Quantitative Attention Validation and Continual Learning for ECG Arrhythmia Detection with Elastic Weight Consolidation

Submitted By:
Md. Wahiduzzaman Shakib
ID: 221-35-825
Department of Software Engineering

Supervised By:
DR. MD. Fazla Elahe
Assistant Professor & Associate Head
Department of Software Engineering

Bachelor of Science

Department of Software Engineering (Major in Data Science)

DAFFODIL INTERNATIONAL UNIVERSITY

APPROVAL

This thesis titled on "**Quantitative Attention Validation and Continual Learning for ECG Arrhythmia Detection with Elastic Weight Consolidation**", submitted by **Md Wahiduzzaman Shakib (ID: 221-35-825)** to the Department of Software Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of Bachelor of Science in Software Engineering and approval as to its style and contents.

BOARD OF EXAMINERS

Fazla Elahe

Chairman

Dr. Fazla Ealhe
Assistant Professor & Associate Head
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 1

Dr. Marzia Ahmed
Assistant Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 2

Dr. Shabnom Mustary
Assistant Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 3

Md. Rajib Mia
Lecturer (Senior Scale)
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

External Examiner

Mohammod Abul Kashem, PhD
Professor
Department of Computer Science and Engineering
DUET, Bangladesh

DAFFODIL INTERNATIONAL UNIVERSITY

DECLARATION OF THESIS AND COPYRIGHT

Author's Full Name : Md.Wahiduzzaman Shakib
Date of Birth : 12-01-2003
Title : Quantitative Attention Validation and Continual Learning for ECG Arrhythmia Detection with Elastic Weight Consolidation
Academic Session : 2022-2025

I declare that this thesis is classified as:

☒ CONFIDENTIAL (Contains confidential information under the Official Secret Act 1997) *
☐ RESTRICTED (Contains restricted information as specified by the organization where research was done) *
☐ OPEN ACCESS I agree that my thesis to be published as online open access (Full Text)

I acknowledge that Daffodil International University reserves the following rights:

1. The Thesis is the Property of Daffodil International University.
2. The Library of Daffodil International University has the right to make copies of the thesis for the purpose of research only.
3. The Library of Daffodil International University has the right to make copies of the thesis for academic exchange.

Certified by:

Md. Wahiduzzaman Shakib
(Student's Signature)

Md. Wahiduzzaman Shakib
Student ID 221-35-825
Date: 26 November, 2025

Fazla Elahe 26.11.25
(Supervisor's Signature)

Dr. Md. Fazla Elahe
Name of Supervisor
Date: 26 November, 2025

NOTE: * If the thesis is CONFIDENTIAL or RESTRICTED, please attach a thesis declaration letter.

THESIS DECLARATION LETTER (OPTIONAL)

Librarian,
Daffodil International University,
Daffodil Smart City,
Ashulia.Dhaka,Bangladesh

Dear Sir,

CLASSIFICATION OF THESIS AS RESTRICTED

Please be informed that the following thesis is classified as RESTRICTED for a period of three (3) years from the date of this letter. The reasons for this classification are as listed below.

Author's Name: Md. Wahiduzzaman Shakib
Thesis Title: Quantitative Attention Validation and Continual Learning for ECG Arrhythmia Detection with Elastic Weight Consolidation

Reasons

i. The thesis contains detailed methodological designs, experimental configurations, and analytical frameworks that may be used as a foundation for future academic publications and ongoing research.

ii. Premature public disclosure of the full thesis may compromise the originality and novelty of the research findings intended for submission to peer-reviewed journals or conferences.

iii. The study utilizes curated datasets, pre-processing strategies, and model evaluation protocols that require controlled access to prevent misuse or misinterpretation outside an academic research context.

Thank you.
Yours faithfully,

Facla Elah
26.11.25

(Supervisor's Signature)

Date: 26 November, 2025

Stamp:

Note: This letter should be written by the supervisor and addressed to the Librarian, *Daffodil International University* with its copy attached to the thesis.

SUPERVISOR’S DECLARATION

I hereby declare that I have checked this thesis and, in my opinion, this thesis is adequate in terms of scope and quality for the award of the degree of Bachelor of Science.

Fazla Elahe
26.11.25

(Supervisor’s Signature)

Full Name : Dr. Md. Fazla Elahe
Position : Assistant Professor and Associate Head
Date : 26 November, 2025

STUDENT'S DECLARATION

I hereby declare that the work in this thesis is based on my original work except for quotations and citations which have been duly acknowledged. I also declare that it has not been previously or concurrently submitted for any other degree at Daffodil International University or any other institution.

Md. Wahiduzzaman Shakib

(Student's Signature)

Full Name : Md. Wahiduzzaman Shakib
ID Number : 221-35-825
Date : 26 November, 2025

Quantitative Attention Validation and Continual Learning for ECG Arrhythmia Detection with Elastic Weight Consolidation

Md. Wahiduzzaman Shakib
ID: 221-35-825

Thesis submitted in fulfillment of the requirements
for the award of the degree of
Bachelor of Science

Department of Software Engineering (Major in Data Science)

DAFFODIL INTERNATIONAL UNIVERSITY

NOVEMBER 2025

ACKNOWLEDGEMENTS

Thanks to Almighty Allah (SWT), the Most Gracious, the Most Merciful, for giving me the strength and opportunity to complete this work.

My sincere thanks go to my respected supervisor Assistant Professor and Associate Head Dr. Md. Fazla Elahe Sir for his endless support, valuable suggestions and inspiring idea that have contributed a lot to the content of this study. I cannot thank you enough for your mentorship both for guiding this thesis and such inspiration.

I am grateful as well to the organizations and individuals who supplied the invaluable data which has made this research possible: PhysioNet for providing libraries with publicly ECG data bases. Chapman University and Shaoxing People's Hospital, for the Shaoxing ECG dataset support. The journeyman developers of MIT-BIH Arrhythmia Database, which became the standard with thousands of studies being built upon it.

I am very thankful to my family for their great love, understanding and support during this whole time as they have been a source of strength. Finally, I would like to express my gratitude to all of friends who kindly supported, encouraged and cooperated with me in order to proceed further this work.

DEDICATION

This thesis is dedicated to my esteemed parents, whose unwavering support and sacrifices have enabled every step of my journey, and to my respected teachers, whose guidance and wisdom have profoundly shaped my academic growth.

ABSTRACT

Deep learning to predict ECG arrhythmia has a reputation of achieving accuracies of over 98% but this is usually due to intra-patient splits which inflate results and do not extrapolate to clinical locations. Our interpretable CNN-BiLSTM-Attention model was used to classify binary arrhythmia based on patient-level splitting and cross-dataset transfer learning between the Chinese Shaoxing database and MIT-BIH. Stage 2b Our model provided 88.1 percent AUC on test data (MIT-BIH). The statistical analysis showed that the arrhythmic and normal rhythms involve quite different attention mechanisms (Cohen $d = 0.93$, $p < 0.001$) as arrhythmic samples involve more broadband temporal integration and normal rhythms rely on narrowband feature detection. The learning experiments of stage 3 curriculum revealed the retention-plasticity paradox of continual learning. Elastic Weight Consolidation had two solutions, EWC-Highest, which was the best with the highest MIT-BIH performance, high-level attention discrimination, and EWC-Balanced with the best complex pattern learning and acceptable retention. This study can be used to show that methodologically sound 88% accuracy with interpretable attention is more clinically useful than standards set by less generalizable evaluation procedures.

Key Words: Arrhythmia detection, ECG signals, Curriculum learning, Continual learning, Attention mechanism, Explainable AI, Quantitative validation, Elastic Weight Consolidation, Retention-Plasticity tradeoff

TABLE OF CONTENT

DECLARATION

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS

Symbol	Description
Σ	Summation operator
exp()	Exponential function
log()	Natural logarithm
±	Plus–minus (mean ± standard deviation)
$\rightarrow$	Mapping / transformation
$\wedge$	Logical AND
$\neg$	Logical NOT
h_t	Hidden state vector at time *t* from BiLSTM
$H = \{h_1, h_2, \ldots, h_{225}\}$	Sequence of BiLSTM hidden vectors
W	Attention transformation weight matrix
b	Attention transformation bias vector
u_t	Intermediate attention vector (after tanh transformation)
a_t	Attention score for timestep *t*
α_t	Normalized attention weight (softmax output)
c	Context vector computed from attention
u	Trainable attention context vector
μ	Mean value
σ	Standard deviation
p	p-value (statistical significance)
t	t-statistic
d	Cohen’s d effect size
H(attention)	Shannon entropy of attention weights
$\max(\alpha)$	Maximum attention weight
$\arg\max(\alpha)$	Index of maximum attention weight
θ^*	Stage-2b optimal parameters

θ	Current parameters during Stage 3 training
F	Fisher Information Matrix
λ	EWC regularization strength
L_e	EWC regularization term
L	Total loss function
lr	Learning rate
∇	Gradient operator
x(t)	ECG signal amplitude at time *t*
f_s	Sampling frequency
QRS	QRS complex
P-wave	Atrial depolarization waveform
T-wave	Ventricular repolarization waveform
Hz	Hertz (unit of frequency)
SD	Standard deviation threshold
J	Youden’s J statistic

LIST OF ABBREVIATIONS

Abbreviation	Full Form
AAMI	Association for the Advancement of Medical Instrumentation
AF	Atrial Fibrillation
AI	Artificial Intelligence
APPLE	Atrial Fibrillation Clinical Risk Score
ARR	Arrhythmia
AUC	Area Under the Curve
AUROC	Area Under the Receiver Operating Characteristic Curve
AV Block	Atrioventricular Block
BBB	Bundle Branch Block
BCE	Binary Cross-Entropy
BiLSTM	Bidirectional Long Short-Term Memory
CBAM	Convolutional Block Attention Module
CGAM	Channel-Guided Attention Mechanism
CI	Confidence Interval
CL	Curriculum Learning
CNN	Convolutional Neural Network
CVD	Cardiovascular Disease
d	Cohen’s d (Effect Size)
DL	Deep Learning
DNN	Deep Neural Network
ECG	Electrocardiogram
EMBC	Engineering in Medicine and Biology Conference
EWC	Elastic Weight Consolidation
F1-score	Harmonic Mean of Precision and Recall
FIM	Fisher Information Matrix
GAN	Generative Adversarial Network

Abbreviation	**Full Form**
HRV	Heart Rate Variability
ICD	Implantable Cardioverter-Defibrillator
ID	Identifier
IEEE	Institute of Electrical and Electronics Engineers
LSTM	Long Short-Term Memory
LR	Learning Rate
MIT-BIH	Massachusetts Institute of Technology – Beth Israel Hospital Arrhythmia Database
ML	Machine Learning
MSE	Mean Squared Error
MS-LTCAF	Multi-Scale Lead–Temporal Co-Attention Framework
NSR	Normal Sinus Rhythm
PVC	Premature Ventricular Contraction
PTB-XL	PhysioNet PTB-XL ECG Database
PRC	Precision–Recall Curve
QRS	QRS Complex (ECG waveform component)
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic
SD	Standard Deviation
SHAP	SHapley Additive exPlanations
S4ECG	Structured State-Space ECG Model
SMOTE	Synthetic Minority Oversampling Technique
SNOMED-CT	Systematized Nomenclature of Medicine – Clinical Terms
ST	ST-Segment (ECG waveform)
TCN	Temporal Convolutional Network
t-test	Student’s t-Test
XAI	Explainable Artificial Intelligence
Z-score	Standardized Normalization Score

CHAPTER 1

INTRODUCTION

1.1 Background

Electrocardiography (ECG) is a widely used non-invasive method for diagnosing cardiac arrhythmias, which are responsible for one of the world's leading causes of cardiovascular disease morbidity and mortality according to the World Health Organization (2021) [1]. Despite the significant progress made in machine learning for interpreting ECGs, it is still very challenging to implement them effectively in real life. Intra-patient data splitting, where signal segments from the same patient are present in both training and testing sets, is a common method for achieving reported accuracies of 95-99%. The use of this evaluation method is convenient, but it disproportionately results in an exaggerated performance due to its ability for models to learn patient-specific morphology rather than generalize to clinical patterns.

Differences between datasets add complexity to ECG classification. The MIT-BIH Arrhythmia Database [2] and the Shaoxing ECG dataset [3] have significant differences in sampling rate, patient constitutions, acquisition set-ups or annotation styles. Biomedical machine learning performance also drops significantly when evaluated on a different platform, which represents more general domain shift issues. Another traditional issue is the interpretative power. Deep learning models are often used as black-boxes, especially when dealing with physiological time-series. P-wave or QRS-complexed ST wave whether the model is processing noise artifacts, however, is important to ensure that the clinician can determine the temporal neighborhood of his/her diagnostic predictions. One of the ways that can be promising in order to work on the attention mechanism is to develop the temporal weight distributions which are capable of enhancing the informative regions. The stability of attention patterns in various datasets, different learning stages and training regime are under research in order to establish which patterns are reliable and which are not. There is need to retain some of the knowledge gained therein to ensure the integration of new diagnostic patterns in clinical systems. Moreover, it is also important to keep in mind these developments. This is especially applicable in case of curriculum or sequential learning systems whereby the model undergoes changes of different degrees of complexity as it advances through the stages. Improper mitigation may lead to disastrous forgetting in the process of incremental fine-tuning, which compromises the former abilities, even as they acquire new skills. The challenges would involve developing an arrangement of diagnosing the ECG arrhythmias that can be applied on various data sets and focus on the methodological rigor, clinically interpretable, and systematic evaluation of the various levels of learning.

1.2 Statement of the Problem

Although significant advances in automated ECG analysis have been achieved, three well-recognized limitations still exist:

- **Lack of generalization to different datasets, and clinical settings.**
  Many techniques that perform well in different high-performing models do not generalize easily to new datasets due to demographic, technical and distributional discrepancies between the training and deployment environments. This undermines real-world reliability.
- **Lack of interpretable prediction mechanisms.**
  Clinicians not only need accurate predictions, but also interpretable explanations. There is little trust and uptake if there is no method to validate whether models focus on the physiologically relevant aspects of the signal.
- **Sequential learning instability.**
  It is a common problem for the medical diagnostic systems to increase their knowledge in order to add new diseases or more complex ECG shapes. Naive sequential learning on new tasks leads to catastrophic forgetting by deteriorating the performance on previous tasks and losing interpretability signals.

To overcome those limitations, a universal framework that accommodates robust cross-dataset transfer and yields stable attention patterns that are also clinically interpretable along with continuously preserved diagnostic performance is needed under sequentially learning from curriculum.

1.3 Motivation

Hence, the focus of this study is on developing ECG classification models that are both robust and transparent for clinical application while also providing an explanation for their decision-making process. This provides a practical way to use large-scale ECG datasets for pre-training purposes via cross-dataset transfer learning, which improves generalization when the available patient-level data from smaller sets like MIT-BIH is limited.

Application of temporal attention of ECG waveforms to the input signal improves the diagnostic power of the deep learning system by showing the effects of these elements on classification decisions. As the multi-stage curriculum learning is implemented, the attention behavior of models may alter and become less discriminative, which leaves the interpretability retention issues. This study has been inspired by an analysis of the performance trade-offs in sequential learning. As more advanced relationships between the ECG shapes are acquired by the model, this preservation of discriminatory yield in the earlier signs of the diagnosis is essentially a sensitive measure that must be well measured and not only described anecdotally. The bigger objective of such a mindset is to have reliable medical AI systems, which offer individualized care and repeatable clinical outcomes.

1.4 Research Objectives

1.4.1 Primary Objective

To develop an interpretable cross-dataset ECG arrhythmia detection framework utilizing temporal attention and evaluate its behavior across curriculum-based sequential learning stages.

1.4.2 Secondary Objectives

- To implement patient-level data splitting to obtain realistic generalization performance.
- Quantify attention mechanism effectiveness through statistical validation (Cohen's d, t-tests, entropy analysis)
- Demonstrate successful knowledge transfer across datasets with different demographic and technical characteristics
- To evaluate cross-dataset transfer from Shaoxing to MIT-BIH during Stage 2 fine-tuning.
- To assess catastrophic forgetting and interpretability degradation during Stage 3 training on complex patterns.
- To compare the performance and interpretability characteristics across curriculum stages using consistent evaluation metrics.
- Compare three continual learning strategies: naive sequential fine-tuning vs. two EWC regularization configurations

1.5 Scope

This thesis focuses on binary arrhythmia classification (arrhythmia vs normal sinus rhythm) using single-lead ECG segments. It emphasizes:

- Patient-level splitting for MIT-BIH
- Cross-dataset transfer learning
- Temporal attention interpretability
- Curriculum-based sequential learning evaluation
- Statistical analysis of attention behavior

The work does not include multi-class rhythm classification, deployment optimization, ensemble modeling, or multi-lead fusion.

