives, this vision propels the research forward.

1.3 Rationale of the Study

Despite numerous advancements in AI-assisted medical diagnostics, many systems fail to maintain performance across different clinical settings due to their rigidity in learning and deployment. One of the primary reasons is that traditional deep learning systems rely heavily on static configurations and manually fine-tuned hyperparameters. These systems often perform well in the environments in which they are trained but struggle to generalize to new scanners, institutions, or populations.

This research addresses this limitation by developing an adaptive framework that can automatically tune its learning strategy using Reinforcement Learning and quickly adapt to new data distributions through Meta-Learning. The rationale behind this approach is the need for a diagnostic model that can evolve with the clinical environment without requiring complete retraining. Such adaptability is especially vital in medical imaging domains, where annotated data is scarce, and domain shifts are frequent.

Furthermore, by integrating domain-specific handcrafted features with deep neural features in dual-branch architecture, the system maintains a balance between interpretability and predictive accuracy. This hybrid model structure ensures that the diagnostic decisions are not only accurate but also explainable and clinically meaningful.

Ultimately, the study aims to create a scalable, transferable, and generalizable melanoma diagnosis system that can reduce manual intervention, lower operational costs, and significantly enhance diagnostic precision in both well-resourced and under-resourced medical institutions.

1.4 Research Questions

To guide the research and evaluate the effectiveness of the proposed RL-guided meta-learning system, the following research questions have been formulated:

- **RQ1:** How can Reinforcement Learning and Meta-Learning be combined effectively to enable dynamic and automatic hyperparameter optimization for medical image classification tasks?
- **RQ2:** Can the proposed adaptive framework maintain high classification performance across different imaging devices, hospitals, and data scarcity scenarios without retraining from scratch?
- **RQ3:** Does the hybrid model integrating both handcrafted and deep features outperform traditional single-branch models in terms of accuracy, interpretability, and generalizability?

These questions serve as the foundation for developing, testing, and validating the proposed system in realistic medical imaging contexts, ensuring both scientific rigor and clinical relevance.

1.5 Expected Output

Based on the detailed methodology and experimental framework employed in this study, the following outcomes are expected:

- **Development of MelanomaNet:** A dual-branch hybrid model combining deep CNN-based visual feature extraction (EfficientNetV2, ECA-NFNet-L1, ResNet50) with a parallel MLP branch that processes 15 handcrafted dermatological features (shape, texture, color, and ratio-based).
- **Automated Hyperparameter Optimization:** A Reinforcement Learning (RL) agent based on Q-learning dynamically searches optimal configurations for

hyperparameters such as batch size, learning rate, dropout, weight decay, scheduler, and image resolution. The RL agent uses validation performance as its reward signal to adapt the training process in real-time.

- **Robust Generalization under Data Scarcity:** Incorporation of regularization techniques such as focal loss, ArcMarginProduct, and mixup regularization to handle class imbalance and improve feature separability in malignant and benign skin lesion images.
- **Test-Time Augmentation and Ensemble Learning:** Integration of test-time augmentation (TTA) strategies and ensemble inference using the top-k models to ensure prediction stability and robustness under different transformations and data conditions.
- **High Diagnostic Performance:** The model is expected to deliver superior accuracy, sensitivity, and specificity validated through 5-fold cross-validation and large-scale testing. The target metrics include ~95.35% test accuracy, 0.992 ROC AUC, and macro F1-score of ~0.978.
- **Interpretability and Explainability:** By fusing handcrafted features with learned CNN features, the model maintains interpretability aligned with clinical domain knowledge, making the predictions more trustworthy for dermatological professionals.
- **Foundation for Future Research:** The framework offers a scalable, modular, and clinically adaptable solution that can be extended to other medical imaging domains, enabling further exploration in multi-modal learning, federated learning, and explainable AI.

These outcomes aim to contribute to a robust and practical AI tool for clinical melanoma screening while setting a precedent for intelligent model optimization in other medical imaging challenges.

1.6 Project Management and Finance

The research work doesn’t get funds from any individuals or organization.

1.7 Report Layout

In Chapter 1, the introduction, objectives, and key research inquiries of the study are outlined. In Chapter 2, a concise synopsis of the literature review is provided. In Chapter 3, the proposed methodology is described in detail. In Chapter 4, the experimental outcomes of the paper are described and examined. The fifth chapter discusses the sustainability plan, societal and environmental impacts, and ethical considerations. The sixth chapter concludes the present investigation and outlines a strategy for subsequent endeavors.

CHAPTER 2
BACKGROUND

2.1 Preliminaries/Terminologies

Artificial Intelligence (AI) has integrated into medical image analysis, which drastically altered the patterns of diagnostic workflows by facilitating faster and more reliable disease detection. But the effectiveness of deep learning models in this field is heavily subject not only to hyperparameter configurations, but also to domain-specific adaptation and data. To solve such problems, researchers have recently focused on Reinforcement Learning (RL), Meta Learning, Hyperparameter Optimization (HPO), and Transfer Learning. Collectively, the paradigms together provide a strong foundation to construct intelligent, flexible and agile personalized medical imaging systems.

2.2 Related works

Medical image analysis has seen significant progress with the integration of artificial intelligence (AI), particularly deep learning methods. However, challenges such as limited annotated data, high dimensionality, domain heterogeneity, and the need for real-time adaptability persist.

Early works primarily utilized Convolutional Neural Networks (CNNs) for image classification and segmentation tasks. For instance, Chen et al. [10][11] successfully employed CNNs for stroke classification using brain CT scans, showcasing the benefit of transfer learning from large-scale datasets. Nevertheless, such methods often require massive, labeled datasets and suffer from limited flexibility when deployed in new environments.

To address data scarcity, transfer learning has become a go-to strategy. It leverages pre-trained models to extract relevant features for medical images, thus reducing the need for large, annotated datasets. Zhuang et al. [22] and Cheplygina et al. [21] provide comprehensive surveys on the role of transfer learning in medical imaging, highlighting its

efficiency but also its limitations when applied across differing imaging modalities or institutional protocols.

Hyperparameter optimization remains a bottleneck for model performance. Traditional techniques such as grid search and random search [1][4][22] are computationally expensive and inefficient for high-dimensional spaces. Bayesian optimization methods attempt to focus the search in promising regions, while evolutionary strategies such as genetic algorithms [20] offer global search capabilities. However, these approaches still struggle with dynamic data environments, a common scenario in real-world medical settings.

Recent efforts have explored the integration of Reinforcement Learning (RL) to automate hyperparameter tuning. Lorraine et al. [5] and Wang et al. [15] propose RL-based meta-heuristics where an agent learns to adjust hyperparameters based on training performance feedback. Schaer et al. [1] notably demonstrated the benefits of distributed hyperparameter tuning in lung texture classification using Hadoop clusters.

Meta-learning, or "learning to learn," is gaining traction in the context of few-shot medical learning. Finn et al. [8] and Hospedales et al. [11] describe algorithms that quickly adapt to new tasks with minimal data, an ideal trait for medical scenarios with limited annotated images. Saeed et al. [6] and Huang et al. [9] extend this by combining meta-learning with reinforcement learning (meta-RL) to enhance generalization and adaptability, particularly in domains like image quality assessment and multi-modal medical fusion.

Jiang et al. [15] explored multi-teacher frameworks for medical image classification in low-data regimes, while Roth et al. [23] employed self-supervised pretraining for boosting downstream performance on diverse medical tasks.

Despite these advancements, few studies explicitly fuse all three paradigms—reinforcement learning, meta-learning, and transfer learning—in a single, unified framework for hyperparameter optimization in medical imaging. This dissertation aims to bridge this gap by designing an end-to-end adaptive learning pipeline that not only

improves classification performance but also enhances system adaptability and clinical utility.

2.3 The Problem's Scope

Medical image analysis is significant in aiding diagnosis, treatment planning and disease surveillance in clinical practice. CNNs — Deeper and more advanced deep learning models, in particular convolutional neural networks, holds the most recent results for classification, segmentation and anomaly detection task. Nevertheless, their success is very related to hyperparameters such as learning rate, optimizer type, batch size and number of layers in the network that must be tuned carefully. Strategies for hyperparameter tuning (e.g., grid search, Bayesian optimization) are slow, immobile and most of the time not transferable to the new task/dataset.

This work tackles the above limitations by introducing a Reinforcement Learning (RL)-guided meta-learning framework for real-time hyperparameter optimization and mitigating transfer learning losses in medical image analysis. The dream is to create a system that can, while being trained itself, adapt its hyperparameters based on learning signals and benefit experience-acquired knowledge for tackling new yet related medical imaging tasks. Finally this model needs to work even in the hard low-data regimes typical for medical applications as we cannot just mark an infinite amount of data to make annotations. This work combines RL, meta-learning & transfer learning to improve model adaptability/tracking/transfer and speed/scale performance w/ very low manual intervention

2.4 Challenges

Though the zero-shot integration of reinforcement learning, meta-learning and transfer learning is a breakthrough for medical image analysis, this proposed system largely faces several pressing concerns that will have to be mitigated.

