

Rohit Beniwal

Readmission Risk Prediction using Explainable AI based on Multimodal EHR Data

 Research Papers

Document Details

Submission ID

trn:oid:::27535:144231299

Submission Date

Jun 25, 2026, 9:40 PM GMT+5:30

Download Date

Jun 25, 2026, 9:47 PM GMT+5:30

File Name

Readmission Risk Prediction using Explainable AI based on Multimodal EHR Data.pdf

File Size

1.6 MB

47 Pages

12,342 Words

71,856 Characters

7% Overall Similarity

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

Filtered from the Report

- Bibliography
- Quoted Text
- Small Matches (less than 14 words)

Match Groups

-  **27 Not Cited or Quoted 6%**
Matches with neither in-text citation nor quotation marks
-  **3 Missing Quotations 0%**
Matches that are still very similar to source material
-  **0 Missing Citation 0%**
Matches that have quotation marks, but no in-text citation
-  **0 Cited and Quoted 0%**
Matches with in-text citation present, but no quotation marks

Top Sources

- 5%  Internet sources
- 1%  Publications
- 5%  Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

Match Groups

- 27 Not Cited or Quoted 6%**
Matches with neither in-text citation nor quotation marks
- 3 Missing Quotations 0%**
Matches that are still very similar to source material
- 0 Missing Citation 0%**
Matches that have quotation marks, but no in-text citation
- 0 Cited and Quoted 0%**
Matches with in-text citation present, but no quotation marks

Top Sources

- 5% Internet sources
- 1% Publications
- 5% Submitted works (Student Papers)

Top Sources

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

1	Student papers	Delhi Technological University on 2026-05-28	1%
2	Student papers	Delhi Technological University on 2026-05-29	<1%
3	Student papers	Delhi Technological University on 2026-05-29	<1%
4	Internet	link.springer.com	<1%
5	Student papers	Delhi Technological University on 2026-06-12	<1%
6	Student papers	Delhi Technological University on 2026-06-09	<1%
7	Internet	openarchive.ki.se	<1%
8	Internet	www.frontiersin.org	<1%
9	Student papers	UCL on 2026-01-30	<1%
10	Student papers	Delhi Technological University on 2026-05-20	<1%

11	Internet	dr.ddn.upes.ac.in:8080	<1%
12	Student papers	Delhi Technological University on 2026-05-25	<1%
13	Internet	ogma.newcastle.edu.au	<1%
14	Student papers	American Intercontinental University Online on 2026-05-20	<1%
15	Publication	Silvio Rizzi, Giulio Saroglia, Violeta Kalemi, Simona Rimoldi, Genciana Terova. "Art...	<1%
16	Student papers	Delhi Technological University on 2026-05-29	<1%
17	Publication	Qianxi Zhao, Yuqing Zou, Xiyang Li. "Spatial-temporal masked modeling pre-train...	<1%
18	Internet	pergamos.lib.uoa.gr	<1%
19	Internet	www.coursehero.com	<1%
20	Student papers	Delhi Technological University on 2026-06-12	<1%
21	Internet	dl.lib.uom.lk	<1%
22	Internet	www.ijert.org	<1%
23	Internet	www.mdpi.com	<1%
24	Student papers	Georgia Southern University on 2026-03-19	<1%

25

Student papers

Indian Institute of Science, Bengaluru on 2026-06-20

<1%

**READMISSION RISK PREDICTION USING EXPLAINABLE AI
BASED ON MULTIMODAL EHR DATA**

12

**A Thesis Submitted
In Partial Fulfillment of the Requirements
for the Degree of
MASTER OF TECHNOLOGY
in
Artificial Intelligence**

**by
Gautam Aggarwal
(24/AFI/02)**

1

**Under the Supervision of
Dr. Rohit Beniwal**
Assistant Professor, Department of Computer Science and Engineering
Delhi Technological University



**Department of Computer Science and Engineering
DELHI TECHNOLOGICAL UNIVERSITY**

(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042, India

May, 2026

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daultapur, Main Bawana Road, Delhi-110042, India

CANDIDATE'S DECLARATION

I, Gautam Aggarwal, Roll No. 24/AFI/02, hereby certify that the work which is being presented in this thesis entitled "Readmission Risk Prediction using Explainable AI based on Multimodal EHR Data" in partial fulfilment of the requirements for the award of the Degree of Master of Technology in Artificial Intelligence, submitted in the Department of Computer Science and Engineering, Delhi Technological University, is an authentic record of my own work carried out during the period from August 2024 to May 2026 under the supervision of Dr. Rohit Beniwal, Assistant Professor, Department of Computer Science and Engineering, Delhi Technological University.

The matter presented in this thesis has not been submitted by me for the award of any other degree of this or any other Institute.

Candidate's SignatureGautam Aggarwal
(24/AFI/02)

This is to certify that the student has incorporated all the corrections suggested by the examiners in the thesis and the statement made by the candidate is correct to the best of our knowledge.

Signature of Supervisor(s)**Signature of External Examiner**

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-110042, India

CERTIFICATE BY THE SUPERVISOR

Certified that Gautam Aggarwal (Roll No. 24/AFI/02) has carried out the research work presented in this thesis entitled "Readmission Risk Prediction using Explainable AI based on Multimodal EHR Data" for the award of Master of Technology from the Department of Computer Science and Engineering, Delhi Technological University, Delhi, under my supervision. The thesis embodies the results of original work, and studies carried out by the student himself, and the contents of the thesis do not form the basis for the award of any other degree to the candidate or to anybody else from this or any other University/Institution.

Signature

Dr. Rohit Beniwal

Assistant Professor

Department of Computer Science and Engineering

Delhi Technological University, Delhi

Place: Delhi

Date: May, 2026

6

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my supervisor, Dr. Rohit Beniwal, Assistant Professor, Department of Computer Science and Engineering, Delhi Technological University, for his invaluable guidance, constant encouragement, and constructive criticism throughout this research work. His insightful suggestions and unwavering support have been instrumental in shaping this thesis.

I am deeply thankful to the Head of the Department of Computer Science and Engineering, the faculty members, and the technical staff of Delhi Technological University for providing the necessary academic environment and infrastructure to carry out this work. I also acknowledge the developers and curators of the MIMIC-IV dataset for making this research possible.

Finally, I extend my heartfelt thanks to my family and friends for their continuous support, patience, and motivation during the course of my M.Tech program. This thesis would not have been possible without their belief in me.

Gautam Aggarwal
(24/AFI/02)

20

10

18

Readmission Risk Prediction using Explainable AI based on Multimodal EHR Data

Gautam Aggarwal

ABSTRACT

Hospital readmission within thirty days of discharge is a major indicator of healthcare quality. Accurate Readmission Risk Prediction (RRP) is therefore important for patient outcomes, mortality, and cost of care. Modern hospitals routinely capture heterogeneous clinical evidence: unstructured clinical narratives, time-series vital signs, and structured International Classification of Diseases (ICD) diagnosis codes. Prior deep-learning approaches to RRP have largely treated these modalities in isolation, ignored their cross-modal dependencies, or omitted explainability, which yields brittle predictions and limited clinical trust.

This thesis proposes a multimodal deep-learning framework that integrates all three modalities for interpretable 30-day RRP. The framework has five new components. Non-linear Ridge ClinicalBERT (NR-ClinicalBERT) performs Semantic Contextual Analysis (SCA) on clinical notes and ICD diagnosis text; the non-linear ridge function augments ClinicalBERT to capture non-linear relationships while controlling feature variance. A Co-Attention Variational AutoEncoder (Co-AVAE) performs Cross-Model Integration Analysis (CMIA) over the clinical, vital, and ICD modalities; the co-attention mechanism captures local cross-modal correlations that the conventional Variational AutoEncoder cannot. A Logarithmic Dual Sigmoid Fuzzy Logic (LDS-FL) module labels risk using the Charlson Comorbidity Index (CCI) and a National Early Warning Score 2 (NEWS2) based Vital Instability Score; the LDS membership function reduces categorisation bias near decision boundaries. A Wald Conditional Random Field (WCRF) models the irregular temporal patterns in vital signs that conventional CRFs miss. A Hessian Smooth Maximum Bidirectional Gated Recurrent Unit (HSM-BiGRU) produces the final readmission prediction using the Smooth Maximum Unit (SMU) activation and Hessian regularization for stability and non-linear feature learning. Local Interpretable Model-agnostic Explanations (LIME) then provide feature-level explanations to support clinician trust.