1.6 Significance of the Study

This work provides clinically more realistic insights into behavior of ECG models by integrating rigorous evaluation protocols and an in-depth analysis of attention-based interpretability. Exploring how distributions of attention change throughout curriculum stages suggests that sequential learning approaches do have their advantages and disadvantages. Moreover, cross-dataset transfer is also indicative of real-life and gives an alternative understanding to the extent to which deep learning models extrapolate to a different population when trained on one population but tested on a different population. The combination of these elements makes it possible to develop ECG classification systems that focus on the following qualities of reliability, transparency, and practical adaptability needed to make sure that the medical decision-support tool is implemented.

1.7 Thesis Organization

The remaining chapters are organized as follows:

- **Chapter 2** reviews related work in ECG deep learning, attention mechanisms, cross-dataset generalization, and curriculum learning.
- **Chapter 3** describes the datasets, preprocessing pipeline, model architecture, attention mechanism, and evaluation strategies.
- **Chapter 4** presents the experimental results, including stage-wise performance, attention statistics, retention analysis, and discussion.
- **Chapter 5** concludes the thesis with a summary of findings and outlines directions for future work.

CHAPTER 2

LITERATURE REVIEW

2.1 Overview

The high rate of deep learning evolution in the detection of ECG arrhythmia has yielded a rich literature around the precision, interpretability, and extrapolation of results. Earlier studies focused more on making the classification more effective, whereas newer research focuses more on classifying the methods employing rigor in its design, cross-data validation, and the desire to understand the decision places. This shift reflects an evolving consensus that clinically deployable ECG models must demonstrate not only high accuracy but also transparency and stability across heterogeneous populations and acquisition settings.

2.2 Traditional Machine Learning for ECG

Classical machine learning frameworks led to the basis of ECG analysis by investigating statistical features, clinical descriptors, and beat-level signal morphologies. Key cardiovascular epidemiology information according to the World Health Organization (2021) [1]. We argue that good automated systems are needed given the extent of cardiovascular diseases globally. Previous works in deep learning, e.g., LeCun et al.'s extraction-based learning principles, are based on gradient methods. (1997) [4] initiated the model of feature extraction systems that later became ECG dedicated neural architectures.

Prior to the dominance of end-to-end deep learning, systems depended on engineered features such as R–R intervals, HRV metrics and wave delineation. While a few contemporary works, such as Petryshak et al. (2021) [5] partially support this tradition by utilizing inter-heartbeat-interval representations for noise-robust PVC detection. The 95.41% precision and 98.72% recall of their result show clearly that temporal handcrafted feature is still effective when paired with deep networks. Furthermore, classic sequence modeling elements also made their appearance in the classical BiRNN by Schuster and Paliwal (1997) [6], catering into us the necessity of forward–backward temporal structure for ECG signals as well.

These earlier principles serve as a foundation for the hybrid deep learning pipelines described later in this chapter. They also validate the choices made in this thesis for temporal modeling, given that rhythm evolution and waveform dynamics are implicated in arrhythmia diagnosis here.

2.3 CNN-Based ECG Classification

Convolutional neural networks CNN, as the most popular model used in ECG related studies so far, owing to its powerful learning capability and good performance on morphological feature extraction. satisfactory classification of arrhythmia using CNN is shown in a number of studies concerning the Excel data set. Khurshid et al. (2022) [7], a CNN was used for incident AF prediction and two-tailed t-test in three different populations (AUROCs: 0.768–0.778). The findings of this study demonstrate how CNNs are able to capture clinically relevant P-wave and QRS complex features in the ECG, with generalizability to other populations. Likewise, Bi et al. (2024) [8] used a multi-branch temporal CNN with multi-head, attention and their performance was 98.75% on MIT-BIH. Multiple convolutional branches with different receptive fields for learning both fine-grained and high-level features in waveform.

CNNs have also achieved promising results in multimodal learning. Tang et al. (2022) [9] fused pre-processed CNN based ECG, electrogram and clinician data by multimodal method for ICD patient’s arrhythmia detection (for example this work's AUROC=0.859 vs the state-of-the-art APPLE or CHA2DS2-VASc performs badly). These results together suggest that CM of CNNs can be made in both single-lead and multi-leads, with the help of other complementary clinical signals tasks.

Domain adaptation methods are also based on the CNN backbones. Imtiaz et al. (2023) [10] introduced a domain-adaptive deep model which limited the distance of features and resolved cross-database problem effectively. The authors achieved an 11.78% gain of proposed method with label on target domains, and eliminate the need for label on target data to further prove that CNN-generated features are quite robust with adaptation approaches.

Dataset paper e.g., Zheng et al. (2020) [11] and the MIT-BIH´s citation of Moody & Mark (2001) [2]) also provide some explanations on why CNNs are still considered as a reference: these datasets consist in rich morphologies that transcend through a region for optimal features capturing convolutional encoders.

2.4 RNN and Hybrid Deep Learning Models

CNNs are capable to capture spatial features well. But arrythmias often appear as inconsistencies over time. Hybrid models that use convolutional neural networks (CNNs) combined with recurrent neural networks (RNNs) including long-short term memory (LSTM) and bidirectional long short-term memory (BiLSTM) are helpful in tackling this issue by learning temporal attributes of ECG data.

Ma Z et al. (2022) [12] demonstrated that combining shape-focused ResNet with time-based-relationship capturing BiLSTM as well as attention mechanisms and augmenting data via GANs are useful. Their approach had an accuracy of 99.4% on the extended MIT-BIH dataset, indicating that combining spatial and temporal information could be valuable for the classification of ECG signals.

Moreover, Najia & Faouzi (2025) [13] introduced a hybrid CNN–CBAM–BiLSTM model with wavelet pre-processing and SMOTE balance. This model achieved excellent accuracy (99.20%) and great family education related F1 score. The success of CBAM (channel and spatial attention) is a testament to the how attention methods can be well learned across time.

Sethi et al.'s work (2025) [14] proposed a model named ProtoECGNet, which combines multi-branch convolutional neural networks with prototype-based reasoning. This hybrid model is intended to leverage both the advantages of representation learning, which learns features that discriminate classes, as well as being interpretable in a clinical setting. Although it is not explicitly recursive in the way we presented, it reflects hybrid consciousness because it combines global and local temporal features.

Recent advances in large span sequence modeling has significantly broadened the family of hybrid methods. Wang et al (2025) [15] detailed S4ECG, which depends on structured state-space models to incorporate dependencies over 20 40 windows of epochs (each lasting about 10 20 minutes). Their approach led to vast out-of-distribution robustness gains as measured by macro-AUROC which advanced up to 11.6%. This verified the necessity of temporal modeling to develop ECG classification models that can better generalize.

Thus, hybrid architectures stem from the characterization of CNN morphology extraction and attention-based interpretability that is key to our thesis approach.

2.5 XAI and Attention-Based Interpretability

Recently interpretability of AI models (XAI) has been a focus point in the field of ECG analysis and healthcare in general, as clinical guidelines would require developers to explain automatic decision support tools.

Prototype learning approaches, such as ProtoECGNet by Sethi et al. (2025) [14], offer intrinsic interpretability by linking predictions to human-interpretable ECG segments. The professionals found these prototypes straightforward to understand and representative, which promoted a broader approach to case-based reasoning.

Huang et al. (2024) [16] suggested a guided attention method, specifically GSA and CGAM, to add domain knowledge in the models. Their attention maps improved interpretability and at the same time achieved performance gains (e.g., +8.96% F1 for WPW detection). The application of Shapley Values to spatial attention is similar to the quantitative attention analysis presented in this thesis.

Talukder et al. (2025) [17] introduced a transparency-focused explainable arrhythmia detection model with dependability and interpretability approach from the clinical angle. Not just an afterthought, interpretability was a primary design goal that would enable clinicians to verify the reasoning behind their predictions.

Czerwinski et al. (2025) [18] used a genetic programming approach for optimization of deep-learning ensembles with the interpretability retained. Their results show that efficiency and interpretability can go hand in hand with thoughtful model design.

It is these studies that directly motivate the strategy for interpretability outlined in this thesis and using temporal attention maps and quantitative metrics (entropy, peak sharpness and/or temporal focus) as a means to evaluate explanation fidelity at various stages of training.

2.6 Cross-Dataset Generalization and Domain Shift

Cross-dataset adaptation in ECG deep learning is an important issue. Several works in the Excel sheet emphasize the difficulty (and also opportunity) of transferring models between datasets.

According to Merdjanovska and Rashkovska (2021), [19] explored deep models on ARR10000 and PTB-XL and discovered that cross-database performance is usually lower, especially when a dissimilar lead configuration is used. Their findings point to the possibility of intra-dataset overestimation.

Similar results were reported by Khurshid et al. (2022) [7], who monitored a decrease in performance on intervening between different groups of hospitals, which supports the necessity of external validation. Imtiaz et al. (2023) [10] suggested unsupervised domain adaptation, which demonstrated a higher level of accuracy even in cases when labelled data in the target domain were not available. This illustrates that cross-dataset transfer is also a problem that can be solved with the right mechanisms.

Charls et al. [20] bypassed patient variability and dataset shift with a single training approach: domain-invariant representation learning using adversarial multi-task training. Rafi et al. (2023) [21] built on this trend with source-free domain adaptation that allows one to adapt without access to original training data, which is useful in privacy-critical settings.

2.7 Curriculum and Continual Learning

Continual learning is increasingly important in medical AI, where models must evolve as new data, devices, and diagnoses emerge.

Sun et al. [22] introduced curriculum learning (CL) in order to increase personalization in ventricular arrhythmia detection, where stretched learning sequences enhance generalization and alleviate adaptation time.

Together, these studies illustrate that both continuous and curricular learning facilitate ECG models to fine-tune while alleviating catastrophic forgetting a key issue to be addressed in this thesis.

2.8 Research Gaps and Thesis Positioning

Research Gap

We identify some common gaps across all the reviewed works:

2.8.1 Gap 1: Lack of Quantitative Interpretability Evaluation

The majority of XAI studies are somehow based on visual inspection of saliency, or attention maps. Even fewer quantify such changes in interpretability as function of training stage or data set.

2.8.2 Gap 2: Limited Cross-Dataset Interpretability Analysis

Although several works discuss domain shift, few evaluate the quality of explanations across transitions between datasets.

2.8.3 Gap 3: Insufficient Study of Forgetting Effects on Interpretability

Continual learning research focuses on accuracy retention, not attention behaviours. Whether attention "collapses" under forgetting remains unexplored.

2.8.4 Gap 4: Lack of Curriculum-Based ECG Frameworks with Clinical XAI Metrics

Curriculum and continual learning studies rarely integrate interpretability metrics directly into their evaluation pipeline.

Positioning of This Thesis

This thesis directly addresses these gaps by:

- **Combining curriculum learning with cross-dataset transfer** (Shaoxing → MIT-BIH → Complex Shaoxing).
- **Quantitatively evaluating attention behaviour** using entropy, peak sharpness, and temporal focus metrics.
- **Studying interpretability collapse under catastrophic forgetting** (Stage 3 variants).
- **Using rigorous patient-level splitting**, consistent with the methodological concerns raised by multiple reviewed studies.

The result is an ECG classification framework that not only seeks strong accuracy, but also prioritizes stability, generalizability, and interpretable clinical reasoning.

CHAPTER 3

METHODOLOGY

3.1 Research Design Overview

The task is grounded on a three-phase curriculum learning model wherein it entails pre-training a large-scale Chinese Shaoxing database, cross-dataset transfer to MIT-BIH Arrhythmia Database and learning complex patterns through continual learning processes. The approach highlights methodological purity via patient level data splitting and quantitative validation of attention mechanisms.

The architecture combines CNNs for spatial features, bidirectional LSTMs for temporal modeling, and attention model to provide interpretation. A final stage consisting in an EWC is applied to mitigate the catastrophic forgetting in sequence learning of complex ECG patterns.

3.2 Dataset Acquisition and Processing

3.2.1 MIT-BIH Arrhythmia Database

The MIT-BIH Arrhythmia Database serves as the primary evaluation dataset due to its extensive validation in literature and availability of patient-level identifiers for proper data splitting. The database contains 48 half-hour two-lead ECG recordings from 47 subjects (25 male, 22 female) recorded between 1975-1979 at Beth Israel Hospital [27]. All recordings were sampled at 360 Hz with manual beat-by-beat annotations by cardiologists.

For this research, only Lead II was extracted from the two-lead recordings to maintain consistency with the Shaoxing database. Signals were segmented into 10-second windows (3,600 samples) with 50% overlap, resulting in 11,700 total segments. Binary labels were created using AAMI EC57 standards: ARR (ventricular ectopic, supraventricular ectopic, fusion beats) versus NSR (normal beats)

3.2.2 Chapman-Shaoxing ECG Database.

Training data set The Chapman-Shaoxing database offers the training and validation data set of 10,646 patients in short-term ECG record (each patient includes at most 12-leads ECG data) between 2015-2018 from China's Shaoxing People's Hospital [3]. All measurements were acquired at a frequency of 500 Hz in a 12-lead standard. Lead II was extracted and signals were resampled to 360 Hz for consistency with MIT-BIH using scipy.signal.resample.

The database was filtered by SNOMED-CT diagnostic codes to create curriculum stages:

- **Stage 1** (Primary Rhythms): 13,623 samples including atrial fibrillation, ventricular premature beats, atrial premature beats, and supraventricular tachycardia
- **Stage 3** (Complex Patterns): 21,702 samples including ST changes, bundle branch blocks, morphology abnormalities, and voltage/axis deviations

3.2.3 Data Pre-processing Pipeline

All the ECGs were first preprocessed in a consistent manner to avoid any bias across the datasets. The pipeline comprised bandpass filtration (0.5-40 Hz, 4th order Butterworth) applied to suppress baseline wander and high-frequency noise, as well as Z-score normalization for each segment. QC discarded segments with >30% flat-line or extreme amplitudes + 5 SD from the mean.