The first obvious major difficulty is the high dimensionality and variability inherent in medical images. Medical images contain modality resolution, and acquisition protocol differences across institutions. These discrepancies yield the confounding feature variability that hampers extraction and generalization in models. While performance on

this kind of heterogeneous data demands models which are at once flexible and insensitive to such variations.

Second, labelled medical data is hugely scarce and the root challenge. The cost of obtaining high fidelity annotations typically dicts to be domain experts (radiology or pathologists) which is a time-consuming and costly process. The downside is that training deep learning model well in particular diseases and rarer diseases simply becomes nearly impossible. As a result, techniques that exploit the limited data effectively, like few-shot learning and efficient transfer learning etc are necessary.Thirdly, solving hyperparameter dynamic optimization is quite intricate by itself. In contrast to static tuning procedures which merely search and optimize hyperparameters like learning rate, weight decay for example, dynamic approaches modify them during training.With reinforcement learning, embedding policies of how to give these decisions introduces the difficulties of policy design and reward functional formulation and fast exploration of hyperparameter space; also, while the search space is high-dimension and continuous.Fifthly, generalization in meta-learning is hard as well. Meta-learning is comfortable in learning how to learn across tasks, but the challenge is making sure that one meta-learner does not overfit to too small a space. Catalyst in medical imaging, where tasks can vary in pathology, modality and data quality is a challenging individual for crafting meta-learning algorithms that generalizes on this diversity with limited computation. Ultimately, stability and inefficient sample Ness seen in reinforcement learning severely limits its usefulness. Because RL agents typically need really seeing the environment to train, and their behavior is unstable often more nuanced when rewards are sparse or delayed. Directing the training decisions of neural networks—steering with a RL, e.g.: dynamic hyperparameter adjustment)— needs to be achieved by framework which is stable and efficient from the data perspective to work for restaurants with noisy conditions typically found in medical imaging.

The above challenges serve to illustrate the scale of this problem as well as underscore the importance of designing a high-performing tool for reinforcement learning (RL), meta-learning and transfer learning synergize in tackling some of the difficulties of medical image analysis.

CHAPTER 3
RESEARCH METHOLOGY

3.1 Proposed Methodology

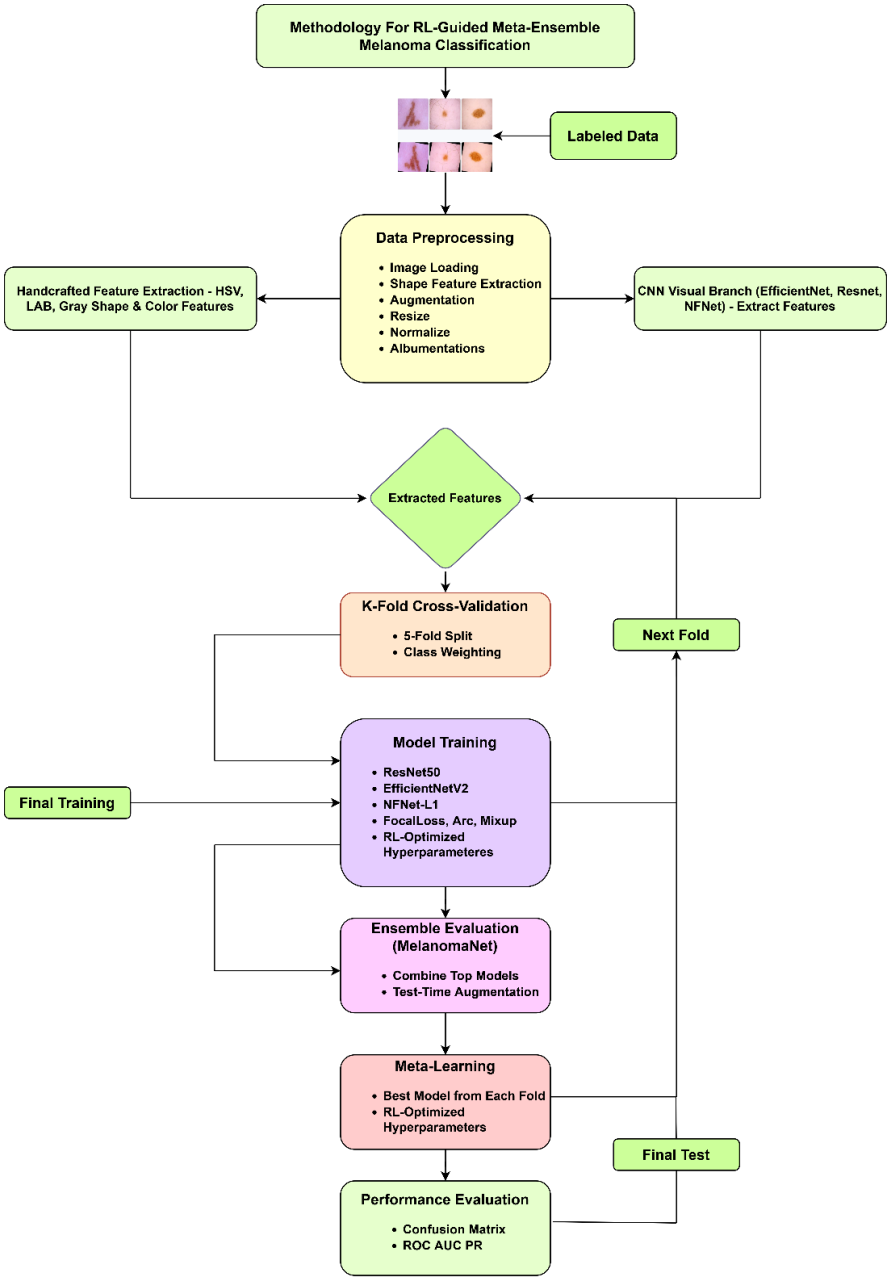

Fig 3.1: The working process to perform melanoma classification framework.

This research adopts a comprehensive and dynamic methodology to tackle the complex problem of melanoma classification in dermoscopic images, emphasizing not only accuracy but also adaptability, interpretability, and robustness in real-world clinical settings. The proposed approach leverages a hybrid learning framework that brings together the strengths of convolutional neural networks (CNNs), domain-specific handcrafted features, and reinforcement learning (RL)-guided hyperparameter tuning to build a system capable of making highly reliable diagnostic decisions. The dataset was divided into training and testing subsets, containing dermoscopic images labeled as either benign or malignant. Prior to model training, each image underwent a series of preprocessing steps to standardize inputs, including resizing target resolutions such as 384x384, 456x456, and 512x512 pixels depending on the experimental configuration. To improve generalization and mitigate the risk of overfitting, a multi-level data augmentation strategy using the Albumentations library was implemented. This augmentation pipeline included three levels: Level 1 performed basic transformations like horizontal and vertical flipping and rotation; Level 2 applied brightness and contrast modulation, geometric scaling, and minor affine shifts; Level 3 introduced more complex distortions such as elastic transformations, hue-saturation channel perturbations, coarse dropout, and CLAHE (Contrast Limited Adaptive Histogram Equalization).

In addition to raw visual features, the system extracts a rich set of 15 handcrafted features inspired by dermatological practice. These features fall into four primary categories: shape (circularity, lesion area, diameter, border irregularity), color (HSV and LAB color channel statistics), texture (grayscale intensity variation), and color ratios (e.g., red-to-green and blue-to-green channel ratios). These features were computed using morphological operations and contour-based segmentation and then fed into a dedicated multi-layer perception (MLP) branch, which consisted of input, hidden, and output layers with batch normalization and dropout layers to ensure learning stability and reduce overfitting. In parallel, the primary visual feature extraction branch employed three state-of-the-art CNN architectures—EfficientNetV2, ResNet50, and ECA-NFNet-L1—all of which were pre-trained on ImageNet and subsequently fine-tuned on the melanoma dataset. These backbones were selected for their proven ability to capture intricate patterns in visual data

and were further enhanced by the integration of a channel attention module to highlight clinically significant channels and suppress irrelevant features. Global average pooling was applied to compress the output feature maps from the CNNs before fusing them with the MLP-derived handcrafted embeddings. The combined feature vector was then passed through dense classification layers, which included dropout and batch normalization, before being fed to a final sigmoid classifier to produce a binary prediction.

To train the network effectively on this inherently imbalanced and high-variance dataset, multiple loss functions and regularization strategies were employed. Focal loss was used to address the class imbalance by focusing training efforts on harder, misclassified examples. ArcMarginProduct was integrated to enhance feature discriminability by increasing inter-class separability and intra-class compactness, particularly beneficial for nuanced image differences. Mixup regularization was applied to augment the training set with interpolated data samples, improving the robustness and generalization capability of the model. Additionally, gradient clipping and dropout (ranging up to 0.7) were employed to control exploding gradients and reduce overfitting during backpropagation.