The framework was implemented in Python and evaluated on the Medical Information Mart for Intensive Care IV (MIMIC-IV) dataset (version 3.1), with 299,712 patient records split 80:20 for training and testing. The proposed HSM-BiGRU classifier reached a readmission prediction accuracy of 98.67%, Root Mean Squared Error (RMSE) of 0.154, Mean Absolute Error (MAE) of 0.118, Cohen's Kappa of 0.961, and Receiver Operating Characteristic Area Under the Curve (ROC-AUC) of 0.9867, beating the BiGRU, BiLSTM, LSTM, and RNN baselines. Co-AVAE obtained a Mean Squared Error (MSE) of 0.011 against 0.092 for the conventional VAE, and NR-ClinicalBERT reached a Mean Reciprocal Rank (MRR) of 0.93 and Mean Average Precision (MAP) of 0.91, ahead of ClinicalBERT, RoBERTa, DistilBERT, and BERT. LDS-FL achieved a defuzzification time of 235

ms. WCRF obtained a log-likelihood of 1.845 with a Sequence Error Rate (SER) of 0.08.

In practice, the framework supports faster and more accurate decision-making for discharge planning and post-discharge resource allocation. Future work can extend the input set to include laboratory reports, medication histories, radiology findings, wearable-sensor data, and genomic information for more complete patient profiling and earlier intervention.

Keywords: *Readmission Risk Prediction; ClinicalBERT; Deep Learning; Clinical Notes; Vital Signs Monitoring; ICD Codes; Natural Language Processing; Multimodal Fusion; Explainable AI.*

LIST OF PUBLICATIONS

- [1] G. Aggarwal and R. Beniwal, "Hospital Readmission Risk Prediction using Explainable AI Based on Multimodal EHR Data," *4th International Conference on Artificial Intelligence and Applications (ICAIA 2026)*, Maharaja Surajmal Institute of Technology, Delhi, 19–20 November 2026. [Communicated]

TABLE OF CONTENTS

Candidate's Declaration	ii
Certificate by the Supervisor	iii
Acknowledgements	iv
Abstract	v
List of Publications	vii
Table of Contents	viii
List of Tables	x
List of Figures	xi
List of Symbols, Abbreviations and Nomenclature	xii
CHAPTER 1 INTRODUCTION	1
1.1 General	1
1.2 Hospital Readmission and its Clinical Significance	1
1.3 Multimodal Clinical Data Sources	2
1.3.1 Clinical Notes	2
1.3.2 Vital Signs	2
1.3.3 ICD Diagnosis Codes	2
1.4 Deep Learning in Clinical Decision Support	3
1.5 Research Gap	3
1.6 Problem Statement	3
1.7 Significance and Scope of the Study	4
1.8 Aim and Key Objectives	4
1.9 Contributions of the Thesis	4
1.10 Organisation of the Thesis	5
CHAPTER 2 LITERATURE SURVEY	6
2.1 Overview	6
2.2 Traditional Machine-Learning Approaches	6
2.3 Deep-Learning Approaches	6
2.4 NLP-Based Approaches for Clinical Notes	6
2.5 Multimodal Approaches	7
2.6 Detailed Review of Recent Works	7
2.7 Comparative Summary of Related Works	8
2.8 Identified Research Gaps	9

2.9 Chapter Summary	10
CHAPTER 3 PROPOSED METHODOLOGY	11
3.1 Overview of the Proposed Framework	11
3.2 Clinical Notes, Vital Signs, and ICD Codes	12
3.3 Pre-Processing and Clinical Keyword Extraction	12
3.4 Feature Extraction	12
3.5 Semantic Contextual Analysis using NR-ClinicalBERT	12
3.6 Cross-Model Integration Analysis using Co-AVAE	13
3.7 Charlson Comorbidity Index and Vital Instability Score	14
3.8 Readmission Risk Labelling using LDS-FL	14
3.9 Sequential Pattern Analysis using WCRF	15
3.10 Readmission Risk Prediction using HSM-BiGRU	15
3.11 Explainability using LIME	17
3.12 Chapter Summary	17
CHAPTER 4 RESULTS AND DISCUSSION	18
4.1 Experimental Setup	18
4.2 Dataset Description	18
4.3 Evaluation Metrics	18
4.4 Performance of Co-AVAE	18
4.5 Performance of HSM-BiGRU	19
4.6 Performance of LDS-FL	21
4.7 Performance of NR-ClinicalBERT	22
4.8 Performance of WCRF	23
4.9 Comparison with Related Works	24
4.10 Discussion	24
4.11 Limitations and Threats to Validity	25
CHAPTER 5 CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT	27
5.1 Conclusion	27
5.2 Practical Implications	27
5.3 Social Impact	28
5.4 Future Scope	28
REFERENCES	29

Plagiarism Report

33

LIST OF TABLES

Table 2.1	Comparative Summary of Recent RRP Works	9
Table 4.1	Comparative Analysis of Co-AVAE	19
Table 4.2	Comparative Analysis of HSM-BiGRU	20
Table 4.3	Comparative Analysis of LDS-FL (in milliseconds)	21
Table 4.4	Comparative Analysis of NR-ClinicalBERT	22
Table 4.5	Comparative Analysis of WCRF	23
Table 4.6	Related Works Comparison regarding RRP	24

LIST OF FIGURES

Figure 3.1	Structure of the Proposed Framework	11
Figure 3.2	Proposed HSM-BiGRU Model	16
Figure 4.1	Graphical Comparison regarding CMIA	18
Figure 4.2	ROC-AUC Analysis of HSM-BiGRU	20
Figure 4.3	Graphical Comparison of LDS-FL	21
Figure 4.4	Comparative Analysis of NR-ClinicalBERT	22
Figure 4.5	Graphical Comparison of WCRF	23

LIST OF SYMBOLS, ABBREVIATIONS AND NOMENCLATURE

Abbreviations

ANFIS	Adaptive Neuro-Fuzzy Inference System
BERT	Bidirectional Encoder Representations from Transformers
BiGRU	Bidirectional Gated Recurrent Unit
BiLSTM	Bidirectional Long Short-Term Memory
CAE	Convolutional AutoEncoder
CCI	Charlson Comorbidity Index
CGAN	Conditional Generative Adversarial Network
ClinicalBERT Transformers	Clinical Bidirectional Encoder Representations from
CMIA	Cross-Model Integration Analysis
CNN	Convolutional Neural Network
Co-A	Co-Attention
Co-AVAE	Co-Attention Variational AutoEncoder
CRF	Conditional Random Field
DAE	Denoising AutoEncoder
DistilBERT Transformers	Distilled Bidirectional Encoder Representations from
DL	Deep Learning
EHR	Electronic Health Record
E-SurP	Environment-aware Survival Predictor
FL	Fuzzy Logic
GLMM	Generalised Linear Mixed Model
GNN	Graph Neural Network
GRU	Gated Recurrent Unit
HMM	Hidden Markov Model
HSM-BiGRU Unit	Hessian Smooth Maximum Bidirectional Gated Recurrent
ICD	International Classification of Diseases
ICU	Intensive Care Unit
LDS-FL	Logarithmic Dual Sigmoid Fuzzy Logic
LIME	Local Interpretable Model-agnostic Explanations
LR	Logistic Regression
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MAP	Mean Average Precision

MCC	Matthews Correlation Coefficient
MIMIC-IV	Medical Information Mart for Intensive Care, version IV
MMD	Maximum Mean Discrepancy
MRR	Mean Reciprocal Rank
MSE	Mean Squared Error
MVI	Missing Value Imputation
NEWS2	National Early Warning Score 2
NLP	Natural Language Processing
NR	Non-linear Ridge
NR-ClinicalBERT	Non-linear Ridge Clinical BERT
PSO	Particle Swarm Optimisation
RBP	Rule-Based Prediction
RMSE	Root Mean Squared Error
RNN	Recurrent Neural Network
RoBERTa	Robust BERT Approach
ROC-AUC	Receiver Operating Characteristic — Area Under the Curve
RRP	Readmission Risk Prediction
SAE	Sparse AutoEncoder
SCA	Semantic Contextual Analysis
SER	Sequence Error Rate
SHAP	SHapley Additive exPlanations
SMU	Smooth Maximum Unit
TFL	Trapezoidal Fuzzy Logic
VAE	Variational AutoEncoder
VOMM	Variable Order Markov Model
WCRF	Wald Conditional Random Field
XGBoost	eXtreme Gradient Boosting