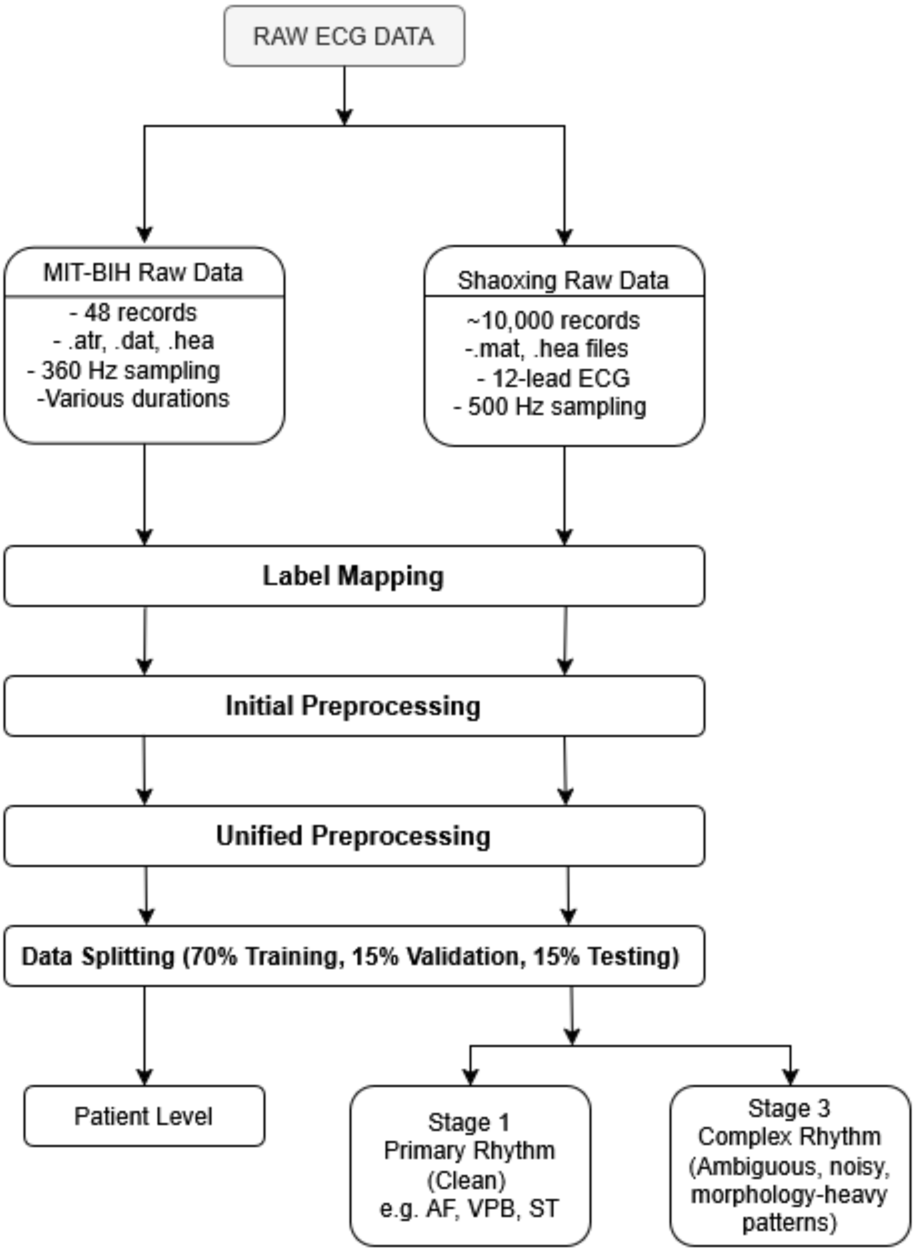

Figure 3.1: Data Preprocessing Pipeline

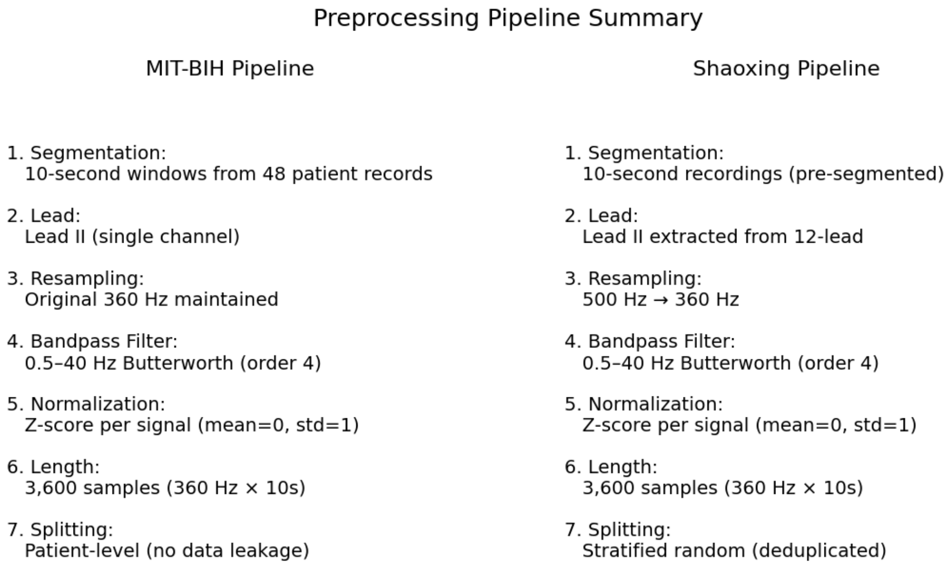

Preprocessing Pipeline Summary

MIT-BIH Pipeline	Shaoxing Pipeline
1. Segmentation: 10-second windows from 48 patient records	1. Segmentation: 10-second recordings (pre-segmented)
2. Lead: Lead II (single channel)	2. Lead: Lead II extracted from 12-lead
3. Resampling: Original 360 Hz maintained	3. Resampling: 500 Hz → 360 Hz
4. Bandpass Filter: 0.5–40 Hz Butterworth (order 4)	4. Bandpass Filter: 0.5–40 Hz Butterworth (order 4)
5. Normalization: Z-score per signal (mean=0, std=1)	5. Normalization: Z-score per signal (mean=0, std=1)
6. Length: 3,600 samples (360 Hz × 10s)	6. Length: 3,600 samples (360 Hz × 10s)
7. Splitting: Patient-level (no data leakage)	7. Splitting: Stratified random (deduplicated)

Figure 3.2: Unified Pre-processing Pipeline Diagram

3.2.4 Data Splitting Protocol

3.2.4.1 MIT-BIH Patient Level Split

Patient-level splitting was implemented for MIT-BIH to prevent data leakage. Training patients (70%) included IDs [100, 101, 103, ..., 222, 223, 228], validation patients (15%) included [105, 108, 111, ..., 221, 230], and test patients (15%) included [102, 104, 107, ..., 231, 232, 234]. Zero patient overlap between splits ensured the model never encountered test patients during training.

3.2.4.2 Shaoxing 2 stage curriculum Split

I. *Disjoint Stage Definition (SNOMED Codes)*

The core principle is defining two mutually exclusive sets of SNOMED codes for classification difficulty:

- **Stage 1: Primary Rhythm Disturbances (Fundamental Arrhythmias)**
  - **Focus:** Clear, fundamental arrhythmias.

  - **Examples:** Atrial Fibrillation (164890007), Atrial Flutter (164889007), Ventricular Premature Beats (17338001), Supraventricular Tachycardia (426761007).
  - **Goal:** Learn fundamental arrhythmia patterns. (13 code categories defined)
- **Stage 3: Complex Cardiac Patterns**
  - **Focus:** Ambiguous rhythms, morphology changes, structural/conduction abnormalities, and voltage/axis deviations.
  - **Examples:** Sinus Bradycardia (426177001), 1st Degree AV Block (270492004), T Wave Changes (164934002), Left Bundle Branch Block (164909002), Left Axis Deviation (89792004).
  - **Goal:** Robustness to subtle/complex patterns. (28 code categories defined)

A critical initial check confirms **zero code overlap** between the two sets.

II. Signal Deduplication

Before splitting, the entire dataset is checked for identical signals to ensure data integrity and prevent information leakage.

- **Method: Hash-based deduplication** is performed on the raw signal data (`X_all`).
- **Result:** A total of **48,935 duplicates** were removed, resulting in a dataset of unique signals.

III. Disjoint Sample Filtering (The Curriculum Separation)

Samples are segregated into Stage 1 and Stage 3 datasets using different rules for Arrhythmia (ARR) and Normal Sinus Rhythm (NSR) samples:

1. **A. ARR Samples (Label = 1)**

- **Disjoint Principle:** An ARR sample is included in a stage **only if** it contains codes from that stage's set and **none** of the codes from the other stage's set.
    - **Stage 1 ARR:** Contains Stage 1 codes ∧ ¬ Stage 3 codes.
    - **Stage 3 ARR:** Contains Stage 3 codes ∧ ¬ Stage 1 codes.
- **Exclusions:** Samples with codes from *both* Stage 1 and Stage 3 (**dual diagnosis**) and samples with *neither* set of codes are excluded from both stage datasets.

2. **B. NSR Samples (Label = 0)**

- NSR samples are used to provide a normal baseline for both stages.
- **Split:** The total unique NSR pool is split into two parts:
    - **40%** of NSR samples are assigned to **Stage 1**.
    - **60%** of NSR samples are assigned to **Stage 3**.

3.2.5 Stratified Train/Validation/Test Split

After combining the ARR and NSR subsets for each stage, the final datasets are split to maintain the label balance (Arrhythmia vs. NSR) across all subsets:

- **Split Ratio: 70% Train / 15% Validation / 15% Test**
- **Method:** The `sklearn.model_selection.train_test_split` function is used twice with the `stratify=y` argument to ensure the proportion of ARR and NSR is consistent.

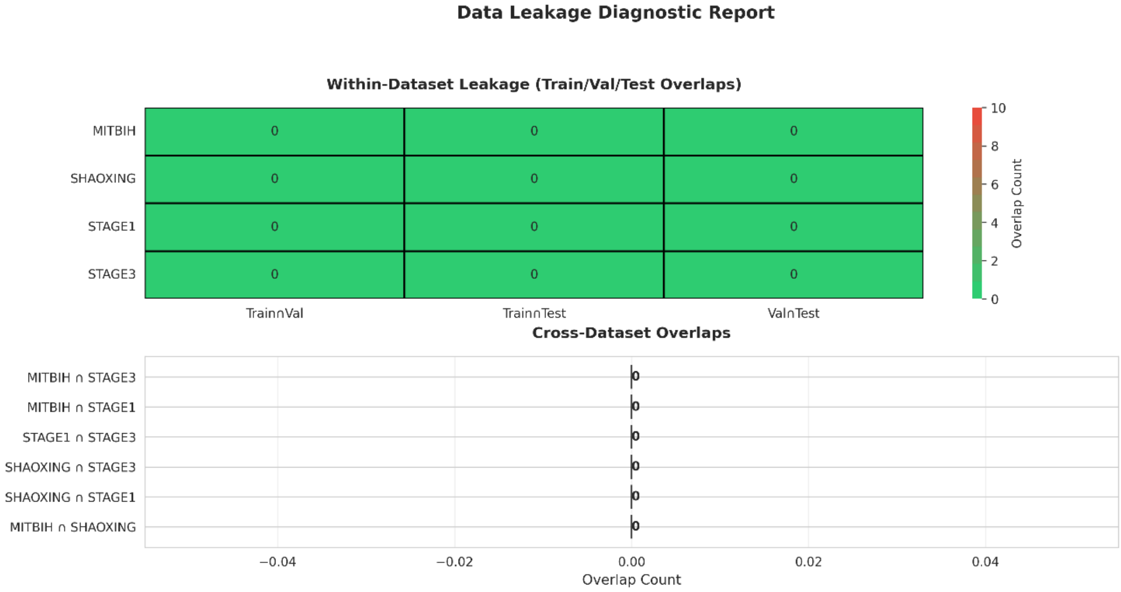

Figure 3.3: Zero Data Leakage Verification

This is an estimate of the structural integrity of the datasets where there may be overlaps of samples. The highest heatmap represents Within-Dataset Leakage, which proves that the training, validation, and testing subsets of each individual dataset are mutually discontinuous, and zero records overlap. The bottom horizontal bar chart, Cross-Dataset Overlaps, confirms that there are no common segments of the ECG between multiple data sources (e.g. MIT-BIH vs. Shaoxing). Such outcomes will make sure that the performance measures of the model will not be artificially boosted due to the existence of the similar information on splits.

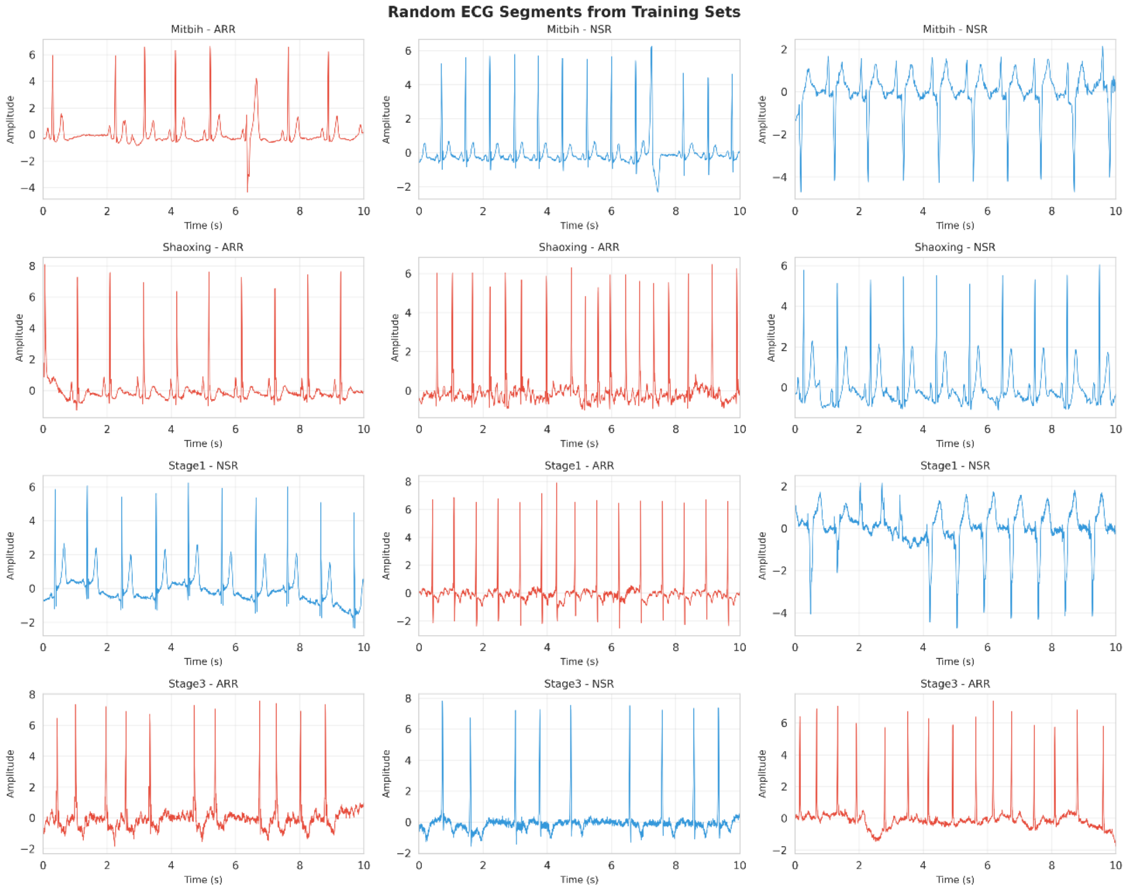

Figure 3.4: Random ECG Signals

This figure is a multi-panel of the statistical and physical properties of the processed signals. Normalization Check and Amplitude Distribution plots indicate that, all datasets have been normalized successfully to a mean of about 0 and standard deviation of 1, which means that there is uniform scaling of features as the input into the neural network. The Signal Length bar chart establishes that each segment of all the four datasets has the identical number of 3, 600 samples, or a 10-second signal at a 360 Hz sampling rate. The table of Signal Specifications that is accompanying gives an overview of total sample counts and the largest and the smallest sample count belong to Shaoxing and Stage 1 respectively.

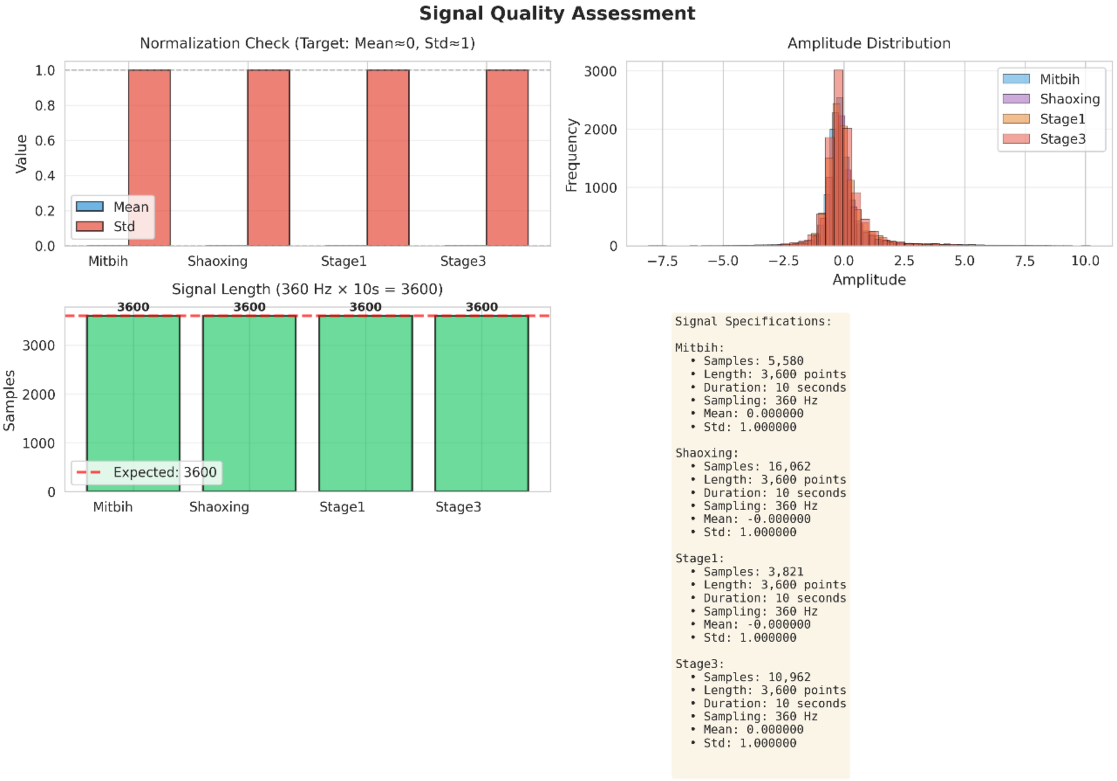

```
Signal Specifications:

Mitbih:
  • Samples: 5,580
  • Length: 3,600 points
  • Duration: 10 seconds
  • Sampling: 360 Hz
  • Mean: 0.000000
  • Std: 1.000000

Shaoxing:
  • Samples: 16,062
  • Length: 3,600 points
  • Duration: 10 seconds
  • Sampling: 360 Hz
  • Mean: -0.000000
  • Std: 1.000000

Stage1:
  • Samples: 3,821
  • Length: 3,600 points
  • Duration: 10 seconds
  • Sampling: 360 Hz
  • Mean: -0.000000
  • Std: 1.000000

Stage3:
  • Samples: 10,962
  • Length: 3,600 points
  • Duration: 10 seconds
  • Sampling: 360 Hz
  • Mean: 0.000000
  • Std: 1.000000
```

Figure 3.5: Signal Quality

3.3 Model Architecture

3.3.1 CNN–BiLSTM–Attention Framework

The proposed architecture integrates multi-scale convolutional feature extraction with sequence modeling and an explicit attention mechanism. The network processes 10-second single-lead ECG windows (3,600 samples) and produces a probability score for binary arrhythmia detection. The architecture consists of four convolutional blocks, a bidirectional LSTM stack, a temporal attention layer, and a compact classification head.