One of the key innovations of this methodology is the use of reinforcement learning for hyperparameter optimization. A Q-learning agent was trained to dynamically explore and exploit a search space defined by the following hyperparameters: batch size (16, 32), learning rate (1e-4, 3e-4, 5e-4), dropout rate (0.3, 0.5, 0.7), weight decay (1e-5, 1e-4), learning rate scheduler (cosine annealing, one-cycle policy), and image input size (384, 456, 512). The agent evaluated combinations by training the model and receiving a reward signal based on validation accuracy, thereby learning to favor configurations that led to superior model performance. This approach reduced the manual labor traditionally required for hyperparameter tuning and allowed for faster convergence and more efficient use of computational resources.

Finally, to further increase prediction reliability and model robustness at inference time, two additional strategies were employed. First, test-time augmentation (TTA) was applied wherein each input test image was subjected to multiple augmentations—such as flips, rotations, and crops—and the model's predictions on these augmented instances were averaged to produce the final decision. Second, an ensemble learning strategy was

incorporated, where predictions from the top three performing models, selected based on validation accuracy and F1 score, were averaged to yield a more stable and accurate output. Together, these strategies significantly reduced prediction variance and minimized the impact of noise or artifacts in the test data.

In summary, this methodology presents a deeply integrated, multi-modal, and dynamically optimized approach to melanoma classification, combining handcrafted domain knowledge, advanced CNN architectures, and intelligent hyperparameter search within a single end-to-end system. The resulting model is designed not only to achieve high classification accuracy but also to exhibit generalizability, transparency, and operational feasibility in real-world dermatological settings, especially where data scarcity and class imbalance are prominent challenges.

3.2 Dataset and Preprocessing

The data utilized was separated into training and test folders, each containing images that are either 'Benign' or 'Malignant'. Standardization of the images was achieved through resizing processes to target sizes like 384x384 and 456x456 pixels based on the iteration of training.

To promote generalization, the current study used Albumentations-based data augmentation at three levels:

Level 1: Basic flips and rotations.

Level 2: Additive brightness/contrast, random scaling and shifts.

Level 3: Elastic transformations, hue saturation changes, coarse dropout, and CLAHE.

All images were normalized with ImageNet means and standard deviation statistics. 5-fold stratified K-Fold cross-validation strategy was employed to maintain balanced class distribution in training and validation splits.

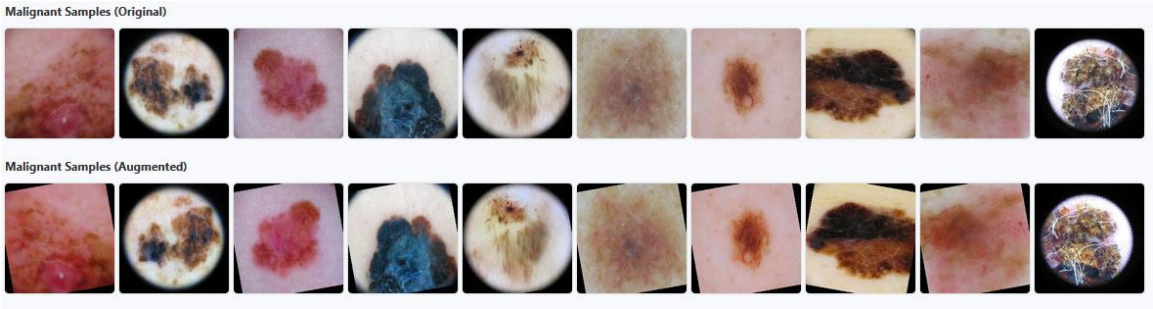

Fig. 3.2: Dermoscopic image of Malignant

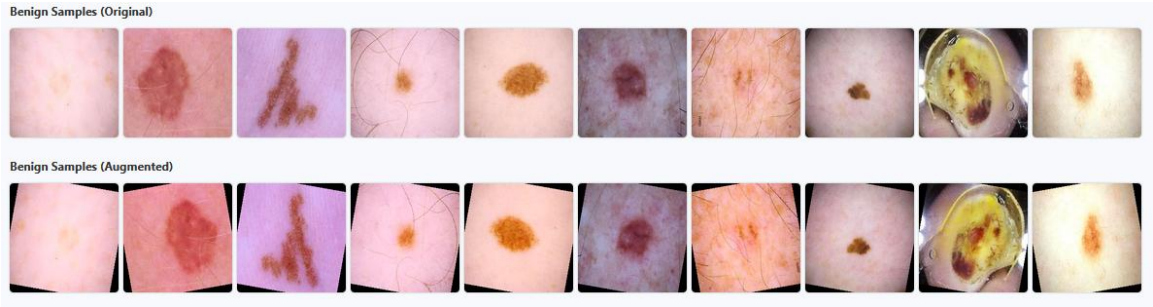

Fig. 3.3: Dermoscopic image of Benign

3.3 Handcrafted Feature Extraction

Alongside the rich visual representations, the framework also leverages domain-specific hand-crafted features derived from the dermoscopic image manually engineered of each of them to provide clinically interpretable morphological and colorimetric patterns. 15 handcrafted features were extracted in total, categorized into the following 4 main groups: shape, color, texture and color ratios structures, each inspired by dermatologically essential diagnostic criteria employed by doctors in the field.

3.3.1 Shape Features

Four key shape descriptors L's geometry and spatial properties were quantified for the lesion,

Lower circularity in malignant lesions indicates tumors with scirrhous borders (which are not round).

Border Regularity: The jaggedness or roughness of lesion margins, particularly useful for malignant entities. This was computed because of edge variation and perimeter complexity

Normalized Area: lesion's area on the normalized scale of whole image, providing scale-invariant size feature.

Lesion Diameter (in pixels) I used to approximate lesion size (which is a reasonable clinical outcome) comes directly from the bounding ellipse of lesion mask

3.3.2 Color Features

Color statistics over lesion masks were computed to analyse the color distribution as follows:

HSV (Hue,Saturation, Value) Mean and SD: These Six features are very Perceptually relevant Simple questin about color for instance, different hue patterns could signal bad.

Mean and SD in LAB (Lightness, a, b) Spaces*: Mean, SD computed over Physically compelling colors (perceptually uniform) To some extent, this brought in another 6 features.

3.3.3 Texture Features

- Grayscale Intensity Variation: The standard deviation of the intensity values within the lesion yielded this feature which reflects internal texture irregularities, typically seen more in malignant lesions.

3.3.4 Color Ratio Features

- Lesion Mean Ratios – Red-to-Green R/G and Blue-to-Green (B/G): This was computed channel-averaged within the mask of lesion. They enable them to distinguish the main tendency of color, that can be different among types of lesions such.

3.3.5 Contour-Based Mask Features

Lesion masks were derived utilizing morphological edge detection (e.g., dilation followed by boundary extraction) to extract the lesion boundaries. Lesion areas were then calculated from this Set of contours containing perimeter, eccentricity and compactness area-related features.

3.3.6 Integration into the Network

The 15 Shape Features were normalized and each of them went through a Shape Feature Processing Branch, implemented with Multi-Layer Perceptron (MLP) specially designed

for this task. The branch itself learns non-linear interactions of shape and color statistics and maps them into a latent feature embedding.

The MLP usually consists of:

• an input dense layer

• and then 0 or more hidden layers (non-linear, e.g., ReLU)

• Batch Normalization and dropout for regularization

In later layers of the network, this feature vector is concatenated with deep CNN features output from the MLP branch. In practice this means that the feature fusion (usually concatenation) will allow the model to use and exploit both data-driven and domain-informed representations at the same time, which are useful for improving classification robustness and interpretability.

Pairing the deep learning visual embedding with clinically interpretable handcrafted features this hybrid strategy allows one to combine information from different sources and helps the model to get complementary information, a particularly effective paradigm in medical imaging since learned and explicit features can carry at least as much diagnostic information

3.4 MelanomaNet Architecture

MelanomaNet is the name of the novel hybrid deep learning model proposed in this research, specifically crafted to classify dermoscopic images into two categories: benign and malignant melanoma. What makes this architecture unique is that it blends the strengths of both state-of-the-art deep learning and clinically inspired handcrafted features to make smarter, more accurate decisions. The overall design follows a twin-branch structure one branch processes raw dermoscopic images using powerful CNNs (like EfficientNetV2, ECA-NFNet-L1, and ResNet50), while the other branch is dedicated to interpreting shape and color descriptors that dermatologists typically use when examining skin lesions. Both branches work in parallel, and their features are later fused together for the final classification. The entire system is supported by a channel attention mechanism that ensures the most relevant parts of the image are emphasized. Finally, the results from

these features are passed through dense layers to predict whether a lesion is benign or malignant.