Principal Symbols

M	Aggregate clinical data (notes, vitals, ICD)
P^{clinical}	Pre-processed clinical notes
P^{vital}	Pre-processed vital-sign measurements
F^{clinical}	Extracted clinical features
F^{vital}	Extracted vital-sign features
F^{ICD}	Extracted ICD code features
S^{clinical}	Semantic contextual output of clinical notes
S^{ICD}	Semantic contextual output of ICD text

C	Cross-model integration output
I	Charlson Comorbidity Index
V	Vital instability score
L	Readmission risk label (low / moderate / high)
T	Sequential pattern of vital signs
R	Final readmission risk prediction
A	LIME-based explainability output

CHAPTER 1

INTRODUCTION

1.1 General

Hospital readmission within thirty days of discharge is one of the most consequential indicators of healthcare quality, post-discharge care effectiveness, and inpatient operational efficiency. Unplanned readmissions reflect gaps in initial treatment, discharge planning, transitional care, or post-acute support. They have been consistently associated with higher patient mortality, prolonged morbidity, greater resource use, and financial burdens on both providers and payers [27], [44]. Accurate, automated, and explainable Readmission Risk Prediction (RRP) is therefore central to the design of clinical decision support systems that can flag at-risk patients early and guide proactive interventions [15], [39].

Modern hospitals generate large volumes of heterogeneous patient data during a single admission. The data span unstructured free-text clinical narratives written by clinicians, physiological time series such as heart rate, oxygen saturation and blood pressure, and structured diagnostic codes such as the International Classification of Diseases (ICD) [7], [30]. The three modalities encode different facets of the patient state. Clinical notes carry semantic context. Vital signs reveal physiological trajectories. ICD codes give diagnostic and comorbidity information. A useful RRP system must reason across all three, capturing both the individual signals and the cross-modal dependencies that emerge when they are considered together.

Despite substantial progress in applying machine learning and deep learning to clinical prediction, most existing RRP approaches focus on a single modality, rely on shallow feature engineering, or treat the modalities independently without modelling their interaction. The absence of model transparency in many deep-learning approaches also limits clinical deployment, since clinicians need interpretable evidence to trust an algorithmic recommendation. This thesis addresses both issues with a multimodal deep-learning framework that integrates clinical notes, vital signs and ICD codes, captures semantic and temporal dependencies, and exposes feature-level explanations for each prediction.

1.2 Hospital Readmission and its Clinical Significance

Unplanned readmissions are treated as adverse events in modern healthcare. Beyond the direct clinical risk to the patient, they impose measurable economic costs on health systems and frequently signal systemic issues such as premature discharge, weak patient education, fragmented post-discharge care, or unrecognised comorbidities at the time of discharge. In several jurisdictions, payer organisations have introduced financial penalties for hospitals with excessive 30-day readmission rates for specific conditions such as heart failure, chronic obstructive pulmonary disease (COPD), pneumonia and myocardial infarction [19]. Early identification of

patients at high risk of readmission has therefore become both a clinical and an economic priority.

Conventional risk-stratification tools used in clinical practice, such as the LACE index [40], the HOSPITAL score [13] and various comorbidity-based instruments, rely on a small set of structured variables and rule-based scoring. They are interpretable and easy to deploy, but their predictive performance is modest, and they cannot exploit the richness of contemporary Electronic Health Records (EHRs). Data-driven approaches that use the full breadth of routinely captured clinical information offer substantially higher accuracy, but they raise new questions about representation learning, modality fusion, temporal modelling, and explainability.

1.3 Multimodal Clinical Data Sources

The proposed framework uses three modalities of clinical evidence that are routinely captured during an inpatient episode. Each modality conveys complementary information about the patient, and the combination supports more accurate risk stratification than any single source.

1.3.1 Clinical Notes

Clinical notes are free-text narratives written by physicians, nurses and allied health professionals throughout an admission. They include admission notes, progress notes, nursing shift reports, discharge summaries and consult notes. Unlike structured fields, these narratives capture context, reasoning, uncertainty and nuance. They document symptoms reported by the patient, examination findings, differential diagnoses considered and discarded, treatment decisions, and responses to therapy. No other modality carries comparable semantic depth, but extracting an actionable signal is hard because of abbreviations, domain-specific jargon, negation, temporality and variation in clinician writing style.

1.3.2 Vital Signs

Vital signs are physiological measurements recorded at regular intervals during a patient's stay. They include respiration rate, oxygen saturation, body temperature, systolic and diastolic blood pressure, heart rate and the consciousness level. Vital signs evolve over time, and it is the trajectory, rather than any individual reading, that usually reveals clinical deterioration or recovery. Early-warning systems such as the National Early Warning Score 2 (NEWS2) [35] aggregate abnormal readings into a single instability score that correlates with adverse events, including unplanned readmission. Vital signs are inherently temporal and irregular: measurements occur at non-uniform intervals, may be missing, and contain noise from monitoring artefacts.

1.3.3 ICD Diagnosis Codes

The International Classification of Diseases (ICD) is the standardised coding system for diagnoses and procedures in administrative and clinical databases. ICD codes are hierarchical and capture the principal diagnosis, secondary diagnoses, comorbidities and complications. They are structured and machine-readable and form the basis for established comorbidity instruments such as the Charlson

Comorbidity Index (CCI) [9], which assigns weighted scores to specific conditions and is a strong predictor of one-year mortality and other adverse outcomes. ICD codes also carry semantic descriptions (short titles and long titles) that can be embedded as text to capture diagnostic context beyond the bare code.

1.4 Deep Learning in Clinical Decision Support

Over the past decade, deep learning has changed how clinical data are analysed. Recurrent Neural Networks (RNNs), in particular Long Short-Term Memory (LSTM) [16] and Gated Recurrent Unit (GRU) [11] variants, are widely used for temporal sequences such as vital signs and laboratory results. Convolutional Neural Networks (CNNs) have been applied to medical imaging and to fixed-window summaries of EHR data. Transformer-based language models such as BERT [12] and its clinical adaptation ClinicalBERT [17] have substantially improved the representation of unstructured clinical notes by capturing bidirectional context across long sequences. Variational AutoEncoders (VAEs) [21] and related generative models give a principled way to learn compact latent representations of heterogeneous data, while attention [41] and co-attention [28] let models focus on the most informative regions of an input.

Two issues limit the clinical deployment of deep models. First, when multiple modalities are involved, naive concatenation or late fusion fails to capture the local interactions between them. Second, deep classifiers are typically opaque: they return a score without indicating which features drove the prediction. Local Interpretable Model-agnostic Explanations (LIME) [34] and SHapley Additive exPlanations (SHAP) [29] provide post-hoc feature attributions that address the second issue, but they are rarely combined with multimodal architectures designed specifically for clinical RRP.

1.5 Research Gap

A systematic review of the most recent literature on hospital RRP shows that none of the prevailing works model cross-modal integration between clinical notes, vital signs and ICD codes [5]. As a result, the relationships across the three modalities are not captured, and the reliability of RRP suffers. The temporal evolution of vital signs is also commonly modelled with uniform-interval assumptions that do not hold in real ICU and ward settings, the contextual meaning of clinical narratives is often reduced to bag-of-words representations that discard semantic dependencies [24], and explainability is treated as an after-thought rather than a design principle.

1.6 Problem Statement

The main limitations of prevailing RRP approaches that this thesis addresses are:

- Conventional works do not adequately capture the contextual information in clinical notes, which reduces the overall performance of the model.

- Existing works often overlook the irregular temporal variations and evolving physiological patterns of vital signs, which lowers predictive performance for time-dependent risks.
- Most prevailing works predict readmission risk without exposing the factors behind the decision, which decreases clinician trust and limits real-world adoption.
- Cross-modal interactions between clinical notes, vital signs and ICD codes are typically ignored, even though they carry significant predictive signal.

1.7 Significance and Scope of the Study

This work develops a single framework for hospital RRP that uses heterogeneous clinical data, captures cross-modal dependencies, models irregular temporal patterns, and explains its predictions. The framework is intended to support real-time clinical decision-making and can be integrated into discharge planning, transitional care coordination, and post-acute resource allocation.