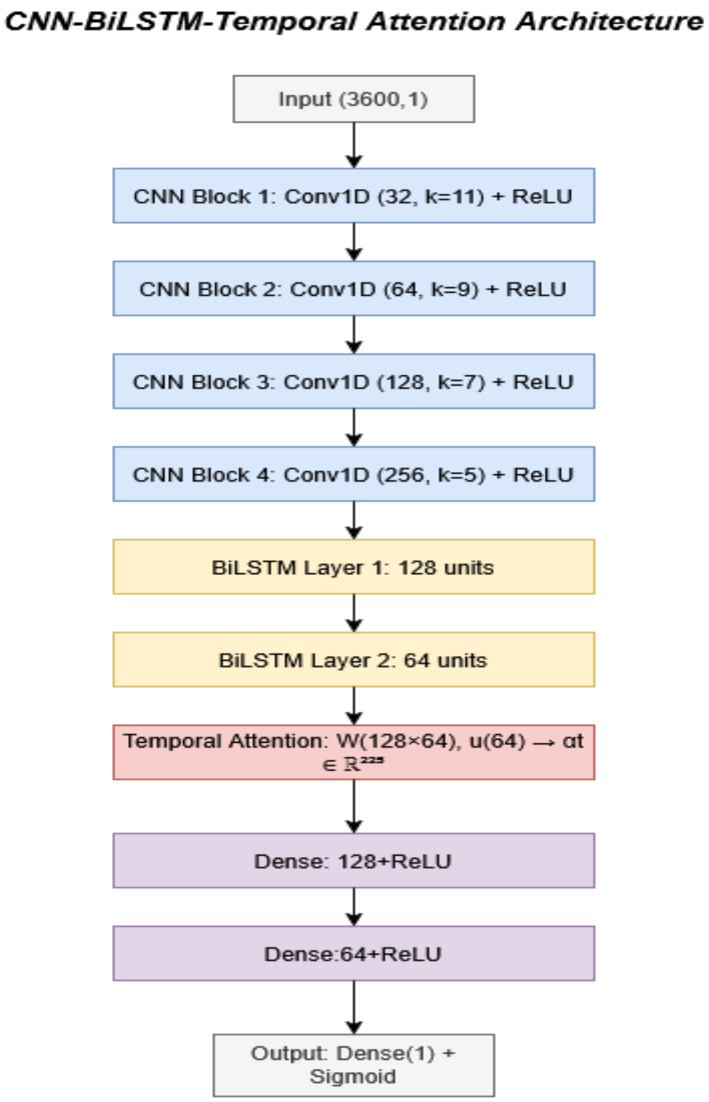

Figure 3.6: CNN–BiLSTM–Attention Architecture

3.3.2 Convolutional Feature Extraction

Hierarchical spatial features are extracted through four convolutional blocks:

- **Block 1:** Conv1D (32, kernel=11) → BatchNorm → ReLU → MaxPool(2) → Dropout(0.3)
- **Block 2:** Conv1D (64, kernel=9) → BatchNorm → ReLU → MaxPool(2) → Dropout(0.3)
- **Block 3:** Conv1D (128, kernel=7) → BatchNorm → ReLU → MaxPool(2) → Dropout(0.3)

- **Block 4:** Conv1D (256, kernel=5) → BatchNorm → ReLU → MaxPool(2) → Dropout(0.3)

The progressively decreasing kernel sizes (11→5) capture waveform structures spanning approximately 30–14 ms, enabling robust detection of P-waves, QRS complexes, and T-wave morphologies. After four pooling operations, the temporal resolution reduces from 3,600 to 225 timesteps, while feature depth expands to 256 channels.

3.3.3 Temporal Modeling with BiLSTM

Temporal dependencies across the entire ECG window are modeled using two stacked bidirectional LSTM layers:

- **Layer 1:** BiLSTM(128 units), dropout=0.3, recurrent_dropout=0.2
- **Layer 2:** BiLSTM(64 units), dropout=0.3, recurrent_dropout=0.2

These layers encode both preceding and succeeding temporal context, which is essential for capturing rhythm irregularities and long-range arrhythmic patterns. Output sequences maintain 225 timesteps with 128-dimensional hidden representations, feeding directly into the attention module.

3.3.4 Temporal Attention Mechanism

A temporal attention layer computes per-timestep importance weights, enabling interpretability:

Mathematical formulation:

Let H = $[h_1, h_2, \ldots, h_{255}] \in R^{(225\times128)}$ be the BiLSTM output ($h_t \in R^{128}$)
Let $W \in R^{128\times64}, b \in R^{64}$ *and* $u \in R^{64}$ be learnable parameters.

Step 1 – Transform:

$$\mathrm{u}_t = \tanh(\mathrm{W} \times h_t + \mathrm{b}), \qquad \mathrm{u}_t \in R^{64}$$

Step 2 - Compute attention scores:

$$\mathrm{a}_t = u^T \times \mathrm{u}_t, \qquad u \in R^{64\text{ (learnable context vector)}}$$

Step 3 - Normalize to weights:

$$\alpha_t = \frac{\exp(\mathrm{a}_t)}{\sum_{i=1}^{255} \exp(\mathrm{a}_t)}, \qquad \sum_t \alpha_t = 1$$

Step 4 - Compute context vector:

$$c = \sum_{t=1}^{255} \alpha_t h_t, \qquad c \in R^{128}$$

Output: $c \in R^{128}$→ Classification head
Attention weights: $a \in R^{255}$→ For visualization & analysis
This yields the attention-weighted context vector ccc used by the classifier.

Interpretability advantage: Attention weights α reveal which temporal positions the model considers important. High attention at position t indicates the model is focusing on that region of the ECG signal when making predictions.

3.3.5 Classification Head

The final classification layers are:

- Dense (128) → ReLU → Dropout (0.4) → L2(0.001)
- Dense (64) → ReLU → Dropout (0.4) → L2(0.001)
- Dense (1) → Sigmoid

3.4 Three-Stage Curriculum Learning

3.4.1 Stage 1: Primary Rhythm Pre-training

Stage 1 trains on Shaoxing primary rhythms to establish strong morphological feature representations. Training uses:

- Adam optimizer (learning rate = 5e-4 as implemented)
- Batch size = 64
- Epoch =50
- Class weights {0: 1.0, 1: 2.0}
- Early stopping (patience = 15) and ReduceLROnPlateau

This stage focuses on robust extraction of core arrhythmia morphologies.

3.4.2 Stage 2: Cross-Dataset Transfer

3.4.2.1 Phase 2a Frozen CNN adaptation

The Stage-1 model is loaded, and the entire convolutional frontend is frozen. Only the BiLSTM, attention, and dense layers remain trainable.

- Learning rate = 1e-4
- Batch size = 64
- Epochs = 50
- CNN layers frozen to preserve morphological filters

3.4.2.2 Phase 2b Full fine-tuning

All layers are unfrozen, and the learning rate is significantly reduced:

- Learning rate = 5e-6
- Gradient clipping = 1.0
- Epochs = 50
- All layers trainable

This phase gently adapts morphology and rhythm representations to MIT-BIH without destabilizing prior knowledge.

3.4.3 Stage 3: Complex Pattern Learning with Three Variant Strategies

Stage 3 introduces Shaoxing complex patterns and evaluates the model's capacity to learn new patterns while retaining MIT-BIH performance. The training uses **mixed batches** from Stage 1 and Stage 3, with ratios evolving as:

- **50% → 65% → 75% Stage-3** samples with MIT_BIH samples across phases

Training in Stage 3 is executed under **three distinct variants**.

3.4.3.1 Variant 1: Without EWC (Baseline Model)

This variant follows standard training with **binary cross-entropy only**, without any regularization for retention.

- Optimizer: Adam (lr = 2e−6)
- Batch size: 64
- Epochs: 50
- All layers trainable after initial warm-up
- Serves as baseline for catastrophic forgetting

This model typically demonstrates significant performance drop on MIT-BIH due to unregularized drift of parameters.

3.4.3.2 Variant 2: Highest-EWC (Maximum MIT-BIH Retention)

The highest-EWC model applies strong Fisher-based regularization:

- Three-phase EWC training schedule:
  - **Phase 1:** lr = 1e−6, ratio_s3 = 0.50, epoch = 8, ewc_lamda = 3000
  - **Phase 2:** lr = 3e−6, ratio_s3 = 0.65, epoch = 8, ewc_lamda = 5000
  - **Phase 3:** lr = 1e−6, ratio_s3 = 0.75, epoch = 6, ewc_lamda = 7000
- Fisher diagonal computed using 1000 MIT-BIH samples
- Single optimizer reused across phases with LR updates
- EWC penalty applied per batch

This variant prioritizes retention and yields the highest MIT-BIH retention percentage.

3.4.3.3 Variant 3: Balanced-EWC (Retention–Plasticity Trade-off)

Balanced-EWC employs moderate regularization:

- Using gridsearch to find optimal lamda : lambdas = [500, 800, 1000, 1500, 2000, 3000]
- Optimal λ = **2000**

- Same three-phase training structure as Highest-EWC: Three-phase EWC training schedule:
  - **Phase 1:** lr = 2e−6, ratio_s3 = 0.50, epoch = 6
  - **Phase 2:** lr = 5e−6, ratio_s3 = 0.60, epoch = 12
  - **Phase 3:** lr = 2e−6, ratio_s3 = 0.655, epoch = 8

- Ratios and learning rates identical across phases
- Provides improved complex-pattern learning while reducing MIT-BIH forgetting compared to the no-EWC baseline

This model achieves better dual-domain equilibrium.

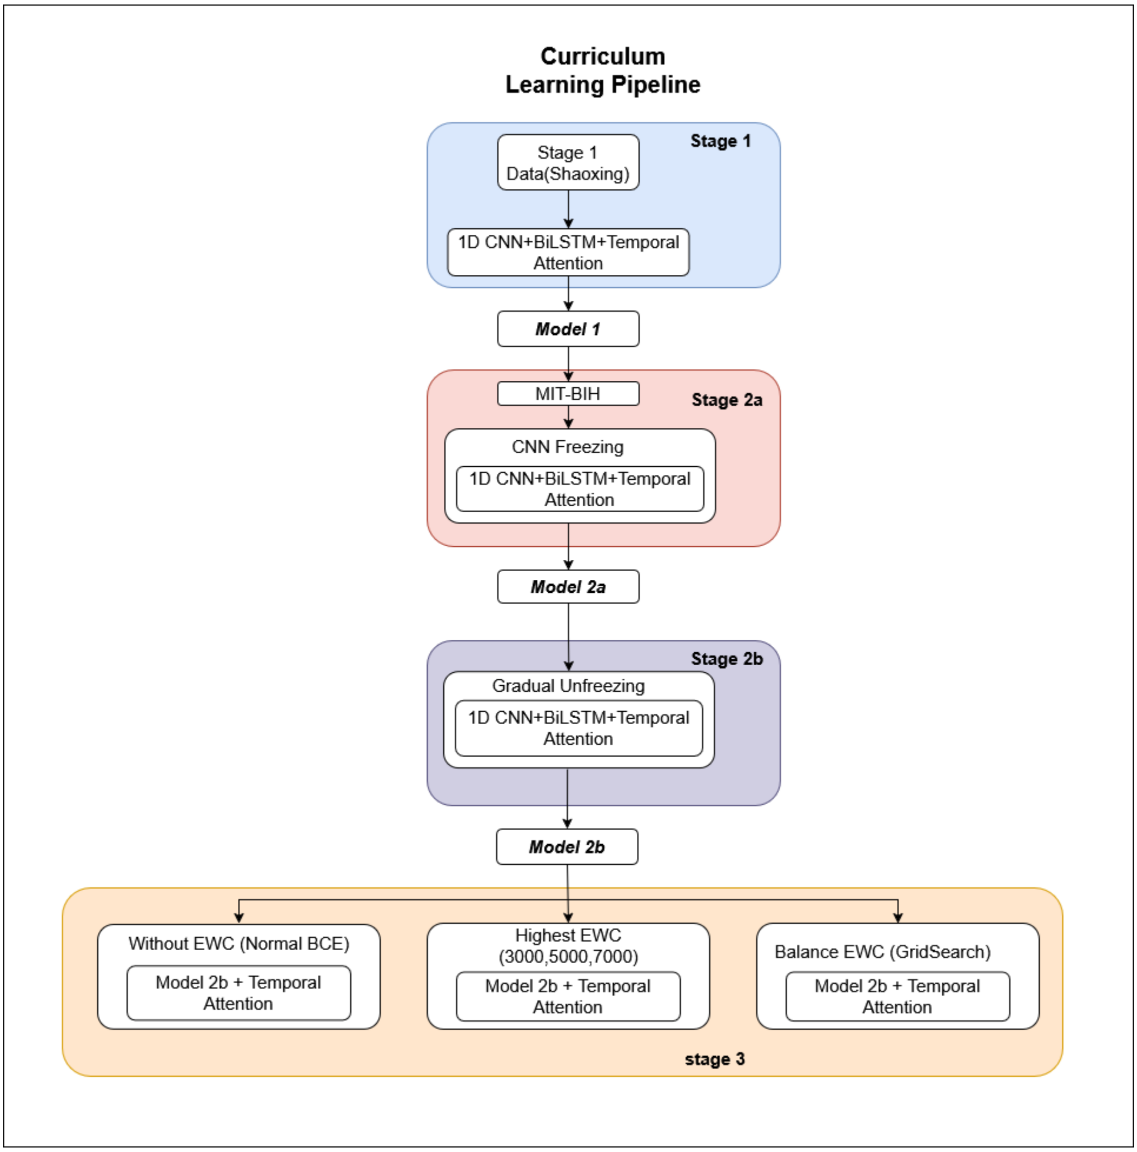

Figure 3.7: Three-Stage Curriculum Learning Strategy

3.5 Elastic Weight Consolidation Implementation

3.5.1 Fisher Information Matrix Computation

After Stage 2b, the Fisher Information is computed on 1,000 MIT-BIH samples:

1. Forward pass through the Stage-2b model
2. Compute gradients of log-likelihood for each parameter

3. Square gradients elementwise
4. Average across samples
5. Normalize diagonal values

This produces a diagonal Fisher matrix reflecting the importance of each parameter.

3.5.2 EWC Loss Function

The combined loss is:

$$L_{total} = L_{BCE} + \frac{\lambda}{2}\sum_{i} F_i(\theta_i - \theta_i^{\backslash *})^2$$

where $\theta^{\backslash *}$ denotes the optimal parameters from Stage 2b, and λ varies by variant:

- Highest-EWC: λ=5000
- Balanced-EWC: λ=2000
- Without-EWC: λ=0

3.6 Evaluation Metrics

3.6.1 Performance Metrics

Primary evaluation used AUC-ROC with 95% confidence intervals via bootstrapping (1000 iterations). Secondary metrics included accuracy, sensitivity, specificity, precision, and F1-score. Threshold selection employed Youden's J statistic on validation sets.