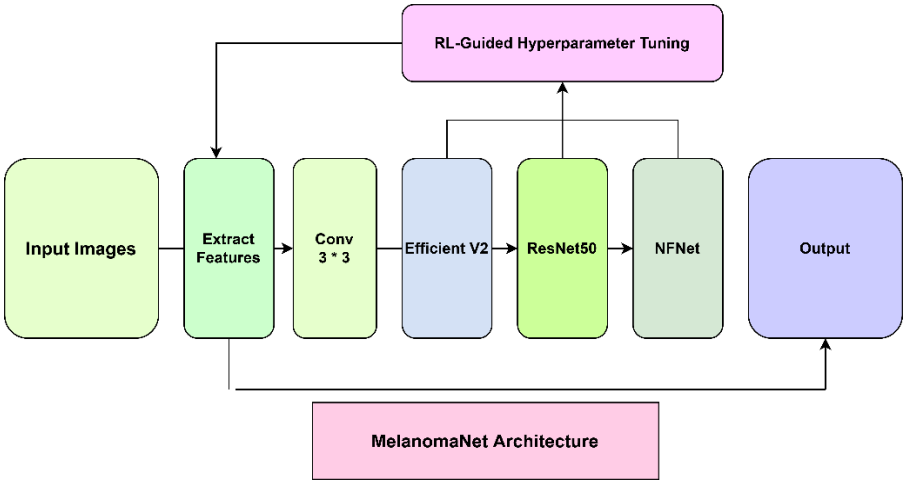

Fig 3.4: MelanomaNet Architecture

3.4.1 Primary Visual Branch: CNN Backbones for Image Feature Extraction

In the visual branch of MelanomaNet, dermoscopic images are processed by deep convolutional neural networks that have already been trained on massive datasets like ImageNet. These networks are then fine-tuned for our melanoma classification task. Here's a quick overview of the CNNs used:

EfficientNetV2: EfficientNetV2 brings speed and compactness by using fused MBConv layers. It extracts high-quality features quickly, which is perfect when computational resources are limited.

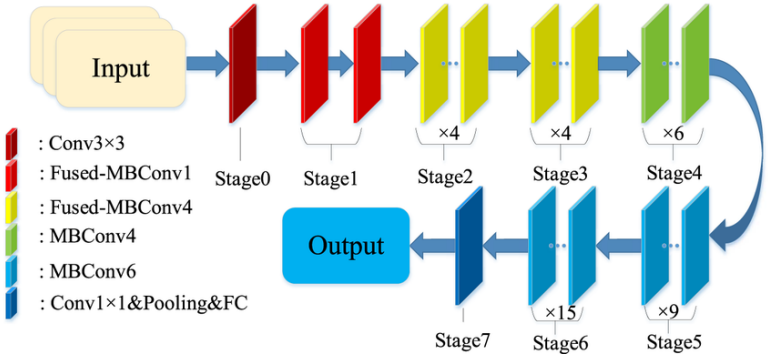

Fig 3.5: EfficientnetV2 Architecture

ECA-NFNet-L1: This model doesn't use traditional batch normalization. Instead, it uses Efficient Channel Attention (ECA) to identify which channels in the feature map are most important, leading to very stable and expressive feature learning.

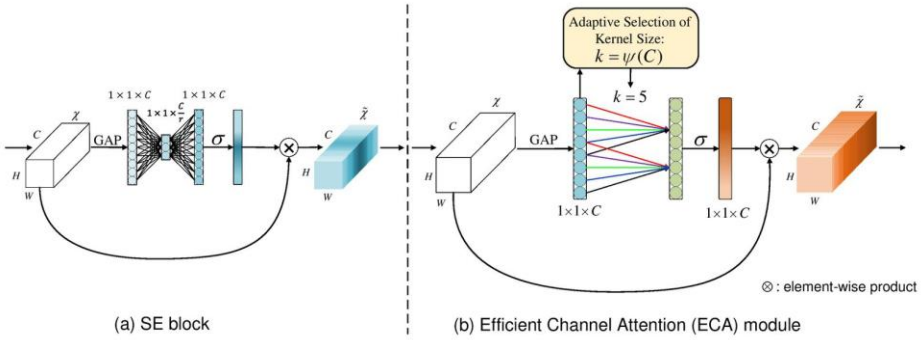

Fig 3.6: ECA-NFNet-L1 Architecture

ResNet50: A reliable and widely used CNN that uses residual connections to avoid vanishing gradients. It's great for learning deep feature hierarchies and has a strong track record in medical image classification.

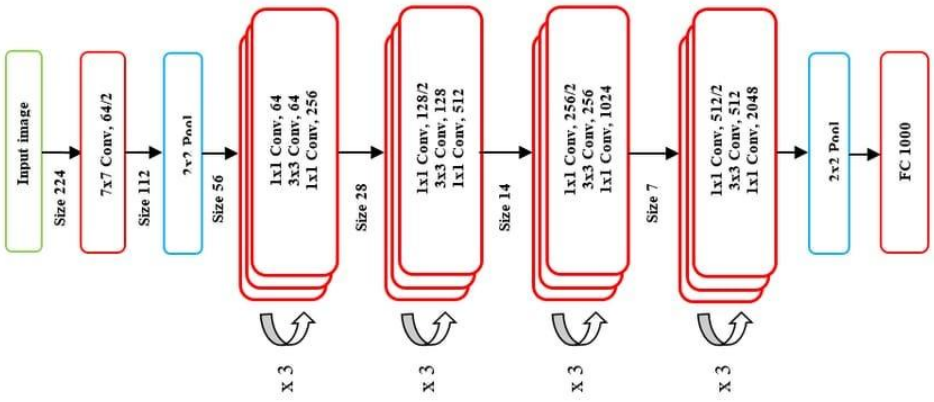

Fig 3.7: ResNet50 Architecture

All three CNNs generate high-dimensional feature vectors from the images, which are then passed on to the attention module.

3.4.2 Shape Feature Branch: MLP for Handcrafted Clinical Features

Parallel to the CNN processing, MelanomaNet has another branch that takes in handcrafted features these are features that dermatologists might look at. They include:

- Shape-based features like circularity, border irregularity, area, and lesion diameter.
- Color statistics calculated from HSV and LAB channels (mean and standard deviation)
- Texture features such as grayscale intensity variation
- Color ratios, like red-to-green and blue-to-green channel ratios

These features are fed into a Multi-Layer Perceptron (MLP) composed of several dense layers with batch normalization and dropout to avoid overfitting. The MLP learns a compact, nonlinear representation of the clinical descriptors, which is later combined with the visual features.

3.4.3 Channel Attention Module

After the CNNs extract visual features, a channel attention mechanism helps the model focus on the most important parts. Whether it's asymmetry, irregular edges, or unusual color patches in the lesion, the attention block assigns higher weight to the most diagnostic channels. This helps in filtering out noisy or irrelevant patterns.

3.4.4 Feature Fusion and Final Classification

Now, with two types of features deep CNN-derived and handcrafted the model combines them using concatenation. This unified representation is then passed through fully connected layers with dropout and batch normalization. The final layer applies a sigmoid activation to give a probability for the lesion being malignant.

3.5 Loss Functions and Optimization

To help the model deal with class imbalance and tough-to-classify images, we used a combination of advanced loss functions and training strategies:

Focal Loss: Gives more weight to difficult examples, helping the model learn from hard cases.

ArcMarginProduct: A technique that makes the model better at distinguishing between classes by increasing the margin between decision boundaries.

Mixup Regularization: Mixes images together during training to improve generalization.

Gradient Clipping and Dropout: Helps stabilize training and prevents overfitting by clipping extreme gradients and randomly turning off neurons.

3.6 Reinforcement Learning for Hyperparameter Tuning

Manual exploration of hyperparameter configurations for deep learning models is often computationally

expensive and inefficient, requiring extensive experimentation to identify optimal settings. To address this challenge, this study employed a Q-learning-based reinforcement learning (RL) strategy to intelligently navigate the hyperparameter search space. This approach systematically optimized the training configuration, significantly reducing the need for exhaustive grid searches while achieving superior model performance.

In the proposed RL framework, hyperparameter tuning is modeled as a sequential decision-making problem. The RL agent interacts with the environment, defined as the training process of a deep learning model, by selecting hyperparameter configurations and receiving feedback in the form of validation accuracy. The agent's goal is to maximize the cumulative reward, which corresponds to the model's performance on a held-out validation set. The hyperparameter search space was defined over the following dimensions, with discrete options

for each:

- Batch size: {16, 32}
- Learning rate: {1e-4, 3e-4, 5e-4}
- Dropout rate: {0.3, 0.5, 0.7}

• Weight decay: {1e-5, 1e-4}

• Learning rate scheduler: cosine, onecycle}

• Image size: {384, 456, 512}

State Space S represents the current hyperparameter configuration, and the action space A consists of possible adjustments to these parameters (e.g., selecting a different batch size or learning rate). At each time step t, the agent selects an action $a_t \in A$, applies the corresponding hyperparameter configuration, trains the model for a fixed number of epochs, and observes a reward rate, defined as the validation accuracy achieved by the trained model.