1.8 Aim and Key Objectives

The proposed system designs a multimodal deep-learning framework for predicting readmission risk using clinical notes, vital signs and ICD codes. The objectives that drive the design are:

- To improve the learning capability of the model by performing Cross-Model Integration Analysis (CMIA) over clinical notes, vital signs and ICD-based disease representations using the proposed Co-Attention Variational AutoEncoder (Co-AVAE).
- To improve the semantic understanding of clinical narratives by performing Semantic Contextual Analysis (SCA) using the proposed Non-linear Ridge ClinicalBERT (NR-ClinicalBERT).
- To analyse temporal variations in vital signs by extracting sequential patterns using the proposed Wald Conditional Random Field (WCRF).
- To label readmission risk into low, moderate and high categories using the proposed Logarithmic Dual Sigmoid Fuzzy Logic (LDS-FL) approach driven by the Charlson Comorbidity Index (CCI) and a NEWS2-based Vital Instability Score.
- To predict the final readmission risk using the proposed Hessian Smooth Maximum Bidirectional Gated Recurrent Unit (HSM-BiGRU) classifier.
- To improve transparency and clinical trust by explaining the factors that influence the RRP using Local Interpretable Model-agnostic Explanations (LIME).

1.9 Contributions of the Thesis

The principal contributions of this thesis are summarised as follows:

- An NR-ClinicalBERT module is introduced for semantic contextual analysis of clinical notes and ICD diagnosis text. Augmenting ClinicalBERT with a Non-linear Ridge function lets the module capture non-linear semantic relationships while controlling excessive feature variance through ridge regularization.
- A Co-AVAE module is introduced for cross-model integration analysis across the three modalities. The co-attention mechanism captures local feature correlations between modalities, which the conventional VAE misses, and the variational encoder/decoder structure produces a single latent representation.
- An LDS-FL module is introduced for readmission risk labelling. The Logarithmic Dual Sigmoid membership function reduces the categorisation bias of conventional fuzzy logic membership functions and gives faster fuzzification and defuzzification.
- A WCRF module is introduced for sequential pattern analysis of vital signs. The Wald distribution captures stochastic temporal intervals that the conventional CRF assumes to be uniform.
- An HSM-BiGRU classifier is introduced for the final RRP. The Smooth Maximum Unit (SMU) activation smooths non-linear learning. Hessian regularization captures second-order information of the loss function and improves the stability of the BiGRU parameters.
- Empirical evaluation on the MIMIC-IV (v3.1) dataset, comparing each proposed module against four to five backbone baselines across five sub-tasks (CMIA, SCA, risk labelling, sequential pattern analysis and final RRP), together with a qualitative comparison to six recently published works.
- Integration of LIME-based explainability that gives feature-level explanations of each prediction for clinical trust and auditability.

1.10 Organisation of the Thesis

The remainder of this thesis is organised as follows:

- Chapter 2 reviews the literature on hospital readmission risk prediction, grouping existing works by methodology and modality and identifying the specific gaps that motivate the present work.
- Chapter 3 presents the proposed methodology in detail and describes each of the five novel modules (NR-ClinicalBERT, Co-AVAE, LDS-FL, WCRF and HSM-BiGRU) together with the LIME-based explainability stage and the underlying mathematical formulations.
- Chapter 4 reports the experimental setup, dataset characteristics, evaluation metrics and comparative results across all five sub-tasks, with a discussion of what the findings imply and how they compare to recently published works.
- Chapter 5 concludes the thesis. It summarises the main findings, sets out the practical and social implications of the work, and outlines concrete directions for future research.

CHAPTER 2

LITERATURE SURVEY

2.1 Overview

This chapter reviews the work most directly related to hospital readmission risk prediction (RRP). The literature is large and varied. It covers traditional machine-learning approaches on structured features, deep-learning methods on sequences or unstructured text, natural-language processing pipelines specialised for clinical narratives, and multimodal architectures that try to fuse two or more data sources. The review is organised first by methodological family and then by recent works directly comparable to the proposed framework. The chapter ends with the specific limitations of these works that the present thesis is designed to overcome.

2.2 Traditional Machine-Learning Approaches

Classical machine-learning models are still widely applied to RRP because they are interpretable, need modest data, and are easy to deploy. Logistic regression remains a strong baseline because its coefficients can be inspected directly. Random forests [8] and gradient-boosted decision trees, especially eXtreme Gradient Boosting (XGBoost) [10], have frequently beaten logistic regression on structured EHR features. These models, though, depend on hand-crafted feature engineering: comorbidity indices, admission counts, length-of-stay summaries, demographic variables and aggregated laboratory values must all be designed by domain experts. They also struggle with unstructured free text and with irregularly sampled physiological time series.

2.3 Deep-Learning Approaches

Deep-learning models have extended the scope of RRP by enabling end-to-end learning from raw or minimally pre-processed data. RNN variants such as LSTM and GRU model temporal dependencies in vital signs and time-stamped events. Bidirectional variants (BiLSTM, BiGRU) [37] further improve performance by processing each sequence in both directions, so that every time step has access to past and future context. CNNs have been applied to short, fixed-length windows of structured EHR data treated as one-dimensional feature maps. More recently, graph neural networks (GNNs) [22] have been proposed for patient-similarity modelling and for fusing structured and unstructured information through heterogeneous graphs.

Deep classifiers are flexible but often struggle with highly non-linear relationships and show training instability on noisy clinical data. They are also prone to overfitting, particularly when the readmission outcome is imbalanced. Regularization strategies such as dropout, weight decay and, more recently, second-order Hessian-based methods have been explored to address these issues.

2.4 NLP-Based Approaches for Clinical Notes

Natural-language processing of clinical notes is hard for specific reasons: domain-specific vocabulary, heavy use of abbreviations, ambiguous negation, and a shortage of large annotated corpora. Pre-trained transformer language models have largely replaced earlier word-embedding approaches. BERT [12] introduced bidirectional self-attention over sub-word units; RoBERTa [25] improved training stability and dataset scale; DistilBERT [36] offered a smaller, faster variant; and ClinicalBERT [17], pre-trained on clinical narratives from the MIMIC corpus, specialised these advances to the medical domain. These models capture contextual semantics well, but in their default form they cannot easily express non-linear relationships among learnt features, and they are sensitive to feature variance from highly heterogeneous clinical text.

2.5 Multimodal Approaches

Multimodal models combine evidence from several data sources, typically by encoding each modality separately and then fusing the resulting representations. Early-fusion methods concatenate raw features before learning. Late-fusion methods train modality-specific classifiers and combine their outputs. Mid-level fusion learns a joint latent representation. Variational AutoEncoders have been used to learn compact joint latent spaces for heterogeneous clinical inputs, but the conventional VAE objective tends to ignore local feature variations across modalities, and the symmetric Kullback-Leibler objective treats all modalities equally regardless of their informativeness. Co-attention mechanisms, which compute attention weights jointly across two or more streams, capture fine-grained cross-modal interactions well in non-clinical settings, but they have rarely been used in clinical RRP frameworks.

2.6 Detailed Review of Recent Works

This section reviews five recent works most directly comparable to the proposed framework. Each entry summarises the data, methodology, main contribution, and the specific limitation that motivates the corresponding component in the proposed framework.

2.6.1 Almeida, Moreno and Barata (2025) — *Clinical Notes with Graph Neural Networks*

Almeida et al. [4] presented a hospital readmission prediction pipeline based on clinical notes. Patient clinical notes were first collected from the EHR. The text was pre-processed and then embedded using transformer encoders. A Graph Neural Network (GNN) was applied to capture relational dependencies between encoded sentences and to predict 30-day readmission. The work showed that contextual embeddings of clinical notes, paired with relational learning, can beat classical baselines on text-only RRP. The model never observed any physiological information, so the deterioration signals carried by vital signs were missed, which limits its overall predictive ceiling.

2.6.2 Pham et al. (2024) — *CNN with SHAP-Based Explainability*

Pham et al. [32] built a forecasting pipeline for early patient readmission using Irish hospital discharge records. Medical data were collected and balanced to address the imbalance between readmitted and non-readmitted classes. A Convolutional Neural Network (CNN) was used to model fixed-window summaries of structured features, and SHAP was applied to add interpretability. The work showed that conventional machine-learning models combined with explanation tools can produce clinically useful predictions. The unstructured clinical narratives were not analysed contextually, which reduces the model's reliability on patients whose risk factors live mainly in free-text fields.