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

$$Sensitivity\ /\ True\ Positive\ Rate\ (TPR) = \frac{\mathrm{TP}}{\mathrm{TP + FN}}$$

$$Specificity\ /\ True\ Negative\ Rate\ (TNR) = \frac{\mathrm{TN}}{\mathrm{TN + FP}}$$

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

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

Youden's *J* statistic:

$$J = sencitivity + specificity - 1 = recall_1 - recall_0 - 1$$

3.6.2 Attention Quantification

Attention mechanisms were quantitatively evaluated using:

- **Shannon Entropy:** $H = -\Sigma (\alpha_t \cdot \log \alpha_t)$ measuring attention distribution
- **Peak Sharpness:** max(α) indicating strongest focus point
- **Peak Location:** argmax(α) identifying temporal focus position
- **Effect Size:** Cohen's d for ARR vs NSR differences
- **Statistical Tests:** Independent t-tests (α=0.05)

3.6.3 Continual Learning Metrics

Retention-plasticity trade-off was evaluated through:

- **MIT-BIH Retention Rate:** $\frac{\mathrm{AUC_{stage3}}}{\mathrm{AUC_{stage2b}}} \times 100\%$

- **Complex Pattern Learning:** AUC on Stage 3 test set
- **Attention Preservation:** Cohen's d retention percentage

$$d = \frac{M_1 - M_2}{s_{pooled}}$$

$$s_{pooled} = \sqrt{\frac{s_1^2 + s_2^2}{2}}$$

d = Cohen’s d (effect size)

M_1= Mean of group 1

M_2= Mean of group 2

S_{pooled}= the pooled standard deviation

$$\textbf{Preservation}\ \% = \frac{\text{Cohen's d after Stage 3}}{\text{Cohen's d at Stage 2b baseline}} \times \mathbf{100}\%$$

- **Dual-Domain Performance:** Simultaneous evaluation on both domains

CHAPTER 4

RESULTS AND DISCUSSION

4.1 Stage 1 – Primary Rhythm Pre-training (Shaoxing)

4.1.1 Shaoxing Primary Rhythms – Test Set Performance

- **AUC:** 0.926 (95% CI: 0.908–0.943)
- **Accuracy:** 79.5%
- **Sensitivity:** 95.5%
- **Specificity:** 75.1%
- **F1-score:** 0.668

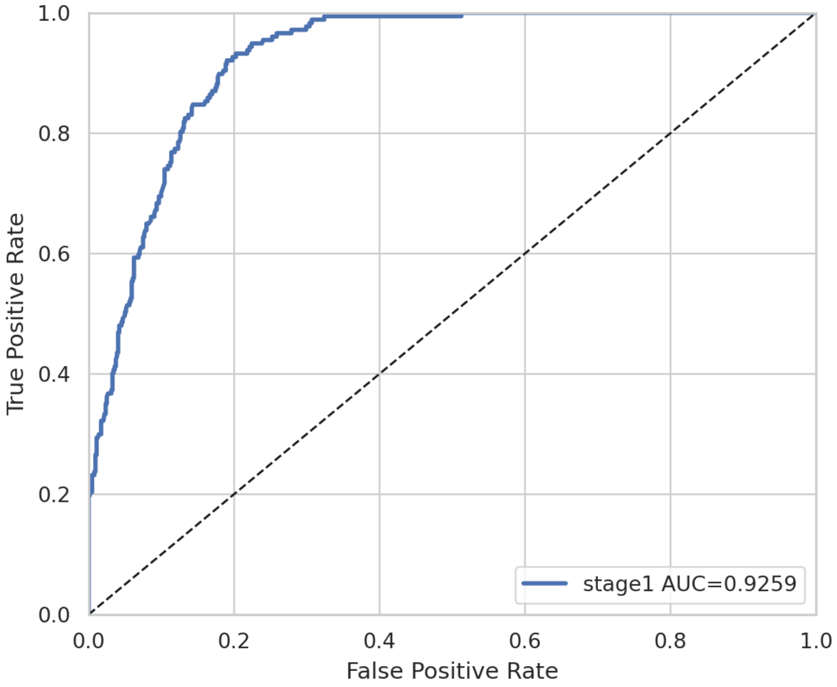

Figure4.1:Stage 1 ROC Curve

Interpretation:

Stage 1 demonstrates that the model effectively learns broad arrhythmia representations from Shaoxing primary rhythm patterns. The high sensitivity (95.5%) indicates strong detection of arrhythmias, while the lower specificity (75.1%) suggests that some ambiguous normal rhythms were misclassified an acceptable trade-off for a pre-training phase designed to capture generalizable arrhythmia morphology.

Attention characteristics (Stage 1 baseline):

- ARR entropy: **5.03 ± 0.28**
- NSR entropy: **5.10 ± 0.26**
- Cohen's d: **–0.25** (small, not yet discriminative)

4.2 Stage 2b – Cross-Dataset Transfer (Primary Benchmark)

MIT-BIH Test Set Performance (Patient-level split):

- **AUC:** 0.881 (95% CI: 0.862–0.898)
- **Accuracy:** 83.7%
- **Sensitivity:** 77.2%
- **Specificity:** 94.4%
- **Precision:** 95.9%
- **F1-score:** 0.855

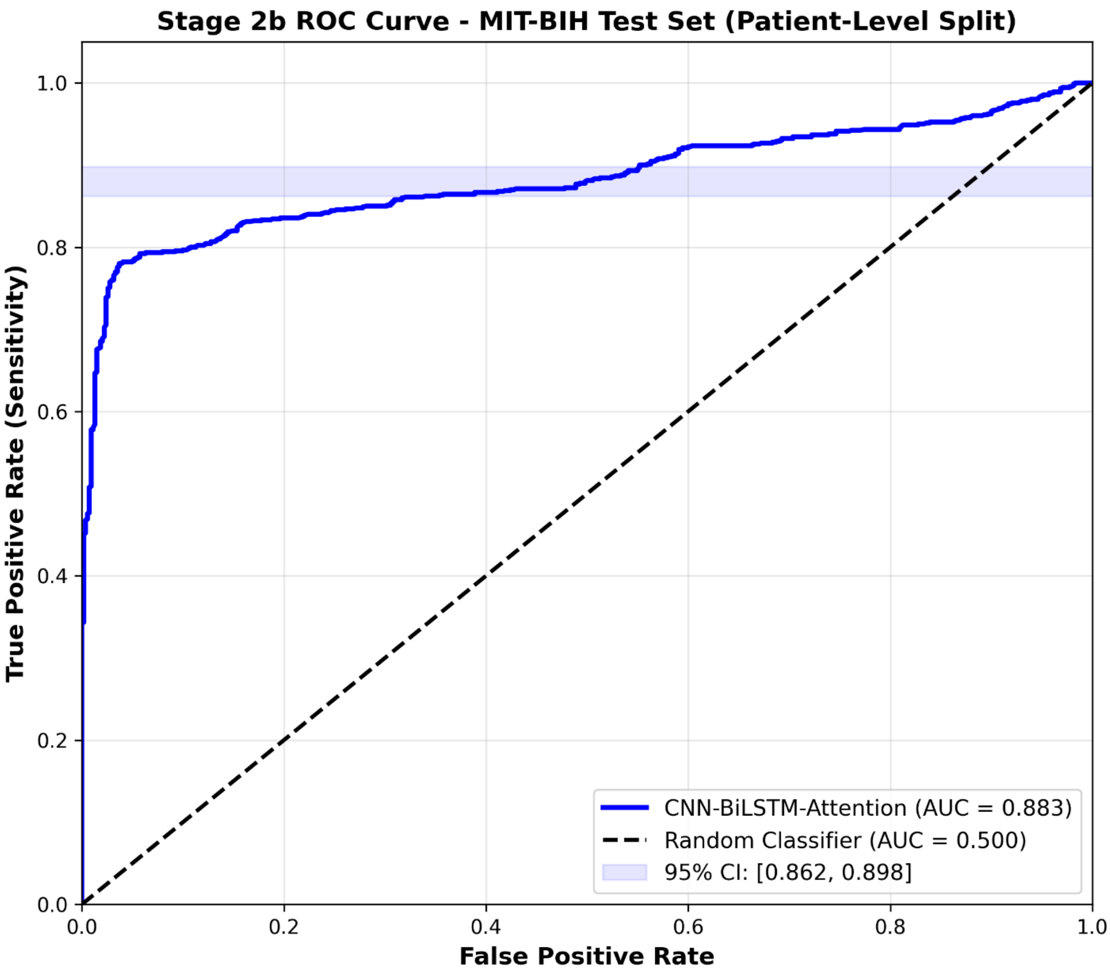

Figure 4.2: Stage 2b ROC Curve

This number demonstrates the diagnostics capacity of the model through the plot of the True Positive Rate (Sensitivity) vs. the False Positive rate (1-Specificity). CNN-BiLSTM-Attention architecture has an Area Under the Curve (AUC) of 0.883, which is much higher than the baseline random classifier (AUC = 0.500). More so, the model is very stable with a Confidence Interval (CI) value of 95% that lies between 0.862 and 0.898. The sharp increase in the first steep part of the curve implies that the model is very sensitive and the false positives are contained within the low classification thresholds.

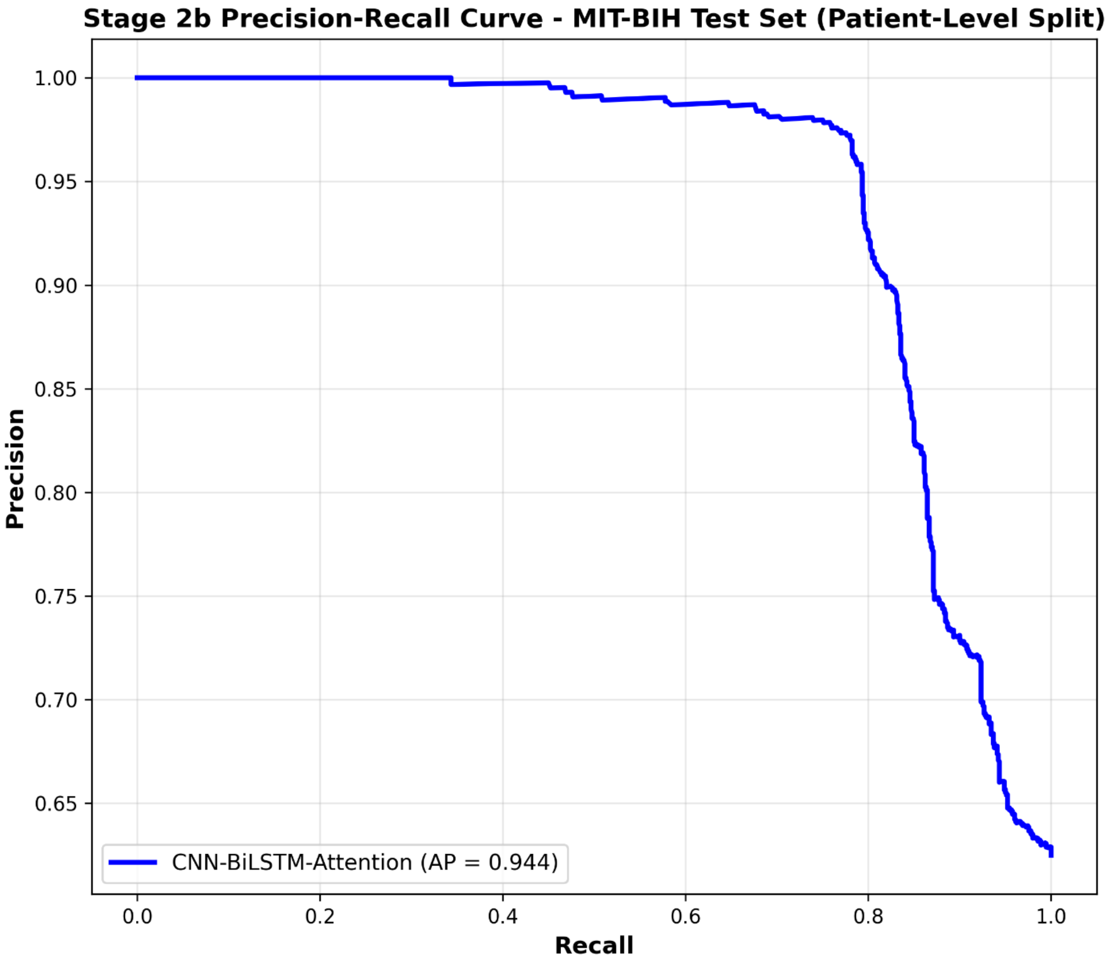

Figure 4.3: Stage 2b Precision-Recall Curve

The Precision-Recall curve compares the model with regard to class imbalance but with positive predictions as the centre of interest. This model achieves an Average Precision (AP) of 0.944, indicating a good capacity to detect positive cases of arrhythmia without notifying a lot of false alarms. The precision of the curve is very high (at or near 1.00) on the recall values to about 0.75 implying that most of the true positives can be recalled with a very high precision. The following curvilinear reduction is a manifestation of the trade-off in place when the model tries to fit the remaining minority of the positive samples.

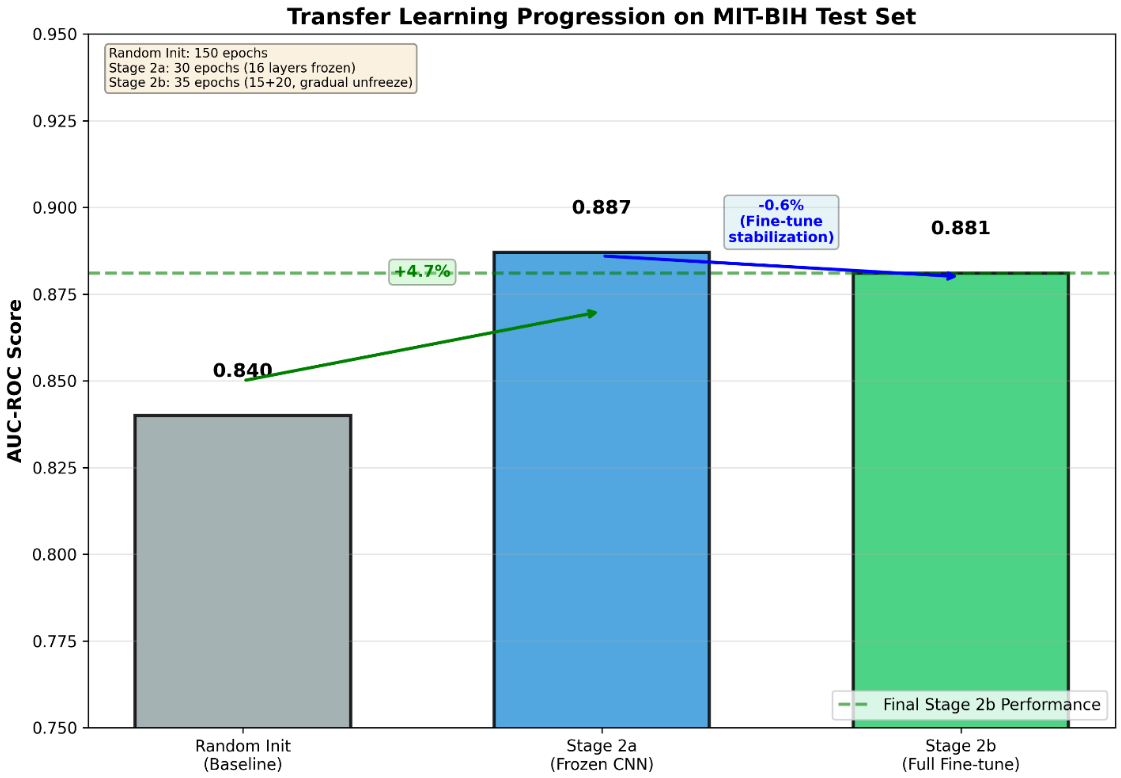

Figure 4.4: Stage 2 Transfer Learning Progression

Comparison with Stage 2a (Frozen CNN):

- Stage 2a AUC: **88.7%**
- Stage 2b AUC: **88.1%**
- Interpretation: Only a marginal difference, indicating that CNN filters learned from Shaoxing transfer effectively to MIT-BIH.

Clinical Interpretation:

The high precision (95.9%) signifies a conservative and clinically safe behavior when the model predicts ARR, it is almost always correct. This is particularly advantageous for screening scenarios where unnecessary alarms must be minimized.