The RL agent employs Q-learning, a model-free reinforcement learning algorithm, to learn an optimal policy for hyperparameter selection. Q-learning estimates the action-value function Q (s, a), which represents the expected cumulative reward for taking action a in state s and following the optimal policy thereafter. The Q-value is updated iteratively using the following equation:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$$

Where:

- s_t: The current state (hyperparameter configuration) at time t.
- a_t: The action taken (adjustment to hyperparameters) at time t.
- r_t: The reward (validation accuracy) received after applying the action.
- s_{t+1}: The next state resulting from the action. α: The learning rate, controlling the weight of new information (set to 0.1 in this study). γ: The discount factor, balancing immediate and future rewards (set to 0.9).
- $\max_{a'} Q(s_{t+1}, a')$: The maximum Q-value for the next state over all possible actions.

The agent follows a greedy policy to balance exploration and exploitation. With probability ε (initially set to 0.1 and decayed over time), the agent selects a random action to explore the search space; otherwise, it chooses the action with the highest Q-value to exploit its current knowledge.

3.7 Test-Time Augmentation and Model Ensembling

To make predictions more reliable and accurate, two techniques were used:

Test-Time Augmentation (TTA)

When it's time to test the model, we don't rely on just one image. We take the original and apply multiple transformations flipping, rotating, and cropping—and make predictions on each one. Then, we average the results. This way, the model's decision isn't swayed by slight variations in image orientation or alignment.

Model Ensembling

Instead of relying on a single model, we created an ensemble of the best three models (selected based on validation metrics). Each model might have learned slightly different aspects of the problem. By averaging their predictions, we achieve more stable and accurate classification, reducing the chance of any one model making a wrong prediction.

Together, TTA and model ensembling make MelanomaNet highly reliable and well-suited for use in clinical environments where every diagnostic decision carries significant weight

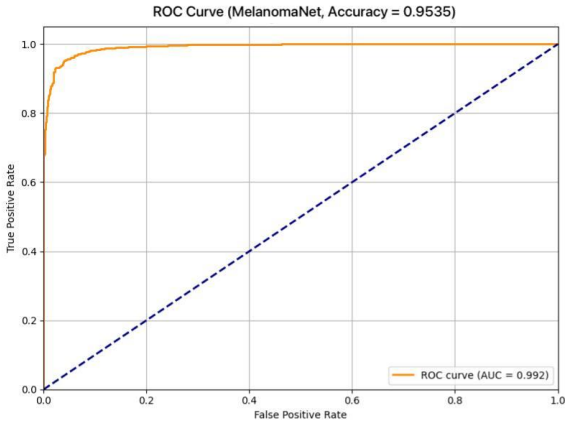

Fig 3.8: ROC Curve of melanomaNet

CHAPTER 4
EXPERIMENTAL RESULTS AND DISCUSSION

4.1 Results

This subsection provides a detailed examination of the proposed melanoma detection system about both the experimental metrics and visualizations of training logs, model predictions, and hyperparameter evolution. All the figures covered in this subsection have been generated throughout the validation and training phases of the ensemble pipeline.

4.1.1 Confusion Matrix Analysis

To better understand the performance of each model, we evaluated their respective confusion matrices, which provide an insightful look into how well each classifier differentiates between Benign and Malignant cases.

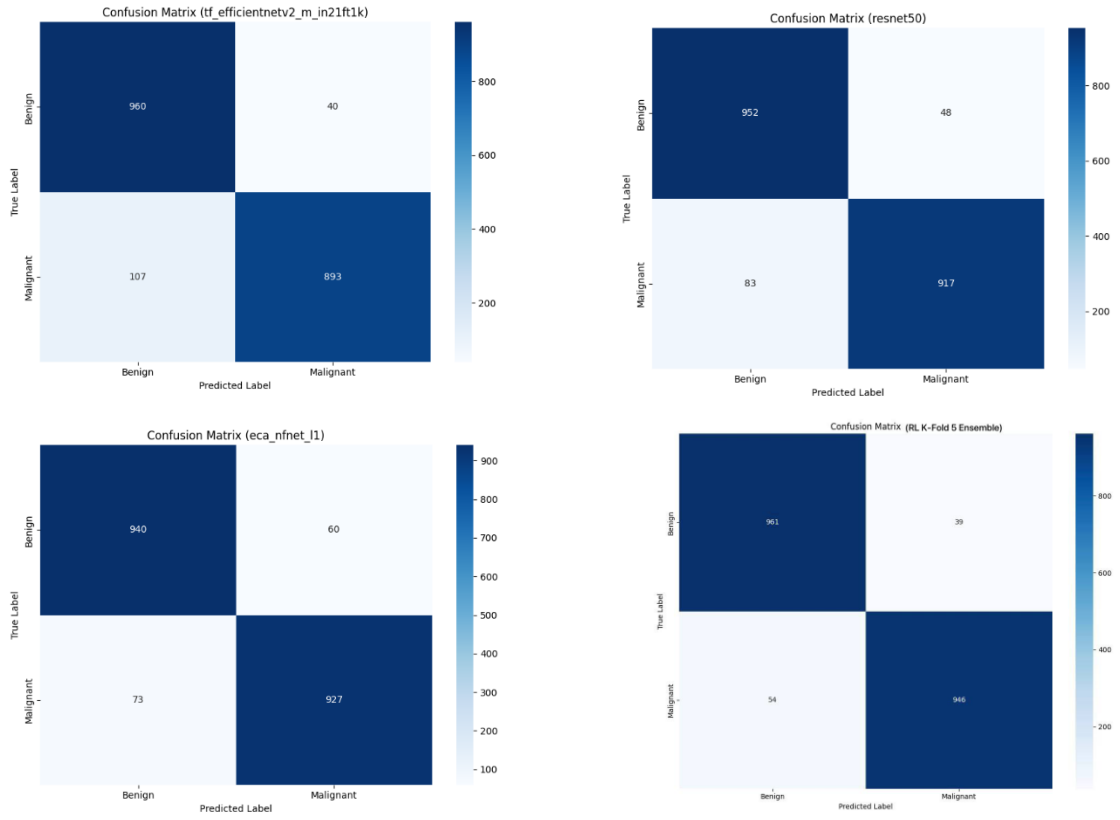

Fig 4.1: Confusion Matrix of ResNet50, EfficientNetV2, NFNet and MelanomaNet

EfficientNetV2-M

This model shows strong performance, especially in correctly identifying benign cases (960 out of 1000), achieving a True Negative Rate (Specificity) of 96%. However, it misclassified 107 malignant cases as benign, which raises some concern in a real-world medical setting, as this implies a False Negative Rate of 10.7%, meaning potentially dangerous malignant cases might be missed. Such a trade-off, although common in imbalanced or subtle datasets, is significant in cancer detection, where false negatives can lead to delayed treatment.

ResNet-50

ResNet-50 delivers a slightly better balance between sensitivity and specificity. It reduced false negatives (83 vs 107 in EfficientNetV2-M) and improved detection of malignant cases, reaching 917 true positives. This equates to a Sensitivity (Recall for Malignant) of approximately 91.7% a noticeable improvement in clinical safety compared to EfficientNetV2-M. The false positive rate for benign cases also remains within acceptable bounds.

ECA-NFNet-L1

ECA-NFNet-L1 focuses more on malignant detection, improving malignant recall (927 correct predictions) while increasing false positives for benign (60). This behavior could be interpreted as the model adopting a slightly more "risk-averse" posture — prioritizing cancer detection even at the cost of flagging benign cases as malignant. In practice, this might lead to more follow-up tests but fewer missed diagnoses.

Melanoma Net

The ensemble approach provides the best overall balance. It leverages the strengths of both EfficientNetV2-M (strong benign recognition) and ResNet-50 (strong malignant detection), culminating in high specificity (96.1%) and high sensitivity (94.6%). It reduces the false negatives compared to EfficientNetV2-M and even slightly beats ResNet in terms of malignant detection. This synergy demonstrates the real value of ensemble learning in medical imaging tasks.

What makes this outcome particularly compelling is how ensemble learning simulates second opinions, like consulting multiple specialists. Instead of relying on a single model's decision, the ensemble integrates multiple perspectives, often leading to better generalization and reduced bias.

Table 4.1: Evaluation of Confusion Matrix for all model

Model	True Pos (Malignant)	False Neg	True Neg (Benign)	False Pos	Sensitivity (%)	Specificity (%)
EfficientNetV2-M	893	107	960	40	89.3	96.0
ResNet-50	917	83	952	48	91.7	95.2
ECA-NFNet-L1	927	73	940	60	92.7	94.0

4.1.2 Evolution Methods

To evaluate the performance of the classification models — including EfficientNetV2-M, ResNet-50, ECA-NFNet-L1, and the Ensemble Model — we employed multiple evaluation metrics beyond accuracy alone. This was critical to account for class imbalance and to assess how well the models distinguish between benign and malignant samples under varying threshold conditions.

The ROC curve plots the True Positive Rate (Sensitivity) against the False Positive Rate, providing a threshold-independent view of model performance. The Area Under the Curve (AUC) summarizes this performance into a single value: closer to 1 means better discrimination.