2.6.3 Woudneh (2025) — *Generalised Linear Mixed Models for Elderly Patients*

Woudneh [42] looked at hospital readmission among elderly patients and combined patient-level and ward-level factors. Hospital health records and surgical logs were collected, and data quality was checked carefully. A Generalised Linear Mixed Model (GLMM) was used to model hierarchical effects across wards and patients. The mixed-model formulation accounted for between-ward variability and produced statistically interpretable coefficients. The model had no explicit explainability layer beyond classical coefficient inspection and did not use unstructured information, which limits both trust and predictive depth.

2.6.4 Song et al. (2025) — *XGBoost with Stepwise Feature Selection*

Song et al. [38] addressed unplanned readmission of elderly patients with coronary heart disease. Clinical data were collected and a stepwise forward feature-selection method was used to identify the most predictive variables. XGBoost was then trained as the classifier, with SHAP-based explanations alongside the predictions. The stepwise approach produced an interpretable and compact feature set. Cross-modal integration between heterogeneous clinical sources was not analysed, and the structure of unstructured narratives and physiological time series was not used.

2.6.5 Liu et al. (2024) — *Random Forest and XGBoost for Post-PCI Patients*

Liu et al. [26] explored readmission risk prediction for patients with non-ST elevation myocardial infarction after percutaneous coronary intervention. Patients were first stratified by infarction type. Clinical variables common across the cohort were selected with a random forest, and XGBoost was then trained to predict 30-day readmission. The ensemble approach reached competitive accuracy on the targeted cohort. The framework ignored vital signs and ICD codes (two of the three modalities considered in the present work), which limits its applicability to broader patient populations.

2.7 Comparative Summary of Related Works

Table 2.1 summarises the most relevant recent works on RRP, including those reviewed individually above and five additional works [1], [2], [18], [33], [43] that appear in the comparative study of Chapter 4. For each work, the table lists the main methodology, the modalities considered, the main advantage, and the limitation that the proposed framework addresses.

Table 2.1: Comparative Summary of Recent RRP Works

Study	Methods	Modalities	Limitation Addressed Here
Almeida et al. (2025) [4]	GNN on clinical notes	Notes only	Vital and ICD information ignored
Pham et al. (2024) [32]	CNN + SHAP	Structured EHR	No semantic contextual analysis
Woudneh (2025) [42]	GLMM	Patient and ward factors	Lacks explicit explainability layer
Song et al. (2025) [38]	XGBoost + stepwise selection	Structured EHR	No cross-model integration
Liu et al. (2024) [26]	Random Forest + XGBoost	Cardiac sub-cohort	Ignores vital signs and ICD codes
Rasheed et al. (2025) [33]	LSTM	ICU time series	Prediction model not stable
Huang et al. (2025) [18]	CGAN + XGBoost	AECOPD cohort	No cross-modality relationships
Albahli (2025) [2]	Random Forest	Diabetic patients	Cannot adapt to real-time monitoring
Alabdulhafith et al. (2023) [1]	PSO + Ensemble ML + XAI	ICU edge/cloud	Generalisability decreased
Zhai et al. (2025) [43]	E-SurP	Cardiac + environment	Non-linear factors hard to model

2.8 Identified Research Gaps

The review above yields four research gaps that the present work directly targets:

- Cross-modal interactions between clinical notes, vital signs and ICD codes are not explicitly modelled in any of the reviewed works, even though all three modalities are routinely captured during inpatient care and carry highly complementary signal.
- Semantic context in clinical narratives is either ignored or handled with bag-of-words and shallow embeddings, which fails to capture the long-range, bidirectional and domain-specific dependencies that ClinicalBERT-class models can express.
- Temporal patterns in vital signs are typically modelled with assumptions of uniform sampling intervals or first-order Markov dynamics, which do not reflect the stochastic and irregular nature of ward-recorded physiology.

- Explainability is treated as an after-thought. Only a minority of works incorporate post-hoc attribution methods such as SHAP or LIME, and none combines explainability with an integrated multimodal architecture.

2.9 Chapter Summary

This chapter has surveyed the recent work on hospital RRP, covering traditional machine learning, deep learning, NLP-based and multimodal approaches. Five directly comparable works were reviewed in depth and a further five were summarised in comparative form. Four specific research gaps were identified: cross-modal integration, semantic contextual analysis, irregular temporal modelling, and integrated explainability. The framework proposed in Chapter 3 is designed to address all four.

CHAPTER 3

PROPOSED METHODOLOGY

3.1 Overview of the Proposed Framework

This chapter presents the proposed Co-AVAE-based Cross-Model Integration Analysis (CMIA) and HSM-BiGRU-based Readmission Risk Prediction (RRP) framework. The framework takes three modalities of clinical evidence (unstructured clinical notes, time-series vital signs and structured ICD diagnosis codes) and produces an interpretable readmission risk prediction together with a feature-level explanation. The end-to-end architecture is summarised in Figure 3.1.




Figure 1: Structure of the Proposed Framework

Figure 3.1: Structure of the Proposed Framework

The framework has ten stages: (i) collection of the three input modalities; (ii) modality-specific pre-processing and clinical keyword extraction; (iii) feature extraction; (iv) semantic contextual analysis of clinical notes and ICD text using NR-ClinicalBERT; (v) cross-model integration of all three modalities using Co-AVAE; (vi) computation of the Charlson Comorbidity Index and the Vital Instability Score; (vii) readmission-risk labelling using LDS-FL; (viii) sequential pattern analysis of vital signs using WCRF; (ix) final readmission-risk prediction using HSM-BiGRU;

and (x) feature-level explanation of the prediction using LIME. Each stage is described in detail in the subsections that follow.

3.2 Clinical Notes, Vital Signs, and ICD Codes

The proposed framework starts with the collection of clinical data M , comprising clinical notes M^{clinical} , vital signs M^{vital} and ICD-based diagnosis information M^{ICD} .

M therefore carries the semantic, physiological and diagnostic information that the downstream stages need.

3.3 Pre-Processing and Clinical Keyword Extraction

Clinical notes and vital signs are pre-processed before feature extraction. For clinical notes, sentence splitting is first performed on M^{clinical} to break the narrative into smaller semantic units. Abbreviations are then extracted to disambiguate domain-specific short forms. Stop-word removal eliminates irrelevant tokens such as "is", "as" and "the". Finally, lemmatisation converts words to their root form, yielding the pre-processed clinical data P^{clinical} .

Meanwhile, Missing Value Imputation (MVI) is performed on M^{vital} to handle gaps in the physiological time series, yielding the pre-processed vital signs P^{vital} . From P^{clinical} , the set of clinical keywords χ — such as symptoms, diseases, treatment procedures and medications — is extracted.

3.4 Feature Extraction

Three feature sets are then extracted, one per modality. From P^{clinical} , the clinical features F^{clinical} (word count, sentence count, average sentence length, unigrams, bigrams and similar surface statistics) are extracted. From P^{vital} , the vital features F^{vital} are extracted: respiration rate, oxygen saturation, body temperature, systolic blood pressure, heart rate and consciousness level. From M^{ICD} , the ICD features F^{ICD} are extracted: the ICD code itself, the ICD version, the short title and the long title.

3.5 Semantic Contextual Analysis using NR-ClinicalBERT

Semantic Contextual Analysis (SCA) is performed over χ and P^{clinical} using the proposed Non-linear Ridge ClinicalBERT (NR-ClinicalBERT). ClinicalBERT [17], a transformer pre-trained on clinical narratives, is the backbone here because it captures the bidirectional semantic meaning of medical terms. The standard ClinicalBERT, however, struggles to express non-linear relationships between its learnt features. A Non-linear Ridge (NR) function is therefore added for feature learning inside ClinicalBERT.

First, the clinically embedded data E is obtained with respect to the embedding matrix δ :

$$E = \chi \times P^{\text{clinical}} \times \delta \quad (3.1)$$

The contextual relationship between the tokens of E is then learnt using a soft-attention mechanism g :

$$g = \text{softmax}[E \times \frac{u \times w^\top}{\sqrt{e}} \times v] \quad (3.2)$$

where u , v and w denote the query, value and key matrices of E , respectively; e is the dimensional scaling factor; and the $\top$ super-script denotes the transpose. In parallel, the NR function ϕ , which extracts non-linear patterns while minimising excessive feature variance, is computed as:

$$\phi = \exp[-\frac{(i^\top E - n)^2}{\sigma^2 + \gamma \|E\|^2}] \quad (3.3)$$

where i and n are the weight and bias of E , σ^2 is the variance-control parameter and γ is the ridge regularization coefficient. The clinical SCA output S^{clinical} is then obtained by combining the attention output g and the non-linear ridge response ϕ :

$$S^{\text{clinical}} = [(g, \phi) \times i] + n \quad (3.4)$$

Likewise, the SCA output S^{ICD} for the ICD diagnosis text is computed from F^{ICD} using the same NR-ClinicalBERT pipeline. Both S^{clinical} and S^{ICD} are forwarded to the cross-model integration stage.