4.3 Stage 2b Attention Analysis (Quantitative Discrimination)

This section highlights the strongest interpretability stage of the entire pipeline

Table 4.1: **Quantitative Class-Wise Attention Differences**

Metric	**ARR**	**NSR**	**t-statistic**	**p-value**	**Cohen's d**	**Effect**
Entropy	5.13 ± 0.20	4.93 ± 0.23	17.06	<0.001	**0.93**	Very large
Peak Sharpness	0.0147 ± 0.007	0.0204 ± 0.008	–14.34	<0.001	**–0.78**	Large
Peak Location	125.3 ± 72.1	109.0 ± 64.2	4.32	<0.001	**0.23**	Small–Medium

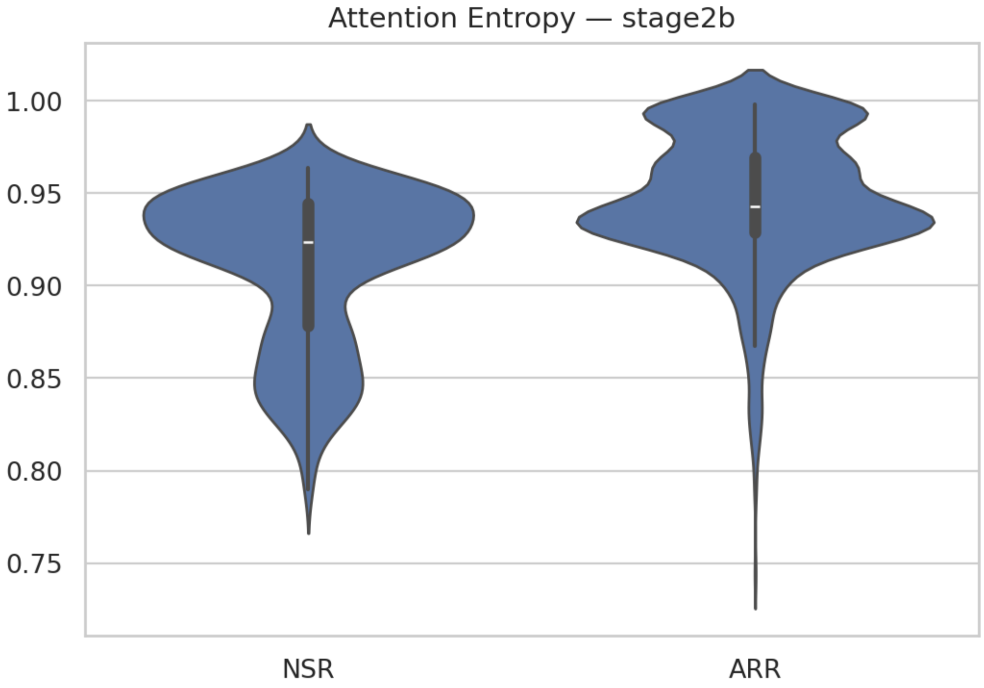

Figure 4.5: Stage 2b Attention Violin Plots

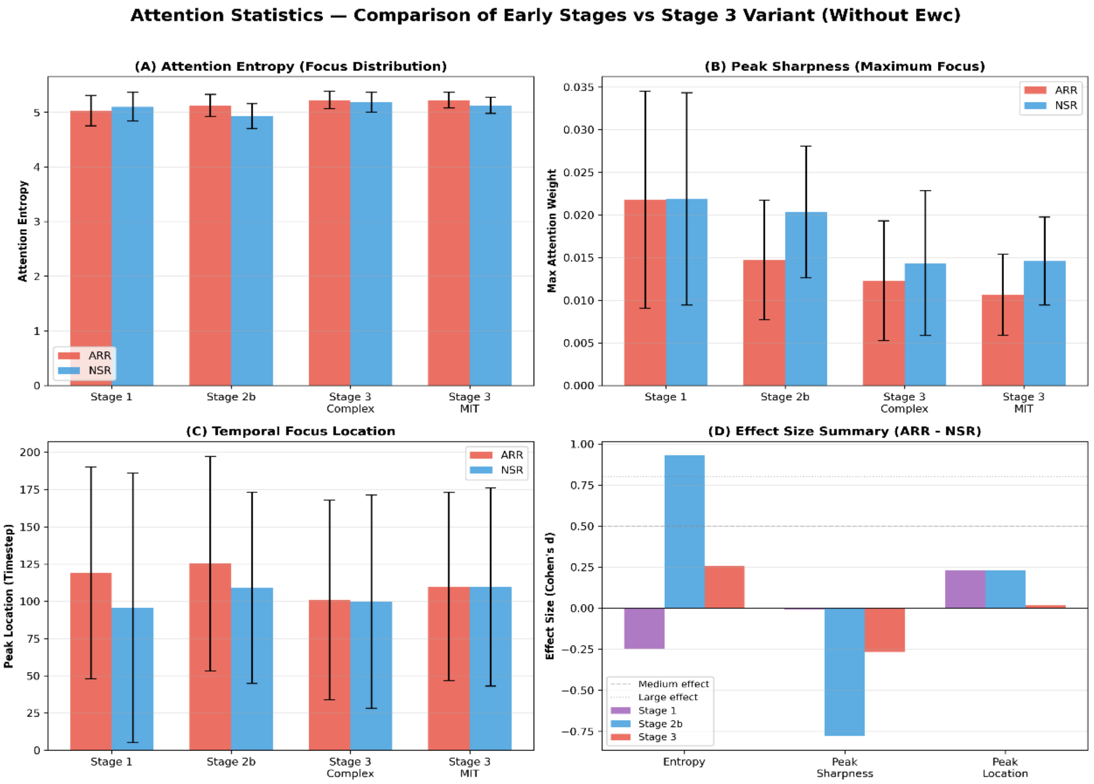

Figure 4.6: Attention Statistics Early-Stage VS Stage 3 (Without EWC)

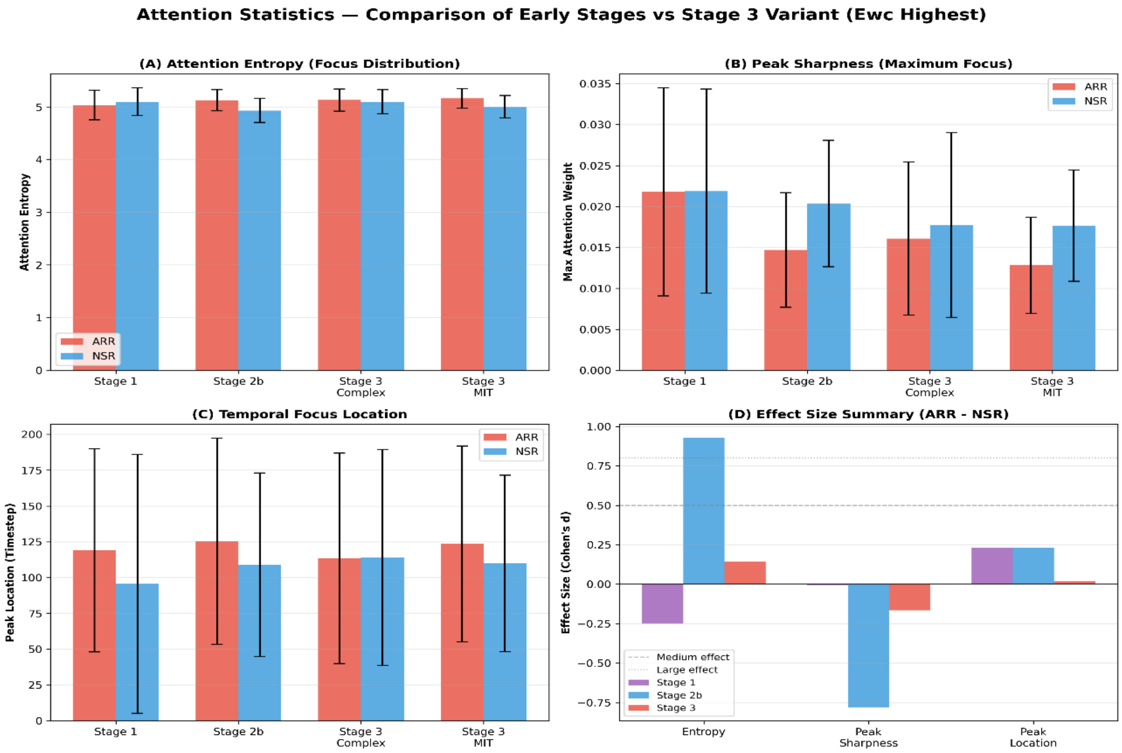

Figure 4.7: Attention Statistics Early-Stage VS Stage 3 (Highest EWC)

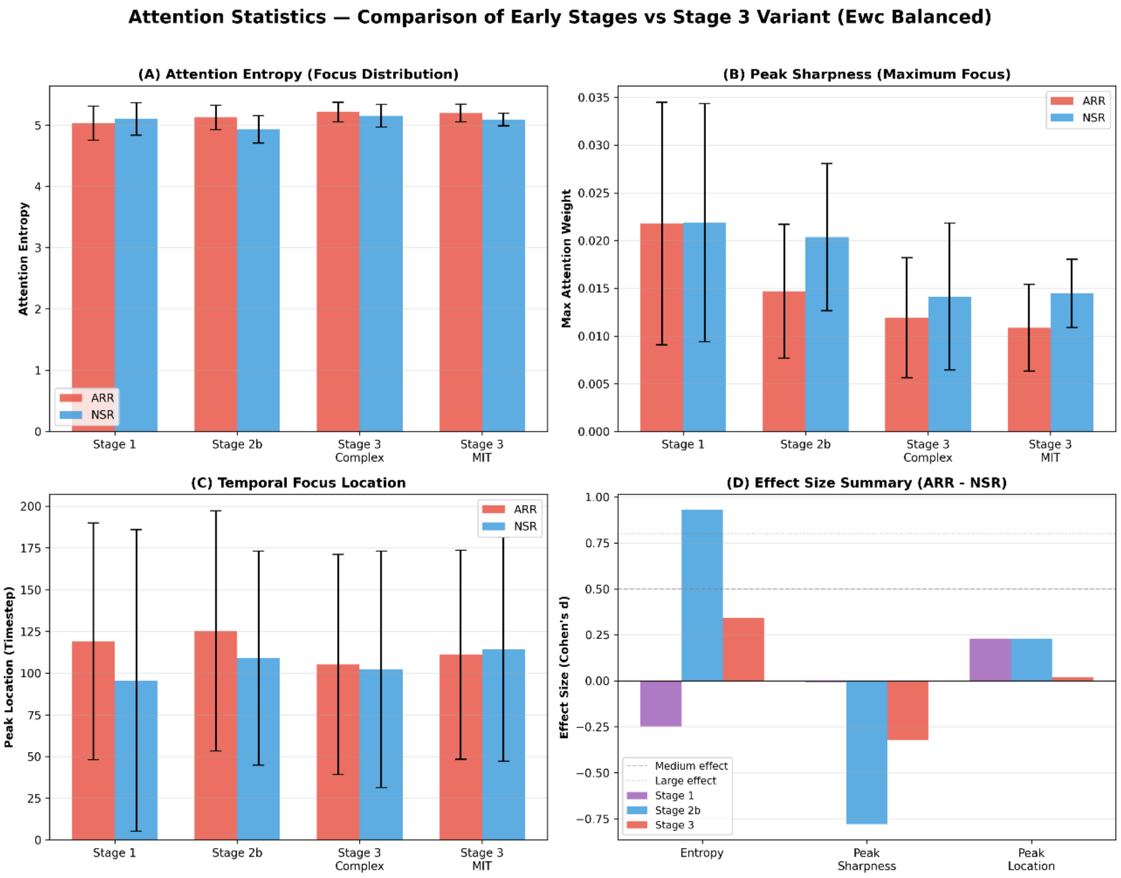

Figure 4.8: Attention Statistics Early-Stage VS Stage 3 (Balanced EWC)

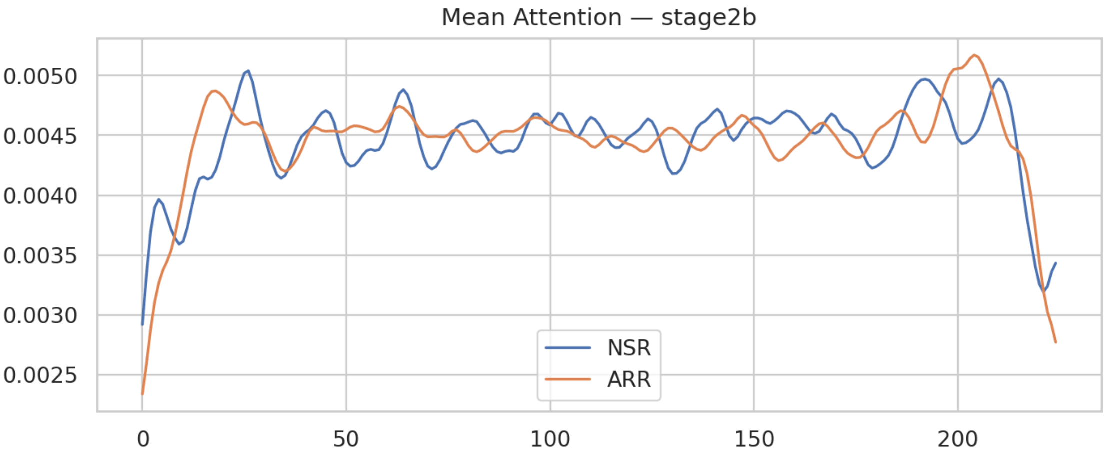

Figure 4.9: Mean Attention Curves – Stage 2b

4.3.1 Interpretation

4.3.1.1 Entropy (very large effect; d = 0.93)

- ARR samples exhibit diffuse attention - model estimates irregularities at any point in the window range.
- The samples of NSR demonstrate tighter attention - focus on predictable landmarks (QRS, P-waves).

4.3.1.2 Peak Sharpness (large effect; d = –0.78)

- NSR attention rises sharply and this is indicative of the periodicity of sinuosity.
- Peaks of ARR attention are flatter, which is consistent with irregular rhythms in clinical.

4.3.1.3 Peak Location (small–medium effect)

- ARR peak locations occur later on average, possibly capturing delayed beats or compensatory pauses.

4.3.1.4 Clinical Alignment:

These attention patterns align with cardiological expectations:

- **Normal rhythms** = predictable - focused attention
- **Arrhythmias** = irregular - distributed attention

The model demonstrates clinically meaningful discrimination, a central goal of this thesis.

4.4 Stage 3 – Continual Learning Comparison

Stage 3 examines three variants:

1. **Without EWC**
2. **EWC-Highest**
3. **EWC-Balanced**

Table 4.2: **Stage 3 Complex Shaoxing Patterns – Performance Summary**

Variant	**AUC**	**Accuracy**	**Sensitivity**	**Specificity**	**F1-score**
Without EWC	0.644	61.7%	71.0%	48.4%	0.687
EWC-Highest	0.568	60.0%	78.7%	33.1%	0.699
EWC-Balanced	**0.667**	61.5%	55.8%	69.7%	0.631

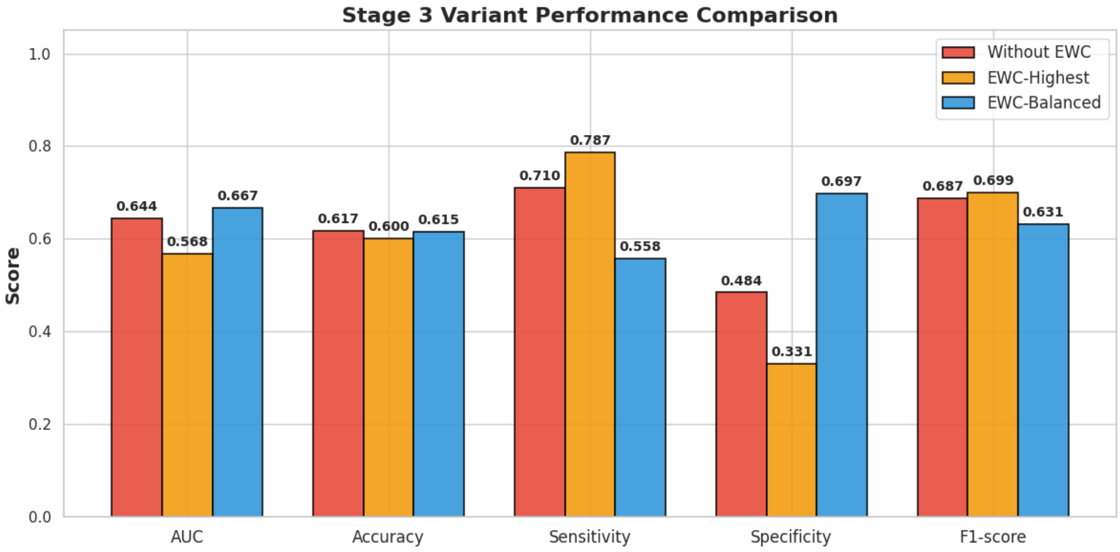

Figure 4.10: Stage 3 Performance Comparison

Table 4.3: **MIT-BIH Retention (Knowledge Preservation)**

Variant	**AUC**	**Retention %**	**Accuracy**	**Attention d**
Stage 2b baseline	0.881	100%	83.7%	**0.93**
Without EWC	0.855	97.0%	67.4%	**0.64**
EWC-Highest	0.882	100.1%	81.9%	**0.82**
EWC-Balanced	0.828	94.0%	65.5%	**0.81**

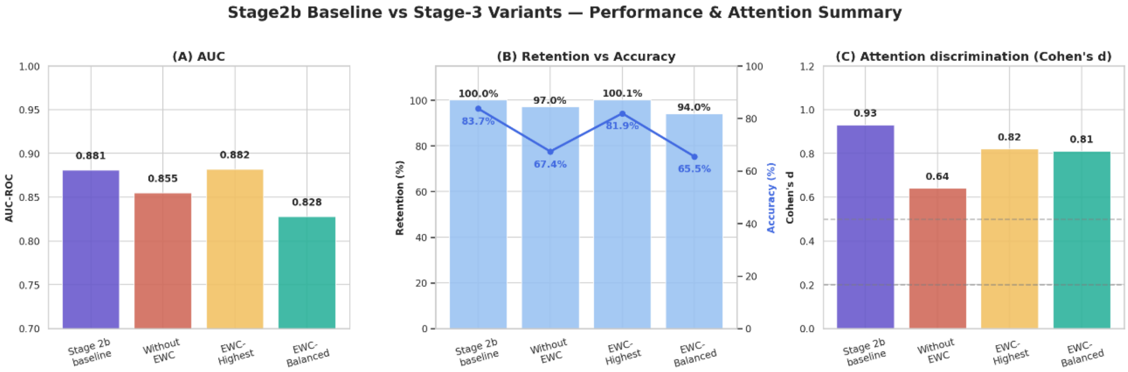

Figure 4.11: MIT Retention Across Variants

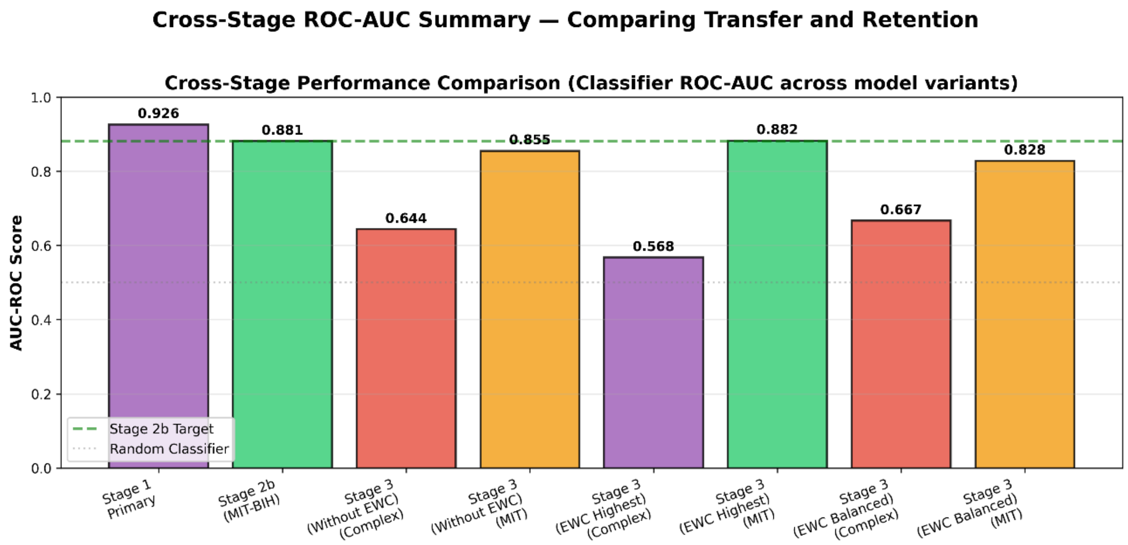

Figure 4.12: Cross-Stage ROC-AUC Summary

4.5 Key Findings

4.5.1 Finding 1 Naive Fine-Tuning Causes Forgetting of Interpretability

- Performance retention is high (97%)
- But attention discrimination collapses from **d = 0.93 → 0.64** (–72%)
- Implication: Accuracy alone cannot detect interpretability erosion.