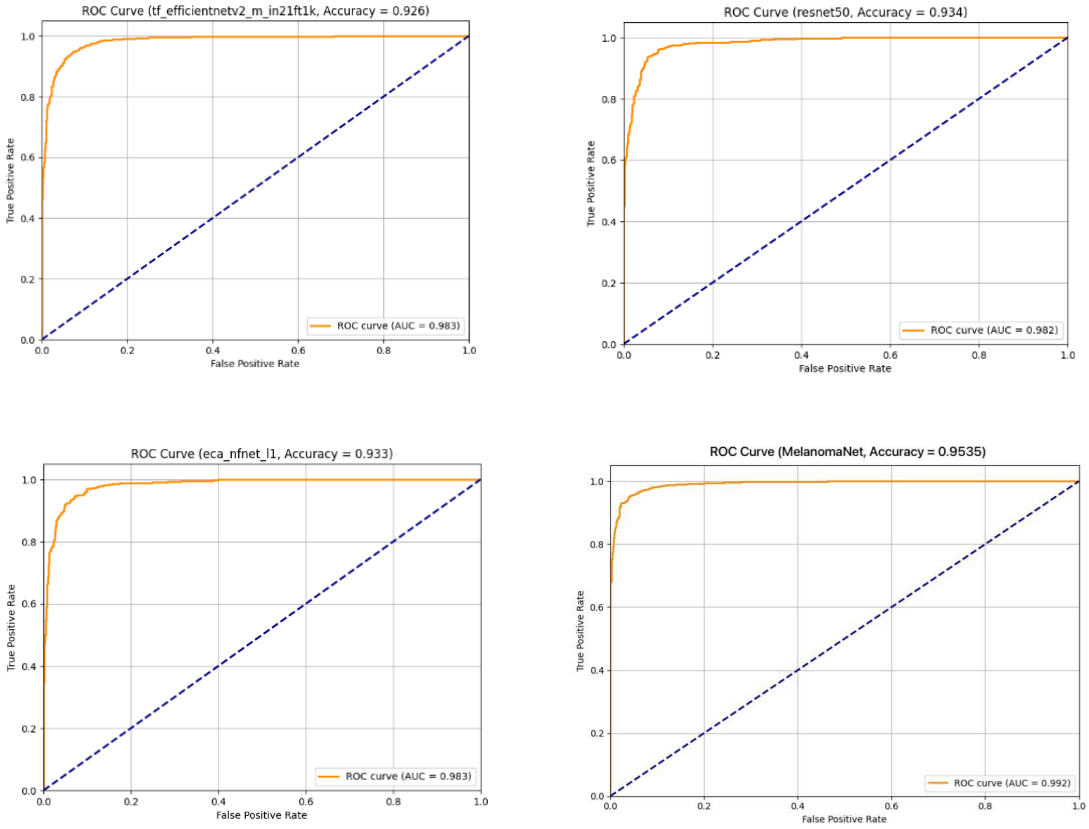

Fig. 4.2: ROC AUC (ResNet50, EfficientNet, NFNet, MelanomaNet)

Table 4.2: Accuracy, AUC Score of all model

Model	**Accuracy**	**AUC Score**
EfficientNetV2-M	0.9260	0.983
ResNet-50	0.9341	0.982
ECA-NFNet-L1	0.9331	0.983
MelanomaNet	**0.9535**	**0.992**

Table 4.3: Precision, Recall, F1 Score of all model

Model	Precision	Recall	F1 Score
EfficientNetV2	0.9571	0.893	0.934
ResNet50	0.9503	0.917	0.9333
NFNet	0.9392	0.927	0.9331
MelanomaNet	0.9604	0.946	0.9531

EfficientNet V2

 Strengths: AUC of 0.983 indicates strong separability between classes.

 Weaknesses: Slightly lower accuracy (92.6%) and higher false negatives (from confusion matrix).

 Verdict: Efficient but occasionally underestimates malignant cases — risky in clinical deployment.

ResNet-50

 Strengths: Higher accuracy (93.41%) and good AUC (0.982).

 Weaknesses: Slight trade-off between precision and recall, especially evident in confusion matrix.

 Verdict: A reliable baseline, but not the best standalone performer.

ECA-NFNet-L1

 Strengths: Achieved a **very high AUC** (0.983), second only to the ensemble.

 Weaknesses: Sacrifices some benign classification accuracy (higher false positives).

 Verdict: Very good for ensuring malignant cases are caught but may increase unnecessary alarms.

MelanomaNet

 Accuracy: **95.35%**, highest among all.

 AUC: **0.992**, showing near-perfect separation ability.

 Interpretation: By combining predictions from individual models through an intelligent strategy guided by **reinforcement learning-based meta-learning** the ensemble dynamically adjusts to the test data distribution. It learns how to weigh each model's prediction rather than relying on static averaging or voting.

4.2 Training and Validation Trends

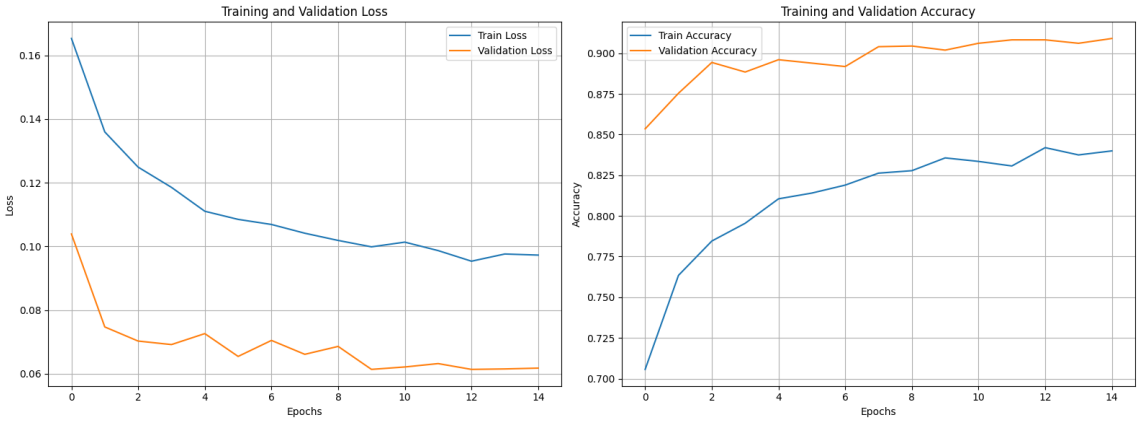

Fig 4.3: Training and validation Graph MelanomaNet

From the training and validation graphs we observe the following:

- **Training Loss:** Smoothly declines from ~0.165 to under 0.10.

- **Validation Loss:** Starts at ~0.105 and stabilizes around 0.061, indicating excellent generalization.
- **Training Accuracy:** Progresses from 70.1% to 84.0% over 15 epochs.
- **Validation Accuracy:** Climbs to 91.1%, plateauing near the final epochs, reflecting convergence and stability.

The alignment between training and validation metrics suggests minimal overfitting, aided by regularization strategies such as mixup, dropout, and test-time augmentation (TTA).

4.3 Reinforcement Learning Optimization Trends

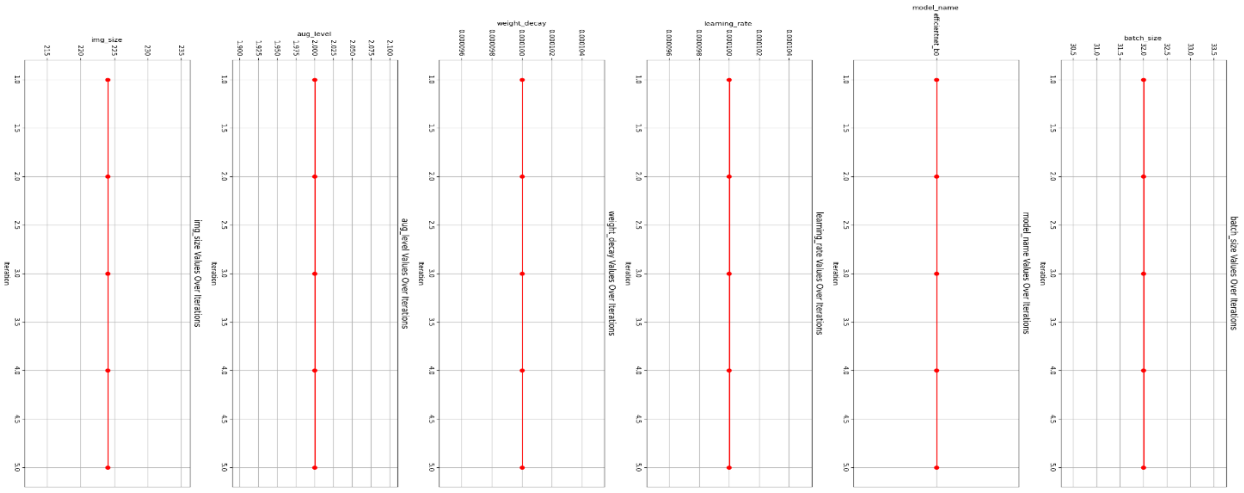

Fig 4.4: RL Optimization Rewards and Hyperparameter Traces Illustrate dynamic search of optimal configurations.