3.6 Cross-Model Integration Analysis using Co-AVAE

The cross-model integration analysis (CMIA) over F^{clinical} , F^{vital} , F^{ICD} , S^{clinical} and S^{ICD} is carried out using the proposed Co-Attention Variational AutoEncoder (Co-AVAE). The Variational AutoEncoder (VAE) [21] is the backbone because it learns compact latent representations from heterogeneous data. The conventional VAE, however, ignores local feature variations across modalities. A Co-Attention (Co-A) mechanism [28] is therefore introduced to capture local cross-modal correlations before the variational encoder.

Let the inputs for CMIA be denoted by G , comprising F^{clinical} , F^{vital} , F^{ICD} , S^{clinical} and S^{ICD} .

Let G^1 and G^2 denote two modalities drawn from G . The Co-A mechanism η , which learns the correlated representation between the inputs, is then evaluated as:

$$\eta = \{\text{softmax}(r \times s^\top)s, \text{softmax}(r^\top \times s)r\} \quad (3.5)$$

where r and s denote the query value of G^1 and the key value of G^2 , respectively. Next, the Co-A-based multimodal feature representation κ is computed with respect to η and the value-representation m of G :

$$\kappa = \sum \eta \times m \quad (3.6)$$

The variational encoder then maps κ into a latent probabilistic space parameterised by a multimodal embedding vector J and the mean p and standard deviation q of G .

The decoder reconstructs the multimodal representation, yielding the decoder output.

Finally, multimodal interaction learning is carried out over the encoder and decoder outputs, yielding the CMIA output C :

$$C = G[\log Q^{encoder} - \lambda Q^{decoder}] \quad (3.7)$$

where λ is a weighting coefficient that balances the encoder and decoder contributions. Fine-grained cross-modal interaction between the heterogeneous clinical inputs is therefore captured in a single latent representation. The pseudo-code of the Co-AVAE module is shown below.

Algorithm 3.1: Co-AVAE

Input: $F^{clinical}$, F^{vital} , F^{ICD} , $S^{clinical}$, S^{ICD}

Output: CMIA output C

Begin

Let $G = [F^{clinical}, F^{vital}, F^{ICD}, S^{clinical}, S^{ICD}]$

For G do

Evaluate Co-A mechanism:

$\eta = \{ \text{softmax}[(r \times s^T) s], \text{softmax}[(r^T \times s) r] \}$

Determine multimodal feature: $\kappa = \sum \eta \times m$

Encode input: $Q^{encoder} = \kappa \times [J; (p, q) G]$

Decode features: $Q^{decoder} = Q^{encoder} \times [\kappa, J]$

Integrate cross-models: $C = G[\log \times Q^{encoder} -$

$\lambda(Q^{decoder})]$

End For

Return C

End

3.7 Charlson Comorbidity Index and Vital Instability Score

Two clinically motivated scores are computed from the modality-specific feature sets to support readmission-risk labelling. From F^{ICD} , the Charlson Comorbidity Index (CCI), denoted I , is calculated as:

$$I = \sum d(F^{ICD}) \times \mu \quad (3.8)$$

where d and μ denote the weight value and the comorbidity indicator of F^{ICD} , respectively. From F^{vital} , the Vital Instability Score V is calculated based on the National Early Warning Score 2 (NEWS2):

$$V = \sum F^{vital} \times \vartheta \quad (3.9)$$

where ϑ denotes the NEWS2-based score assigned to each vital-sign feature in F^{vital} . Together, I and V summarise the patient's comorbidity burden and current physiological stability respectively, and are passed to the labelling module.

3.8 Readmission Risk Labelling using LDS-FL

Given I and V , the readmission risk is labelled into three discrete categories using the proposed Logarithmic Dual Sigmoid Fuzzy Logic (LDS-FL) approach. The backbone is a conventional Fuzzy Logic (FL) system because it can integrate clinical knowledge through rule-based reasoning. The standard membership functions used

in conventional FL (triangular, trapezoidal, Gaussian) introduce categorisation bias near the decision boundaries. The Logarithmic Dual Sigmoid (LDS) membership function replaces them.

First, a set of fuzzy rules B is defined using if-then conditions:

$$B = \begin{cases} L^1, & \text{if } I + V \leq 4 \\ L^2, & \text{if } 5 \leq I + V \leq 8 \\ L^3, & \text{if } I + V \geq 9 \end{cases} \quad (3.10)$$

where L^1 , L^2 and L^3 denote low, moderate and high readmission-risk labels respectively. The LDS membership function y , which provides an accurate representation of uncertainty, is then evaluated as:

$$y = \frac{1}{2} \times \left[\frac{1}{1 + \exp^{-a[(I,V)-c]}} + \frac{\ln(1+b|[(I,V)-c]^2|)}{1 + \ln(1+b|[(I,V)-c]^2|)} \right] \quad (3.11)$$

where a , b and c are scaling parameters. Fuzzification is then performed to convert crisp data into fuzzy data f :

$$f = (V, I) \times B \times y \quad (3.12)$$

followed by defuzzification, which yields the discrete readmission-risk label L :

$$L = \frac{\sum y \cdot f}{\sum y} \quad (3.13)$$

3.9 Sequential Pattern Analysis using WCRF

In parallel with risk labelling, sequential patterns in F^{vital} are extracted using the proposed Wald Conditional Random Field (WCRF). The backbone is the Conditional Random Field (CRF) [23] because it captures relationships between consecutive clinical observations. Conventional CRFs, however, assume regular sampling intervals and cannot represent the irregular temporal variations typical of ward-recorded vital signs. The Wald distribution ς is therefore introduced to learn stochastic temporal intervals:

$$\varsigma = \sqrt{\frac{h}{2\pi U^3}} \exp\left(-\frac{h(U-j)^2}{2j^2 U}\right) \quad (3.14)$$

where h is the shape parameter controlling temporal variation, π is the mathematical constant, U is the time interval between consecutive vital-sign observations and j is the mean temporal interval of F^{vital} . The sequential pattern T is then evaluated with respect to ς :

$$T = \frac{1}{z(F^{\text{vital}})} \times \exp[\sum (N \times F^{\text{vital}} \times U) + \varsigma] \quad (3.15)$$

where z is a normalisation factor and N is the weight value of F^{vital} . WCRF therefore keeps the relational reasoning of CRFs while accommodating the stochastic temporal intervals that arise in real clinical settings.

3.10 Readmission Risk Prediction using HSM-BiGRU

The final readmission-risk prediction is performed using the proposed Hessian Smooth Maximum Bidirectional Gated Recurrent Unit (HSM-BiGRU). The Bidirectional Gated Recurrent Unit (BiGRU) [11], [37] is the backbone classifier

because it can capture disease-progression patterns in clinical sequences. The standard BiGRU, however, struggles with highly non-linear relationships, and its parameters are sensitive to noise during training. The Smooth Maximum Unit (SMU) activation function [6] and Hessian regularization are therefore introduced for feature learning and model stability, respectively. The architecture of HSM-BiGRU is shown in Figure 3.2.




Figure 2: Proposed HSM-BiGRU Model

Figure 3.2: Proposed HSM-BiGRU Model

Let D denote the resultant feature vector that aggregates the CMIA output C , the Charlson Comorbidity Index I , the Vital Instability Score V , the readmission-risk label L and the sequential pattern T . First, Hessian regularization, which captures second-order information of the loss function ℓ , is applied to regularize the model parameters $\mathfrak{F}$. The regularized output ε is given by:

$$\varepsilon = \frac{\partial^2 \ell}{\partial \mathfrak{F}^2} \quad (3.16)$$

Next, for each Gated Recurrent Unit (GRU) in both the forward and backward layers, the reset gate Z^1 and the update gate Z^2 are computed as:

$$Z^1 = \alpha \otimes \{[(D, H^{(k-1)}) \times t] + o\} \quad (3.17)$$

$$Z^2 = \alpha \{[Z^1 \otimes (D, H^{(k-1)})] \times t + o\} \quad (3.18)$$

where α denotes the sigmoid activation function, $H^{(k-1)}$ is the previous hidden state at time step k , t is the weight value and o is the bias of the model. The present hidden state H^k is then estimated based on Z^2 and the SMU activation function β , which smoothens the non-linear learning:

$$H^k = \langle \beta \{[(D, H^{(k-1)}) \times t] + o\} \otimes |\alpha \{[(D, H^{(k-1)}) \times t] + o\} - 1| \rangle \oplus Z^2 \quad (3.19)$$

$$\beta = \frac{D + \sqrt{D^2 + \tau}}{2} \quad (3.20)$$

where τ is the smoothness parameter. Let $H^{\rightarrow k}$ and $H^{\leftarrow k}$ denote the hidden states obtained from the forward and backward processing respectively. The final RRP output R is obtained by concatenating these two hidden states:

$$R = \vec{H}^k \oplus \overleftarrow{H}^k \quad (3.21)$$

where Θ denotes the concatenation operator. The final output R classifies the readmission risk as low, moderate or high. The pseudo-code of HSM-BiGRU is given below.