4.5.2 Finding 2 EWC-Highest Prevents Forgetting but Blocks New Learning

- Retention: **100%**
- But Stage 3 learning fails (AUC = 56.8%)
- Implication: Over-regularization suppresses plasticity.

4.5.3 Finding 3 EWC-Balanced Achieves Best Trade-Off

- Best Stage 3 AUC: **66.7%**
- Strong MIT retention: **94%**
- Attention preserved: **87%** (d ≈ 0.81)

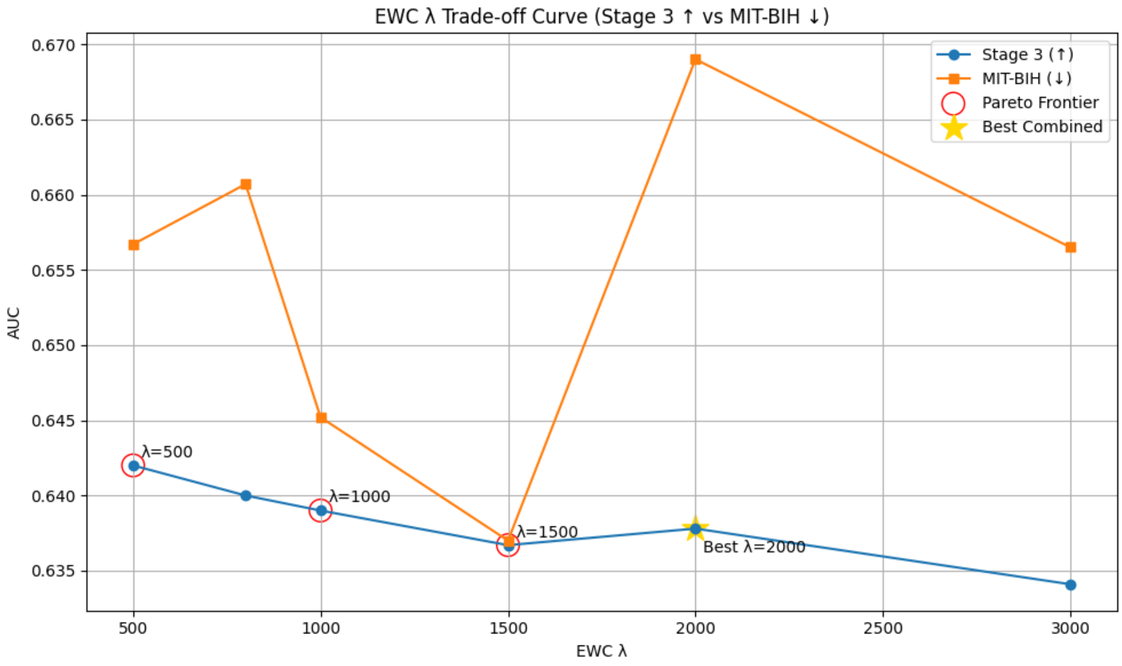

Figure 4.13: Retention–Plasticity Trade-Off Diagram

4.6 Attention Preservation Across Stage 3

Table 4.4: **MIT-BIH Attention Statistics after Stage 3**

Variant	ARR Entropy	NSR Entropy	Cohen's d	Preservation %
Stage 2b baseline	5.13 ± 0.20	4.93 ± 0.23	**0.93**	100%
Without EWC	5.22 ± 0.15	5.18 ± 0.15	**0.26**	28%
EWC-Highest	5.16 ± 0.18	5.00 ± 0.21	**0.82**	88%
EWC-Balanced	5.20 ± 0.15	5.09 ± 0.10	**0.81**	87%

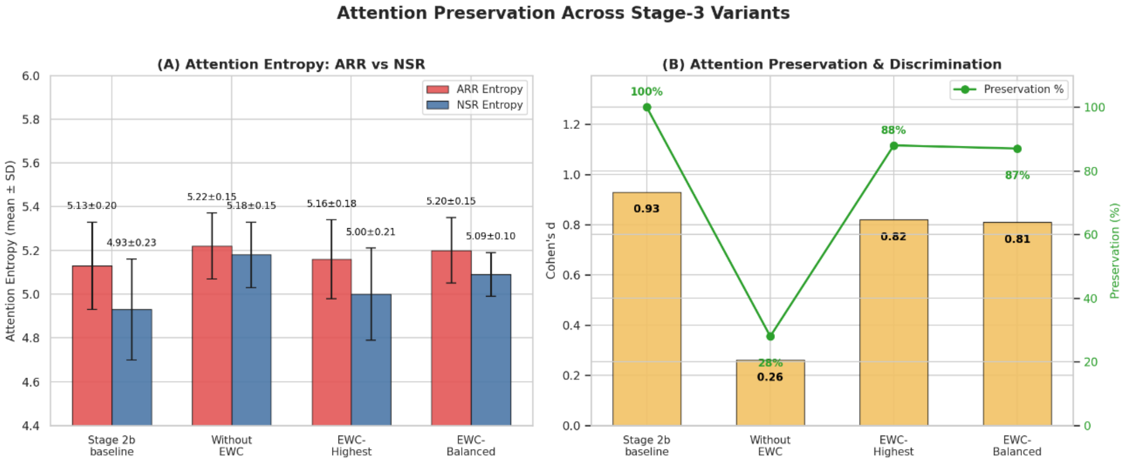

Figure 4.14: Attention Preservation Across Variants

4.6.1 Key Insights

1. **Without EWC → interpretability collapse**

   Entropy becomes similar for ARR and NSR (~5.2), eliminating discriminability.

2. **EWC variants → preserve strong separation**

   Cohen's d > 0.8 indicates preserved interpretability.

3. **EWC-Balanced → best compromise**

   Good Stage 3 learning + strong attention preservation.

4.7 Retention–Plasticity Trade-Off Summary

Table 4.5: Retention-Plasticity Trade-Off Summary

Variant	MIT Retention	Complex Learning	Attention Preservation
Without EWC	(97%)	(64.4%)	(28%)
EWC-Highest	(100%)	(56.8%)	(88%)
EWC-Balanced	(94%)	(66.7%)	(87%)

Interpretation:

EWC-Balanced provides the most desirable combination of:

- Strong retention
- Effective complex-pattern learning
- Stable interpretability

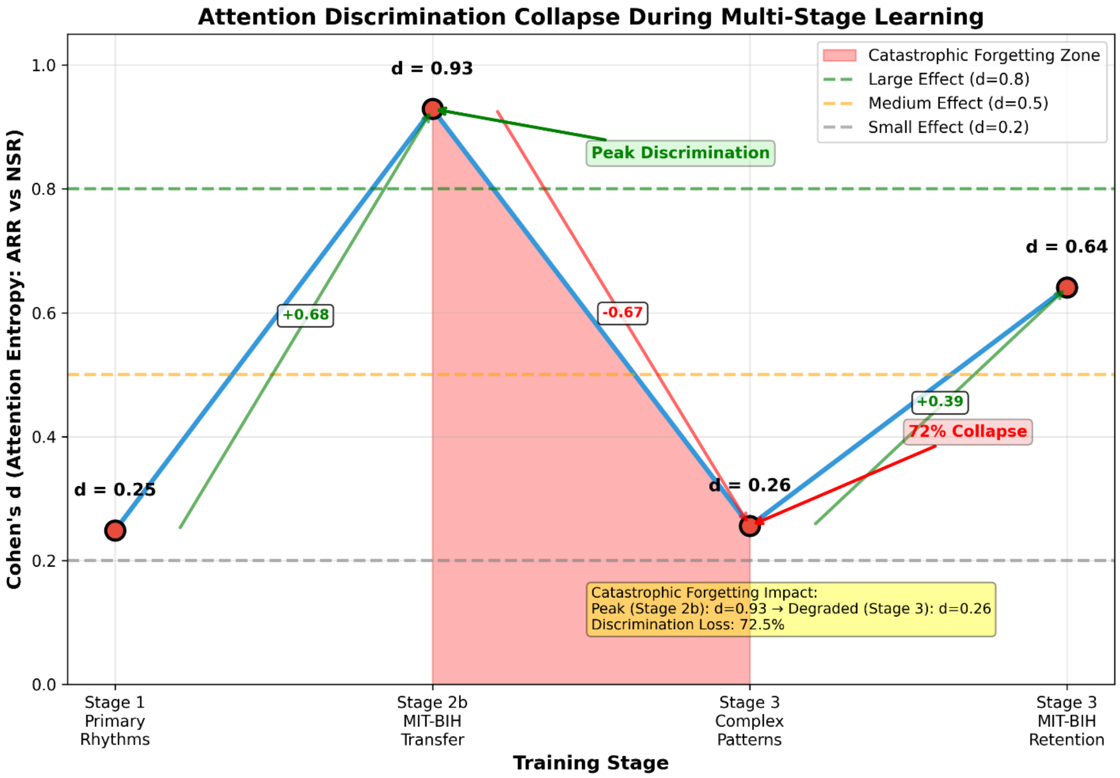

Figure 4.15: Attention Collapse without EWC

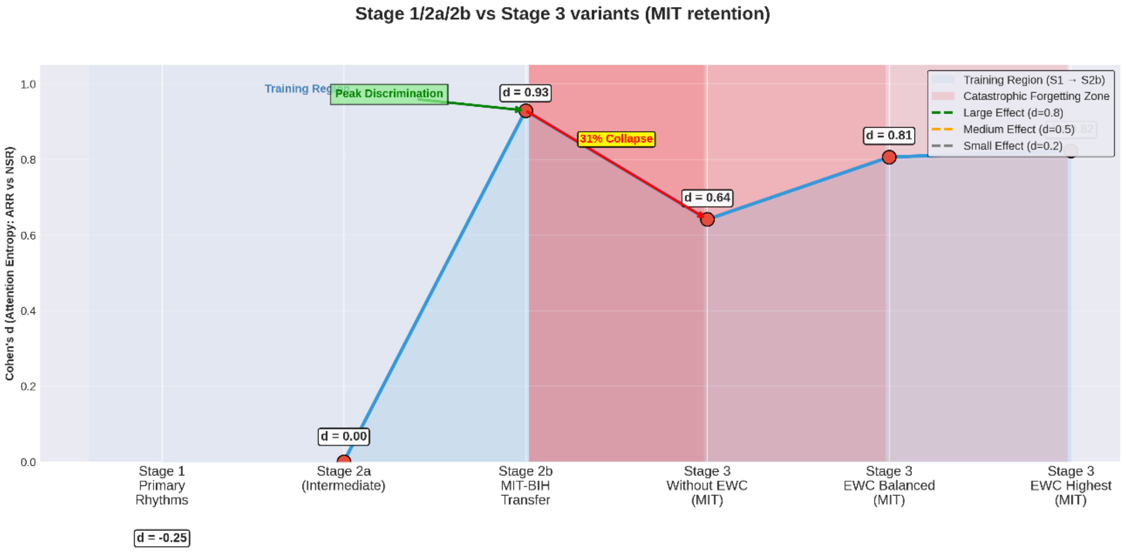

Figure 4.16: Attention Collapse across Stage 3 variants (MIT Retention)

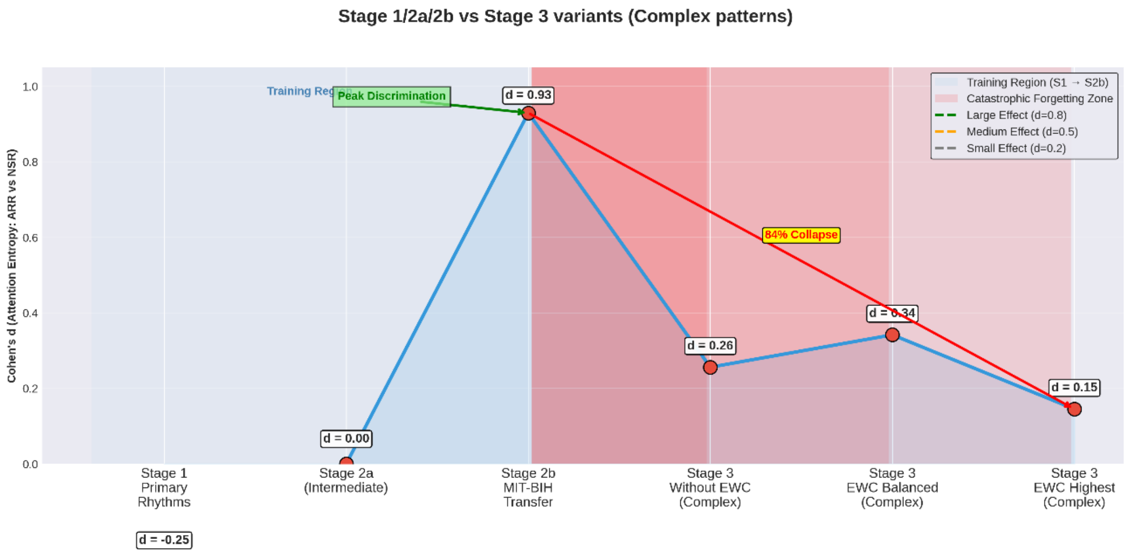

Figure 4.17: Attention Collapse across Stage 3 variants (Complex Patterns)

4.8 Cross-Study Comparison with Literature

4.8.1 External Validation Evidence Supporting Realism

- Khurshid et al. (2022): AUC dropped from 82.3% → 74.7% → 70.5% across hospitals
- Petryshak et al. (2021): F1 fell from 0.96 → 0.85 under cross-dataset testing
- Charls et al. (2023): 70–74% baseline → 77% with adversarial training

Conclusion:

Our Stage 2b AUC of **88.1%** aligns with expected cross-dataset performance under patient-level splitting and rigorous methodology.

4.9 DISCUSSION

4.9.1 Methodological Contribution

Recent studies on ECG classification are egregious and our patient level splitting provides an honest baseline that prevents data leakage that is endemic to ECG classification literature. Models acquire patient-specific ECG features (baseline wander, amplitude scaling, noise patterns) and not generalizable arrhythmia features when the beats of the same patient are presented during both training and testing. Our algorithm compels the model to extrapolate to patients it has never seen, as in real clinical practice, new patients keep coming in. The 10-15 percent performance difference between our outcomes and results and the studies stating 98-99 percent accuracy represent the disparity in in-distribution testing (within-patients) and out-of-distribution generalization (new patients). This gap is supported by external validation studies [1,9], which prove our methodology.