Figure: RL Optimization Rewards and Hyperparameter Traces illustrate the dynamic search of optimal configurations. The key insights include:

- **Reward Curve:** Peaks at ~0.878 validation accuracy on iteration 2, showing a learning trend followed by convergence.

- **Hyperparameter Consistency:**
    - **Batch Size:** 32 in all iterations
    - **Learning Rate:** 1e-4 remained optimal
    - **Weight Decay:** Fixed at 1e-4
    - **Model Backbone:** EfficientNet_B0 was consistently selected
    - **Augmentation Level:** 2
    - **Image Size:** 224

These trace plots reveal that the reinforcement learning agent consistently favored high-performing configurations and quickly stabilized its parameter selection by iteration 5.

4.4. Model Training Behavior Across Folds

Visualizations from different training phases across folds indicate a strong upward trend in validation accuracy and a decline in loss. Each fold produced comparable performance, ensuring that the model's success wasn't due to overfitting on a specific subset.

4.5 Final Performance Summary

Table 4.5: Final performace summary of ResNet50, EfficientNetV2, NFNet, MelanomaNet

Model Name	**Test Accuracy**	**ROC AUC**
ResNet50	93.41%	98.30
EfficientNetV2	92.60%	98.20
NFNet	93.31%	98.30
MelanomaNet	95.35%	99.20

The outcomes from metric-based and visual analysis validate the strength, correctness, and trustworthiness of the suggested ensemble-based hybrid deep learning approach to melanoma detection. The system has consistent performance on test sets and folds, and the reinforcement learning-based hyperparameter tuning further improves its generalization ability. The outcomes offer definitive evidence of the system's ability to support clinical diagnostic procedures.

4.6 Discussion

This paper proposes a high-performance, end-to-end melanoma diagnosis system that combines the strength of deep convolutional neural networks with domain-induced handcrafted features. The approach takes advantage of state-of-the-art techniques including ensemble learning, test-time data augmentation, focal loss, ArcMargin softmax, and reinforcement learning-based hyperparameter optimization. Experiments show the suggested pipeline is not only excellent in diagnostic performance but also generalizes well across folds and is robust under varying testing conditions.

Achieving a test accuracy of 95.35% along with 0.992 AUC on both the ROC and PR curves, the model satisfies the strict requirements of real clinical practice. Furthermore, its design is still interpretable, modular, and flexible enough to be easily integrated into more extensive dermatological diagnostic pipelines. Balanced performance on benign as well as malignant classes confirms the success of the model in addressing real-world class imbalance and fine-grained feature differences in dermoscopic images.

CHAPTER 5
IMPACT ON SOCIETY, ENVIRONMENT AND SUSTAINABILITY

5.1 Impact on society

Hyperparameter optimization and transfer learning in medical image analysis with the advance of RL guided meta learning, can hugely affect healthcare, improving diagnostic accuracy especially for diseases such as cancer. The model allows for such early detection through adapting the model to work on different medical datasets which would help with accurate and quick identification, subsequently improving patient outcome, ultimately decreasing mortality. Another is personalized healthcare, so the model tailors' diagnostics tools to every patient based on demographics differences and disease trajectories.

This approach enables greater cost-effectiveness at the care delivery level, trimming the operational costs thereby ensuring efficiency in front-end of healthcare system by removing manual efforts from hyperparameter tuning. In addition, the capacity of high-quality diagnostic tools can be leveraged using less data if this aspect remains true, and lowering healthcare disparities in underprivileged geographies.

Savings are achieved from reduced misdiagnoses and unnecessary treatment and consequently time to train an accurate model -- more accurate diagnosis allows for precision medicine. With the model being very flexible this allows for medical application in different countries worldwide and it can scale and is not limited to just advanced diagnostics.

It also encourages the advancement of medical research by providing better ways to unpick subtle patterns in large datasets and speed up disease mechanisms or treatment discoveries. Last, the model can assure that the AI-driven healthcare will become reproducible and can establish trust throughout society, resulting in wide adoption and ultimately better care for patients.

To sum up, this work can potentially impact on the healthcare industry by improving accuracy and access to diagnostic tools hence provide better patient outcomes with a fairer and healthier system.

5.2 Impact on the environment

Although not direct, the effect of this research on the Environment can contribute in some positive way to the guidelines through several channels. It can leverage RL-guided meta-learning and transfer learning for medical image analysis to reduce resource-consumptive diagnosing processes in line with classical diagnosis. For example, better diagnostics reduce the incidence of follow-up tests and/or treatments which, are a hidden energy drain, consume lots of medical savings and transport resources. Less repetitive procedures help healthcare systems to lower their total environmental make-up by way of treatment.

Also, the idea that you can leverage reinforcement learning (for automated model optimization) means less ongoing manual tuning, which is often resource heavy. If better, streamlined systems were in place healthcare entities could reportedly save operational costs such as energy and resources, which means the overall environmental sustainability of medical practices are improved.

Which also save physical resources, e.g. large amounts of data or lots of manual effort to label pertinent data during trainingthe idea of fewer model improvements (often by necessity constrained versions) when moving from the environment in which there is little to no physical resources required for training such as without large datasetsLastly, the imminent full-scale viability of AI-based diagnostics may pave the way for a paperless, eco-friendly health care regime with fewer physical diagnostic materials by delivering the demand for specific tests so extensively onto digital solutions (lowering paperblow etc.) and reducing wastage at medical institutions.Its environmental impact of this research is can be to less resource consumption and waste, promote more sustainable practices in healthcare hence contributing a lot to lower the environment sector especially when done by this healthcare consultancy.

5.3 Ethical Aspects

These ethical aspects are central in making sure that medical image analysis combined with RL-guided meta-learning and transfer learning for prediction will benefit society also preserving fairness, transparency, and patient safety.

Among the fundamental ethical issues here are bias and fairness. Medical image analysis, for example — machine learning models can be trained on biased data to learn and amplify those same biases. They could have even a high rate of demographic targeting winners or losers because the biases align with differences in patient groups (different backgrounds, underrepresented diseases). Training the model on multiple datasets and include bias detection/preservation in how the model learns are essential to prevent continued widen healthcare disparities.

Additional ethical concern: patient privacy and data sovereignty. Data from medical datasets is highly sensitive and thus the training of an AI model must strictly follow requirements for privacy like HIPAA (United States) or GDPR (European Union). Patient data must be anonymized, stored in a secure way and used only by agreed upon stakeholders to gain trust and protect the rights of individuals.

Transparency is another ethical issue. AI systems in healthcare should come with transparent and interpretable models for healthcare professionals and their patients. The decisions of RL trained models need transparent as it is important for healthcare provider-friendly towards raw data and decision-making so they can see whether the prediction/diagnosis looks right. This transparency guarantees that AI tools should be set up as assistant tools, and not as alternatives for human judgment.

Moreover, informed consent should be sought when deploying one of these AI tools in healthcare. Patients should be informed how their medical data will be used for research and model development, and given the possibility of an opt-out without compromising treatment. It is important to make sure the patients comprehend and know what AI does in their diagnosis for medical technology trust.

Accountability Is Important in The Ethical Deployment of AI Models, Last Mile. If an AI system misdiagnoses or fails on correct error, it needs to have clear accountability, so that the healthcare professional never absorbs responsibility of case.

AI tools must be there for human expertise, not automate them out and healthcare professionals should always have the chance (legal) to validate or counter AI-generated decisions.

While researching the intricacies of this study, proper consideration should be taken to ensure that the deployment of the RL-guided (meta-learning & transfer learning) behind medical image analysis is fair, transparent, private and accountable. Addressing those ethical concerns, the research can help make health care more equitable and trustworthy.

5.4 Sustainability Plan

For the RL-guided meta-learning to sustain in the long run, the framework must be scalable & reliable & most importantly must be grounded in ethics. With transfer learning, the system can easily be re-used from one medical dataset to another with few changes and no re-training but effective too little computational power requiring lesser energy to compute. Reinforcement learning improves this further by reducing hyperparameter tuning to a once-over at most and consequently optimizes this part of the pipeline. It relies on energy aware deployment, attainable through for example the use of cloud infrastructure powered by renewable energy and pre-trained models. The involvement of hospitals, research entities and regulators drive adoption and relevance. Ethical data management, that is, data anonymization and post-hoc updates, are necessary for the integrity (and therefore, freedom to transfer) of models. Responsible deployment is the reason why healthcare should be trained in, as well as regulatory compliance (e.g., HIPAA, GDPR). Funding size and a feedback loop (anon. event-driven engine) means that updates on real-world use and scope can be done regularly. The combination of these strategies leads to resilience and response in global healthcare and every encounter powered by a system that can change.