Algorithm 3.2: HSM-BiGRU

Input: Resultant factor D

Output: RRP output R

Begin

Initialize weight t and bias o

Apply Hessian-based parameter regularization: $\varepsilon = \partial^2 \ell / \partial \mathfrak{S}^2$

While (time = k) do

For each D do

Evaluate reset gate: $Z^1 = \alpha \otimes \{ [(D, H^{(k-1)}) \times t] + o \}$

Determine update gate: $Z^2 = \alpha \times \{ [Z^1 \otimes (D, H^{(k-1)})] \times$

$t + o \}$

Calculate SMU activation: $\beta = (D + \sqrt{D^2 + \tau}) / 2$

Estimate present hidden state H^k

Predict readmission risk: $R = H \rightarrow^k \Theta H \leftarrow^k$

End For

End While

Return R

End

3.11 Explainability using LIME

The explainability of the features that influence R is then evaluated using Local Interpretable Model-agnostic Explanations (LIME) [34]. LIME approximates the local decision boundary of the trained HSM-BiGRU classifier with an interpretable surrogate model, which identifies the contribution of each feature to the prediction. The explainability output A is computed as:

$$A = \arg \min \{ R(F^{clinical}, F^{vital}, F^{ICD}) \times (\xi + \wp) \} \quad (3.22)$$

where ξ is the local fidelity loss and $\wp$ is the penalty factor of the explanation model. The resulting explanations let clinicians inspect which clinical-note phrases, vital-sign trends and ICD codes contributed most to a particular readmission prediction, which improves model trust in real-world deployments.

3.12 Chapter Summary

This chapter has detailed the proposed methodology: five novel modules (NR-ClinicalBERT, Co-AVAE, LDS-FL, WCRF and HSM-BiGRU) together with a LIME-based explainability stage. The mathematical formulation of each module has been given, and pseudo-code is included for the two most involved modules (Co-AVAE and HSM-BiGRU). Chapter 4 reports the empirical evaluation of the framework on the MIMIC-IV dataset and compares its performance against established baselines.

22

CHAPTER 4

RESULTS AND DISCUSSION

4.1 Experimental Setup

The proposed framework was implemented in Python using standard deep-learning libraries on a GPU-equipped platform for model training and inference. Each module (NR-ClinicalBERT, Co-AVAE, LDS-FL, WCRF and HSM-BiGRU) was implemented as a self-contained block and integrated through the end-to-end pipeline described in Chapter 3. For reproducibility, an 80:20 train/test split was applied to the dataset, and all comparative experiments used identical splits, evaluation protocols and metric definitions.

4.2 Dataset Description

All experiments were run on the Medical Information Mart for Intensive Care IV (MIMIC-IV) dataset, version 3.1, released through the PhysioNet platform [20]. MIMIC-IV is a large, freely available, de-identified clinical database of patient stays from the Beth Israel Deaconess Medical Center. It covers demographic data, hospital stays, ICU stays, laboratory measurements, charted vital signs, prescriptions, procedures and ICD-coded diagnoses. The total number of records used in this study was 299,712: 239,770 (80%) for training and the remaining 59,942 (20%) for testing. The dataset gives a realistic, large-scale and heterogeneous evaluation environment for multimodal readmission risk prediction.

4.3 Evaluation Metrics

Each module was evaluated using metrics appropriate to its sub-task. Co-AVAE was assessed using Mean Squared Error (MSE), Perceptual Loss, and Maximum Mean Discrepancy (MMD), which measure reconstruction accuracy, perceptual similarity and distributional alignment, respectively. HSM-BiGRU was assessed using classification Accuracy, Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Cohen's Kappa, and the Receiver Operating Characteristic Area Under the Curve (ROC-AUC). LDS-FL was evaluated using the time taken for rule generation, fuzzification and defuzzification. NR-ClinicalBERT was assessed using Mean Reciprocal Rank (MRR) and Mean Average Precision (MAP). WCRF was evaluated using Log-Likelihood, Matthews Correlation Coefficient (MCC) and Sequence Error Rate (SER).

4.4 Performance of Co-AVAE

Table 4.1 reports the cross-model integration performance of the proposed Co-AVAE against four conventional autoencoder baselines: the Variational AutoEncoder (VAE), Sparse AutoEncoder (SAE), Denoising AutoEncoder (DAE)

8

4

and Convolutional AutoEncoder (CAE). The corresponding perceptual loss is plotted in Figure 4.1.

Table 4.1: Comparative Analysis of Co-AVAE

Algorithm	MSE	Perceptual Loss	MMD
Proposed Co-AVAE	0.011	0.128	0.052
VAE	0.092	0.193	0.123
SAE	0.234	0.213	0.475
DAE	0.477	0.536	0.731
CAE	0.796	0.941	1.362




Figure 4.1: Graphical Comparison regarding CMIA

As shown in Table 4.1 and Figure 4.1, the proposed Co-AVAE captured cross-model integration between the clinical modalities with an MSE of 0.011, a perceptual loss of 0.128 and an MMD of 0.052. The conventional VAE, SAE, DAE and CAE all gave higher MSE, perceptual loss and MMD values. The result supports the central design hypothesis: the co-attention mechanism lets Co-AVAE capture the local feature correlations across modalities that the conventional VAE objective misses, while preserving the variational structure needed for a compact latent representation.

4.5 Performance of HSM-BiGRU

Table 4.2 reports the final readmission risk prediction performance of the proposed HSM-BiGRU classifier against four conventional sequence baselines:

BiGRU, BiLSTM, LSTM and RNN. Figure 4.2 plots the corresponding ROC-AUC curves.

Table 4.2: Comparative Analysis of HSM-BiGRU

Techniques	Accuracy (%)	RMSE	MAE	Cohen's Kappa
Proposed HSM-BiGRU	98.67	0.154	0.118	0.961
BiGRU	96.02	0.411	0.213	0.938
BiLSTM	93.38	0.743	0.399	0.899
LSTM	90.12	0.984	0.474	0.842
RNN	88.57	1.284	0.524	0.817




Figure 4.2: ROC-AUC Analysis of HSM-BiGRU

The proposed HSM-BiGRU improved both model stability and feature learning through Hessian regularization and the SMU activation function. As shown in Table 4.2 and Figure 4.2, the classifier predicted readmission risk with an accuracy of 98.67%, an RMSE of 0.154, an MAE of 0.118, a Cohen's Kappa of 0.961 and a ROC-AUC of 0.9867. The conventional BiGRU, BiLSTM, LSTM and RNN baselines gave higher error rates and lower agreement metrics. The bidirectional architecture combined with the Hessian-regularized SMU activation let the model capture the highly non-linear progression patterns that underlie readmission risk and

limited the parameter sensitivity that usually degrades vanilla BiGRUs on noisy clinical data.

4.6 Performance of LDS-FL

Figure 4.3 reports the rule-generation, fuzzification and defuzzification times of the proposed LDS-FL approach in comparison to four conventional fuzzy-logic baselines: Fuzzy Logic (FL), Adaptive Neuro-Fuzzy Inference System (ANFIS), Trapezoidal FL (TFL) and Rule-Based Prediction (RBP).