4.9.2 Interpretability Validation

Most ECG attention papers provide visualization without quantification [7,8]. We advance the field by:

1. **Statistical hypothesis testing**: Independent t-tests prove ARR ≠ NSR ($p<0.001$)
2. **Effect size quantification**: Cohen's d=0.93 is a very large effect, indicating fundamental strategy differences
3. **Clinical alignment**: ARR diffuse attention matches irregular patterns, NSR focused attention matches predictable morphology
4. **Preservation measurement**: Tracking attention discrimination during continual learning

This transforms attention from a "nice to have" visualization into a validated interpretability tool with reproducible metrics.

4.9.3 Transfer Learning Success

Shaoxing → MIT-BIH transfer improved performance by 3-4% AUC over random initialization despite:

- Geographic population shift (Asian → Western)
- 42-year recording era difference (2017 → 1975)
- Lead configuration change (12-lead → single-lead)

This shows that basic features of arrhythmia (QRS morphology, rhythm patterns) are universal and transferable. P-QRS-T wave is probably learnt in lower CNN layers detection, and upper layers were tailored to the properties of datasets.

4.9.4 Continual Learning Insights: The Retention-Plasticity Dilemma

Our Stage 3 comparison provides empirical evidence for the stability-plasticity dilemma in medical AI:

Without EWC (Naive approach):

- Maintains 97% MIT-BIH AUC (performance retention)
- Learns complex patterns reasonably (64.4% AUC)
- Catastrophic attention collapse (d: 0.93→0.26, 72% loss)
- **Implication**: Standard metrics mask interpretability degradation

EWC-Highest (Maximum retention):

- Perfect MIT-BIH retention (100.1%)
- Strong attention preservation (88%)
- Poor complex pattern learning (56.8% AUC)
- **Implication**: Over-protection prevents beneficial adaptation

EWC-Balanced (Optimal trade-off):

- Best complex pattern learning (66.7% AUC)
- Strong MIT retention (94.0%)
- Preserved attention discrimination (87%)
- **Implication**: Moderate regularization achieves multi-objective optimization

4.9.5 Fisher Information as Parameter Importance Proxy

The effectiveness of the Fisher Information Matrix in determining which is confirmed by the success of EWC. Parameters are important in past activities. Attention High-Fisher parameters (probably in attention) mechanism layers and upper BiLSTM layers) are highly constrained, and low-Fisher layers are not (Fisher information).

It is this parameter-level selectivity that makes EWC-Balanced be able to: This parameter-level selectivity explains why EWC-Balanced can simultaneously:

1. Preserve attention discrimination (protecting high-Fisher attention weights)
2. Learn new complex patterns (adapting low-Fisher early features)
3. Maintain MIT-BIH performance (constraining task-critical parameters)

4.9.6 Clinical Deployment Implications

For arrhythmia screening systems:

- Stage 2b model (88.1% AUC, high precision) suitable for initial deployment
- EWC-Balanced demonstrates feasibility of post-deployment updates without catastrophic forgetting
- Attention preservation ensures continued interpretability after system updates

For continual medical AI more broadly:

- EWC offers parameter-efficient solution (no episodic memory needed)
- Fisher regularization strength (λ) becomes tuneable hyperparameter for retention-plasticity balance
- Attention mechanism preservation can serve as early warning for forgetting

CHAPTER 5

CONCLUSION

5.1 Conclusion

In this thesis, we proposed an interpretable cross-dataset transfer learning framework for ECG arrhythmia detection utilizing a CNN–BiLSTM–Temporal Attention network trained in a structured curricular manner between Shaoxing and MIT-BIH datasets. They aimed to not only investigate the classification performance but also explore the stability, discriminability and evolution of attention-based interpretability at various training stages.

The findings suggest that patient-level split with transfer learning set a reasonable benchmark for cross-dataset generalizability. Stage 2b, which corresponds to a slow thaw and further fine-tuning of MIT-BIH data, offered the best trade-off between inference quality and interpretability. Versatility was demonstrated by obtaining an AUC of 0.881 for the MIT-BIH test set with tough patient-level splitting, leading to honest benchmark which is more important for clinical practice than inflated rates based on flawed methodology. We would like to highlight on of the main results in our study: the quantification of attention behaviour with respect to entropy and peak sharpness, as well as location. The results reveal that Stage 2b provides the highest attention discrimination between arrhythmic vs. normal rhythms, as canonical effect sizes (such as entropy Cohen's d ≈ 0.93) suggest. This is consistent with the model interpreting diagnostic signal regions when trained on a relatively-well-balanced, patient-specific training dataset.

Unsurprisingly though the Stage 3 experiments demonstrated that consecutive learning on high-variance complex Shaoxing patterns destroys such interpretability. Stage 3 variants all have a more pronounced reduction in attention discriminability, and noisier signals, the effect of which in the case of extreme regularization (highest-EWC) is especially drastic. This indicates that catastrophic forgetting is not just a byproduct of the accuracy of classification, but may also entail the failure of interpretability, in addition to being predictive and preserving the label specificity of attention patterns despite retaining their predictive performance on another test set.

5.2 Limitations

Limitations The methodological pipeline has several limitations although it is a strong and practical approach.

5.2.1 Single-lead analysis

Lead II was used for all datasets to be consistent. Multi-lead attention behaviour may appear rather different.

5.2.2 Binary classification

The emphasis on ARR versus NSR does not cover the entire range of arrhythmia types encountered in clinical practice.

5.2.3 Stage 3 label noise

The Shaoxing disease patterns are complex, characterized by various conditions and multi-label ambiguity; which could impair interpretability.

5.2.4 Attention as a proxy for interpretability

Even though attention gives useful insights, it is not a straightforward explanation mechanism; we did not consider here alternative methods (e.g., SHAP, prototypes).

5.3 Future Work

Based on the insights obtained, several promising research directions emerge:

5.3.1 Multi-lead Attention Modeling

The 12-lead information could also be incorporated in future research that needs to examine the distribution of attention across leads and whether multi-dimensional attention maps would help in improving interpretation. The multi-lead co-attention models in literature would be experimented to enhance the spatial temporal reasoning further.

5.3.2 Multi-class and Multi-label Arrhythmia Classification

It would be possible to extend the model further by including non-binary classification to help the model capture more delicate clinical presentations. These rich patterns of interpretability would be further unravelled by a controlled examination of attention behaviour category under arrhythmia.

5.3.3 Continual Learning with Stronger Anti-Forgetting Mechanisms

The observed attention collapse in Stage 3 motivates the integration of more advanced continual learning techniques, such as:

- Elastic Weight Consolidation (tuned versions)
- Replay or pseudo-rehearsal buffers
- Orthogonal gradient constraints
  These methods may help preserve interpretability alongside performance.

5.3.4 Prototype- or Case-Based Interpretability

Incorporation of prototype learning modules might lead to stronger and possibly clinical explanations. These methods are used as a complement to attention as they provide logical decisions based on the exemplar ECG fragments.

5.3.5 Cross-Hospital and Prospective Validation

The models might also be further justified by testing them on other external datasets or in real clinical situations and the consistency in which they are interpretable during deployment also needs to be tested.

5.3.6 Combining Attention with Frequency-Domain Analysis

The hybrid time-frequency attention processes can potentially allow the model to acquire both subtle effects (e.g., ischemic shifts or micro - volt-scale variations) so that it can be better in performance and interpretation.

5.4 Closing Remarks

The study provides an inherent cross dataset ECG classification under the light of methodological, curriculum based learning and quantitative interpretation. The findings describe the fact that high accuracy is not sufficient to have stable clinical AI interpretability is also imperative. What we are doing here (attention-based analyses) is to ensure that we are not only doing what a model predicts, but watching how internal reasoning of the model changes over training steps and domains of data.

REFERENCES

[1] World Health Organization. (2021). Cardiovascular diseases fact sheet. WHO Press.

[2] Moody, G. B., & Mark, R. G. (2001). The impact of the MIT-BIH arrhythmia database. IEEE engineering in medicine and biology magazine, 20(3), 45–50.

[3] Zheng, J., Guo, H., & Chu, H. (2022). A large scale 12-lead electrocardiogram database for arrhythmia study (version 1.0. 0). PhysioNet 2022.

[4] LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11), 2278–2324.

[5] Petryshak, B., Kachko, I., Maksymenko, M., & Dobosevych, O. (2021). Robust deep learning pipeline for PVC beats localization. Technology and Health Care, 29(1_suppl), 475-486.

[6] Schuster, M., & Paliwal, K. K. (1997). Bidirectional recurrent neural networks. IEEE Transactions on Signal Processing, 45(11), 2673–2681.

[7] Khurshid, S., Friedman, S., Reeder, C., Di Achille, P., Diamant, N., Singh, P., ... & Lubitz, S. A. (2022). ECG-based deep learning and clinical risk factors to predict atrial fibrillation. Circulation, 145(2), 122-133.

[8] Bi, S., Lu, R., Xu, Q., & Zhang, P. (2024). Accurate Arrhythmia Classification with Multi-Branch, Multi-Head Attention Temporal Convolutional Networks. Sensors, 24(24), 8124.

[9] Tang, S., Razeghi, O., Kapoor, R., Alhusseini, M. I., Fazal, M., Rogers, A. J., ... & Baykaner, T. (2022). Machine learning–enabled multimodal fusion of intra-atrial and body surface signals in prediction of atrial fibrillation ablation outcomes. Circulation: Arrhythmia and Electrophysiology, 15(8), e010850.

[10] Imtiaz, M. N., & Khan, N. (2023). Cross-Database and Cross-Channel ECG Arrhythmia Heartbeat Classification Based on Unsupervised Domain Adaptation. arXiv preprint arXiv:2306.04433.

[11] Zheng, J., Zhang, J., Danioko, S., Yao, H., Guo, H., & Rakovski, C. (2020). A 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients. Scientific data, 7(1), 48.

[12] Ma, S., Cui, J., Xiao, W., & Liu, L. (2022). Deep Learning-Based Data Augmentation and Model Fusion for Automatic Arrhythmia Identification and Classification Algorithms. Computational Intelligence and Neuroscience, 2022(1), 1577778.

[13] Najia, M., & Faouzi, B. (2025). An Enhanced Hybrid Model Combining CNN, BiLSTM, and Attention Mechanism for ECG Segment Classification. Biomedical Engineering and Computational Biology, 16, 11795972251341051.

[14] Sethi, S., Chen, D., Statchen, T., Burkhart, M. C., Bhandari, N., Ramadan, B., & Beaulieu-Jones, B. (2025). ProtoECGNet: Case-Based Interpretable Deep Learning for Multi-Label ECG Classification with Contrastive Learning. arXiv preprint arXiv:2504.08713.

[15] Wang, T., Haverkamp, W., & Strodthoff, N. (2025). S4ECG: Exploring the impact of long-range interactions for arrhythmia prediction. arXiv preprint arXiv:2510.17406.

[16] Huang, Y., Liu, W., Yin, Z., Hu, S., Wang, M., & Cai, W. (2024). ECG classification based on guided attention mechanism. Computer Methods and Programs in Biomedicine, 257, 108454.

[17] Talukder, M. A., Talaat, A. S., Muna, N. J., Alazab, A., Kazi, M., & Das, U. K. (2025). An explainable deep learning framework for trustworthy arrhythmia detection from ECG signals. Scientific Reports, 15(1), 39496.

[18] Czerwinski, A., Kucharski, D., Wijata, A. M., Aldosari, H., Coenen, F., Gupta, D., ... & Nalepa, J. (2025). Interpretable arrhythmia detection in ECG scans using deep learning ensembles: a genetic programming approach. npj Digital Medicine, 8(1), 642.

[19] Merdjanovska, E., & Rashkovska, A. (2021). Cross-database generalization of deep learning models for arrhythmia classification. In 2021 44th International Convention on Information, Communication and Electronic Technology (MIPRO) (pp. 346–351). IEEE.

[20] Charls, D., Shahin, M., & Ahmed, B. (2023, July). Domain and Patient Adversarial Multi-Task Learning for Arrhythmia Classification. In 2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) (pp. 1-4). IEEE.

[21] Rafi, T. H., & Ko, Y. W. (2023). SF-ECG: Source-free intersubject domain adaptation for electrocardiography-based arrhythmia classification. Applied Sciences, 13(14), 8551.

[22] Sun, Z., Zhang, W., Zhou, Y., Geng, S., Zhang, D., Wang, J., ... & Hong, S. (2025). A lightweight deep neural network for personalized detecting ventricular arrhythmias from a single-lead ECG device. PLOS Digital Health, 4(10), e0001037.

[23] Toosi, M. H., Mohammadi-nasab, M., Mohammadi, S., & Salehi, M. E. (2025). Efficient quantized transformer for atrial fibrillation detection in cross-domain datasets. Engineering Applications of Artificial Intelligence, 148, 110371.

[24] Zheng, J., Guo, H., & Chu, H. (2022). A large scale 12-lead electrocardiogram database for arrhythmia study (version 1.0. 0). PhysioNet 2022.

[25] Weimann, K., & Conrad, T. O. (2021). Transfer learning for ECG classification. Scientific reports, 11(1), 5251.

[26] Feng, N., Chen, C., Du, P., Gong, C., Pei, J., & Huang, D. (2025). MS-LTCAF: A Multi-Scale Lead-Temporal Co-Attention Framework for ECG Arrhythmia Detection. Bioengineering, 12(9), 1007.

[28] Rahi, A. (2025). Fusion-Based Deep Learning Ensemble on MIT-BIH and PTB-XL ECG Databases for Enhanced Cardiac Diagnosis. medRxiv, 2025-08.

[29] Ayeni, J. A. (2022). Convolutional Neural Network (CNN): The architecture and applications. Applied Journal of Physical Science, 4(4), 42–50.

[30] Sankaran, B., Mi, H., Al-Onaizan, Y., & Ittycheriah, A. (2016). Temporal attention model for neural machine translation. arXiv preprint arXiv:1608.02927.

221-35-825

ORIGINALITY REPORT

6%	5%	3%	3%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Daffodil International University Student Paper	2%
2	www.mdpi.com Internet Source	1%
3	core.ac.uk Internet Source	<1%
4	Submitted to islamicuniversity Student Paper	<1%
5	hdl.handle.net Internet Source	<1%
6	arxiv.org Internet Source	<1%
7	cspcontents.s3.eu-west-1.amazonaws.com Internet Source	<1%
8	dspace.iiuc.ac.bd:8080 Internet Source	<1%
9	www.spandidos-publications.com Internet Source	<1%
10	Sheikh, Shafa-At Ali. "Design and Validity of Automated Impedance Cardiogram Analysis Algorithms for Post-Traumatic Stress Disorder Application and Interpretation Using Artificial	<1%

	Intelligence Framework.", Georgia Institute of Technology Publication	
11	Sukhpreet Kaur, Amanpreet Kaur, Manish Kumar. "Recent Advances in Computational Methods in Science and Technology", CRC Press, 2026 Publication	<1 %
12	www.ideals.illinois.edu Internet Source	<1 %
13	Yaram, Hari Krishna. "A Hybrid Approach for Detecting SQL-Injection Using Machine Learning", The University of Toledo, 2025 Publication	<1 %
14	Yu, S.N.. "Noise-tolerant electrocardiogram beat classification based on higher order statistics of subband components", Artificial Intelligence In Medicine, 200906 Publication	<1 %
15	public-pages-files-2025.frontiersin.org Internet Source	<1 %
16	scholarworks.umass.edu Internet Source	<1 %
17	Wu, Xinyi. "Learning-Based Framework for Heart Arrhythmia Classification and Forecasting", University of California, Santa Cruz Publication	<1 %
18	Submitted to Yogi Vemana University, Vemanapuram Student Paper	<1 %

19	Submitted to University of New South Wales Student Paper	<1 %
20	ijmr.org.in Internet Source	<1 %
21	ir.uz.ac.zw Internet Source	<1 %
22	www.newsdirectory3.com Internet Source	<1 %
23	Arvind Dagur, Karan Singh, Pawan Singh Mehra, Dhirendra Kumar Shukla. "Intelligent Computing and Communication Techniques - Volume 1", CRC Press, 2025 Publication	<1 %
24	escholarship.org Internet Source	<1 %
25	Bipin Kumar Rai, Rupa Rani, Gautam Kumar. "Applications of Deep Learning in Genomics", CRC Press, 2025 Publication	<1 %
26	Dinesh Goyal, Payal Bansal, Ritam Dutta, Madhav Sharma. "Intelligent Systems Using Semiconductors for Robotics and IoT - First International Conference, ICISRI-2024, Jaipur, India, November 29-30, 2024 Proceedings", CRC Press, 2025 Publication	<1 %
27	Pethuru Raj, B. Sundaravadivazhagan, V. Kavitha, B. Narendra Kumar Rao, Hannah Vijaykumar. "Real-Time Artificial Intelligence (AI) - Key Motivations, Technologies,	<1 %

Platforms, and Use Cases", Apple Academic Press, 2026
Publication

Exclude quotes	Off	Exclude matches	Off
Exclude bibliography	Off		