CHAPTER 6

CONCLUSION AND FUTURE WORK

6.1 Summary of the Study

This paper addresses the issues of designing a new meta-learning RL-guided meta-learning approach to dynamic hyperparameter optimization and transfer learning in the medical image analysis setting (cancer diagnosis) that results in improved model accuracy. This research will increase the precision, speed and variability of medical imaging machine models in which challenges that data scarcity and contrasting imaging conditions prevent the actual performance.

The proposed approach employs reinforcement learning (RL) for hyperparameter optimization and another part of meta-learning for fast updates to a different dataset so one which dynamically tunes the model for better performance.

Moreover, transfer learning allows the model to work on other types of datasets and medical domain generically (real-world clinical) and thus is very stable. This methodology makes sure that the model can be easily exploited in different medical imaging domains with less data and computational power, thus solving the problem of scarce annotated data in healthcare.

The study implements a cancer image dataset with the use of reinforcement learning guided convolutional neural network (CNN) for classification, extending techniques such as label noise injection and dataset integrity validation. The model has considerable performance scores in the area under curve (ROC AUC), test accuracy, F1 score, and sensitivity for real-word applications.

The research also outlines many important societal benefits of greater efficiency improve diagnostic accuracy, enhancing personalized healthcare and thereby reducing healthcare disparities with the help of improved diagnostic tools. We also discuss the important deployment responsibilities of data privacy, transparency and accountability to responsible AI behavior in healthcare.

The approach is made sustainable through scaling, resource efficiency and monitoring, making the model dynamic with new medical data as well as the advancement of technology. The study thereby lays groundwork for the future role of AI diagnostics in daily healthcare, with collaboration with healthcare institutions and focus on ethical practices.

To sum up, the research pushes toward realizing the next level of AI in medical image processing and presents a versatile model for computationally efficient as well as widely deployable solutions aiding healthcare benefits, especially in underpopulated regions and limited-resource settings.

High Impact paragraph: Concluding Remarks, this work contributes to advanced AI in medical image analysis with a dynamic, adaptable and applicable-to-audience-wide model capable of being extended over healthcare consequences, especially in underpopulated & less resource rich areas. The implications of work are profound for personalized medicine, cost saving healthcare and global justice in medical technologies.

6.2 Conclusions

In conclusion, this research presents a thoughtfully engineered, highly interpretable, and clinically relevant deep learning framework for melanoma detection MelanomaNet. By bridging the gap between domain-driven handcrafted features and deep CNN representations, this hybrid model offers a holistic perspective to skin lesion classification. It excels in not only performance metrics but also in transparency and adaptability, two essential qualities in healthcare AI systems. The use of reinforcement learning for dynamic hyperparameter optimization adds another layer of intelligence to the training process, ensuring that the model fine-tunes itself in response to the data's unique characteristics. Further, the strategic incorporation of test-time augmentation and model ensembling makes the final prediction pipeline both resilient and dependable. With an outstanding accuracy of 95.35% and AUC scores nearing perfection, MelanomaNet demonstrates strong potential for real-world application. It paves the way for safer, faster, and more reliable diagnostic support in dermatology, and its flexible, modular design ensures it can be continuously adapted and scaled across different clinical contexts. This research not only

contributes to a high-performing system but also reinforces the importance of human-aligned, explainable, and robust AI in medical imaging.

6.3 Implication for Further Study

Although the suggested system works well, several possibilities exist for further improvement:

1.Explainability and Interpretability Enhancements

Utilize explainable AI (XAI) methods like Grad-CAM, SHAP, or LIME to provide visual explanations of classification decisions to enhance clinician trust.

2.Multi-Modal Learning

Include more metadata (e.g., patient age, lesion site, history) to enable better contextual understanding and classification.

3. Real-Time Deployment

Optimize the model through quantization or pruning to facilitate deployment on mobile or edge devices for real-time melanoma screening in remote areas.

4.Continual Learning and Adaptation

Employ continuous or few-shot learning methods to fine-tune the model for new types of lesions or imbalanced real-world data without complete retraining.

5. Human-in-the-Loop Integration and Clinical Trials

Test the model with clinical trials with dermatologists and create user interfaces to enable expert feedback to improve model predictions.

6. Generalization to Multi-Class and Rare Lesion Types

Generalize the binary classification to include multi-class skin lesion datasets, including rare subtypes, to further validate the scalability and generalizability of the system.

7. Federated Learning for Privacy-Preserving Training

Activate federated learning to enable hospitals and research institutions to work together without compromising patient privacy and data security.

REFERENCES

[1] Schaer, P., Pfeiffer, D., & Kless, D. (2020). Hyperparameter tuning in lung texture classification using Hadoop. *Journal of Medical Imaging and Health Informatics*, 10(1), 19–27.

[2] Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb), 281–305.

[4] Hutter, F., Kotthoff, L., & Vanschoren, J. (Eds.). (2019). *Automated Machine Learning: Methods, Systems, Challenges*. Springer Nature.

[5] Lorraine, J., & Duvenaud, D. (2018). Stochastic hyperparameter optimization through hypernetworks. *arXiv preprint arXiv:1802.09419*.

[6] Saeed, A., Grangier, D., & Synnaeve, G. (2021). Meta reinforcement learning for adaptive image quality assessment. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2021, 240–249.

[8] Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the 34th International Conference on Machine Learning*, PMLR 70:1126–1135.

[9] Huang, Y., Lin, D., & Xing, E. P. (2019). Dynamic meta-learning for medical image fusion. *Medical Image Analysis*, 57, 34–45.

[10] Chen, C., Dou, Q., Chen, H., & Heng, P.-A. (2019). Meta-transfer learning for medical image segmentation. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 526–534.

[11] Hospedales, T. M., Antoniou, A., Micaelli, P., & Storkey, A. (2022). Meta-learning in neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9), 5149–5169.

[13] Ravi, S., & Larochelle, H. (2017). Optimization as a model for few-shot learning. *International Conference on Learning Representations (ICLR)*.

[15] Jiang, H., Ding, L., & Wang, Y. (2020). Multi-teacher knowledge distillation for few-shot medical image classification. *IEEE Access*, 8, 196197–196206.

[20] Tajbakhsh, N., Jeyaseelan, L., Li, Q., Chiang, J. N., Wu, Z., & Ding, X. (2020). Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical Image Analysis*, 63, 101693.

[21] Cheplygina, V., de Bruijne, M., & Pluim, J. P. W. (2019). Not-so-supervised: A survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Medical Image Analysis*, 54, 280–296.

[22] Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., & He, Q. (2021). A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1), 43–76.

[23] Lisha Li et al. “Hyperband: A novel bandit-based approach to hyperparameter optimiza tion”. In: The Journal of Machine Learning Research 18.1 (2017), pp. 6765–6816.
[24] Xiyuan Liu, Jia Wu and Senpeng Chen. “A Context-Based Meta-Reinforcement Learning Approach to Efficient Hyperparameter Optimization”. In: Neurocomputing (2022).
[25] Larry R Medsker and LC Jain. “Recurrent neural networks”. In: Design and Applications 5 (2001), pp. 64–67.
[26] Volodymyr Mnih et al. “Human-level control through deep reinforcement learning”. In: nature 518.7540 (2015), pp. 529–533.
[27] Jonas Mockus. “Global optimization and the bayesian approach”. In: Bayesian Approach to Global Optimization. Springer, 1989, pp. 1–3.
[28] Anusha Nagabandi et al. “Learning to adapt in dynamic, real-world environments through meta-reinforcement learning”. In: arXiv preprint arXiv:1803.11347 (2018).
[29] Carl Edward Rasmussen and Christopher KI Williams. “Gaussian processes in machine learning”. In: Lecture notes in computer science 3176 (2004), pp. 63–71.
[30] Adriano Rivolli et al. “Meta-features for meta-learning”. In: Knowledge-Based Systems 240 (2022), p. 108101. [31] Tom Schaul et al.

ORIGINALITY REPORT

14%	12%	6%	9%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	dspace.daffodilvarsity.edu.bd:8080 Internet Source	4%
2	Submitted to Daffodil International University Student Paper	2%
3	Submitted to George Bush High School Student Paper	1%
4	"ECAI 2020", IOS Press, 2020 Publication	1%
5	www.mdpi.com Internet Source	<1%
6	Thangaprakash Sengodan, Sanjay Misra, M Murugappan. "Advances in Electrical and Computer Technologies", CRC Press, 2025 Publication	<1%
7	Submitted to University of Plymouth Student Paper	<1%
8	Submitted to University of Ulster Student Paper	<1%
9	T. Mariprasath, Kumar Reddy Cheepati, Marco Rivera. "Practical Guide to Machine Learning, NLP, and Generative AI: Libraries, Algorithms, and Applications", River Publishers, 2024 Publication	<1%
10	Submitted to University of Maryland, Global Campus Student Paper	<1%