Table 4.3: Comparative Analysis of LDS-FL (in milliseconds)

Techniques	Rule Generation (ms)	Fuzzification (ms)	Defuzzification (ms)
Proposed LDS-FL	354	242	235
FL	642	475	563
ANFIS	854	614	655
TFL	996	843	784
RBP	1053	929	928




Figure 4.3: Graphical Comparison of LDS-FL

As shown in Figure 4.3 and Table 4.3, the proposed LDS-FL and the conventional FL, ANFIS, TFL and RBP labelled the readmission risk with rule-generation times of 354 ms, 642 ms, 854 ms, 996 ms and 1053 ms; fuzzification times of 242 ms, 475 ms, 614 ms, 843 ms and 929 ms; and defuzzification times of 235 ms, 563 ms, 655 ms, 784 ms and 928 ms, respectively. The Logarithmic Dual

Sigmoid membership function reduced categorisation bias near decision boundaries and supported a more compact rule base, which together produce substantially lower processing times in all three phases.

4.7 Performance of NR-ClinicalBERT

Figure 4.4 reports the semantic-contextual-analysis performance of the proposed NR-ClinicalBERT against four transformer-based baselines: ClinicalBERT, RoBERTa, DistilBERT and BERT.

Table 4.4: Comparative Analysis of NR-ClinicalBERT

Techniques	MRR	MAP
Proposed NR-ClinicalBERT	0.93	0.91
ClinicalBERT	0.88	0.86
RoBERTa	0.84	0.82
DistilBERT	0.81	0.78
BERT	0.78	0.73




Figure 4.4: Comparative Analysis of NR-ClinicalBERT

The proposed Non-linear Ridge function captured non-linear relationships among the learnt features. As shown in Figure 4.4 and Table 4.4, the proposed NR-ClinicalBERT reached an MRR of 0.93 and a MAP of 0.91. The baseline models ClinicalBERT, RoBERTa, DistilBERT and BERT reached MRR values of 0.88, 0.84, 0.81 and 0.78 and MAP values of 0.86, 0.82, 0.78 and 0.73, respectively. The result shows that augmenting ClinicalBERT with non-linear ridge feature learning

produces richer semantic representations of clinical narratives and ICD diagnosis text without the instability that comes with excessive feature variance.

4.8 Performance of WCRF

Figure 4.5 reports the sequential-pattern-analysis performance of the proposed WCRF against four conventional sequential baselines: Conditional Random Field (CRF), Hidden Markov Model (HMM), Variable Order Markov Model (VOMM) and Logistic Regression (LR).

Table 4.5: Comparative Analysis of WCRF

Techniques	Log-Likelihood	MCC	SER
Proposed WCRF	1.845	0.92	0.08
CRF	1.142	0.89	0.19
HMM	0.931	0.85	0.37
VOMM	0.736	0.81	0.59
LR	0.493	0.77	0.73




Figure 4.5: Graphical Comparison of WCRF

As shown in Figure 4.5 and Table 4.5, the proposed WCRF and the conventional CRF, HMM, VOMM and LR analysed sequential patterns in the vital signs with log-likelihoods of 1.845, 1.142, 0.931, 0.736 and 0.493; MCC values of 0.92, 0.89, 0.85, 0.81 and 0.77; and SER values of 0.08, 0.19, 0.37, 0.59 and 0.73, respectively. Learning stochastic temporal intervals through the Wald distribution improved the overall sequential modelling performance and gave a substantially

lower sequence error rate than any baseline. This confirms that the uniform-sampling-interval assumption embedded in the conventional CRF is harmful in real ward-based settings.

4.9 Comparison with Related Works

Table 4.6 compares the proposed framework against six recent works on hospital RRP. For each work, the table lists the main methodology, the main advantage attributed to it in the literature, and the limitation that the proposed framework overcomes.

Table 4.6: Related Works Comparison regarding RRP

Study	Methods	Advantages	Limitations
Proposed Work	Co-AVAE, HSM-BiGRU, LDS-FL, NR-ClinicalBERT, WCRF	Accurately predicted the readmission risk.	Medication history was not analysed.
Rasheed et al. (2025) [33]	LSTM	Readmission risk was effectively determined.	The prediction model was not stable.
Huang et al. (2025) [18]	CGAN and XGBoost	Generalised RRP was attained.	Relationship between the clinical data was not determined.
Albahli (2025) [2]	Random Forest algorithm	Improved clinical outcomes were attained.	Could not adapt to real-time monitoring.
Alabdulhafith et al. (2023) [1]	PSO and Ensemble Machine Learning	Global and local explanation of RRP was enhanced.	Generalisability decreased.
Zhai et al. (2025) [43]	Environment-aware Survival Predictor (E-SurP)	Accuracy in RRP was improved.	Non-linear clinical factors were hard to analyse.

Table 4.6 summarises the comparison between the proposed work and existing works on RRP. The proposed framework analyses clinical notes, vital signs and ICD codes of the patient and predicts the readmission risk accurately. The conventional works have several disadvantages: instability of the prediction model, inability to capture relationships between clinical data, lack of real-time adaptability, weaker generalisability, and difficulty modelling non-linear clinical factors. The proposed framework addresses these limitations and gives improved RRP.

4.10 Discussion

The empirical results support the design choices made in Chapter 3. Co-AVAE reached an MSE of 0.011, roughly one order of magnitude lower than the conventional VAE, consistent with the design intent that the co-attention mechanism captures local feature correlations across modalities which the conventional VAE

does not. NR-ClinicalBERT improved MRR from 0.88 (ClinicalBERT) to 0.93, showing that augmenting transformer-based clinical text models with non-linear ridge feature learning is beneficial. LDS-FL reduced rule-generation, fuzzification and defuzzification times by roughly 45%, 49% and 58% relative to conventional FL, which illustrates the practical impact of the Logarithmic Dual Sigmoid membership function. WCRF improved the log-likelihood from 1.142 (conventional CRF) to 1.845 and reduced the SER from 0.19 to 0.08, which supports the Wald distribution as a more faithful model of irregular temporal intervals in ward-recorded vital signs. HSM-BiGRU reached 98.67% accuracy with a Cohen's Kappa of 0.961, an improvement of 2.65 percentage points over the BiGRU baseline under the same training setup. This relative gain is consistent with the expectation that Hessian regularization and the SMU activation produce a more stable classifier, but because no ablation isolates the contribution of these two components, it remains an interpretation rather than a demonstrated decomposition. As noted in Section 4.11, the absolute values of the accuracy and ROC-AUC are inflated by the dependence between the prediction target and the model inputs, so the relative comparison against the baselines is more informative than the absolute figures.

Across all five sub-tasks, each proposed module improved on its backbone baselines, and the end-to-end framework reached an overall readmission accuracy of 98.67% on the MIMIC-IV dataset, a figure that should be read together with the limitations in Section 4.11. The LIME-based explainability also lets clinicians interrogate individual predictions and see the specific note phrases, vital-sign trends and ICD codes that drove the model's decision in a given case. The main advance offered by this thesis is the integrated multimodal architecture that combines cross-modal fusion, sequential temporal modelling and feature-level explainability in a single pipeline, rather than the absolute accuracy reported here.

4.11 Limitations and Threats to Validity

The results of this chapter should be read alongside several limitations. The first concerns the construct being measured and a dependence between the prediction target and the model inputs. The final prediction is a three-level readmission-risk stratification (low, moderate, high) rather than the observed binary 30-day readmission outcome. This stratification is defined by the deterministic rule of Eq. (3.10) applied to the Charlson Comorbidity Index I and the NEWS2-based Vital Instability Score V . Because I , V , and the resulting risk label L are also supplied to the HSM-BiGRU classifier as input features (Section 3.10), the prediction target is, to a large extent, a function of the model's own inputs. The accuracy of 98.67%, Cohen's Kappa of 0.961, and ROC-AUC of 0.9867 reported in Section 4.5 should therefore be interpreted as an upper bound for this label construction rather than as the difficulty of predicting genuine readmission, and they are not directly comparable to studies that predict the observed readmission event, for which reported ROC-AUC values typically lie in the range 0.65 to 0.80. The positive-class base rate and the class distribution of the label across the cohort are not characterised in this work.

A second limitation concerns the evaluation metrics. Performance is summarised using accuracy, RMSE, MAE, Cohen's Kappa, and ROC-AUC. Minority-class metrics, namely precision, recall, F1-score, and the area under the

precision-recall curve (PR-AUC), together with a confusion matrix, are not reported, even though these are the metrics most informative under class imbalance. In addition, no ablation study isolates the marginal contribution of each module to the final result; the comparative experiments in Tables 4.1 to 4.5 evaluate each module only against its own backbone baselines. The observation in Section 4.10 that the module-level improvements reinforce one another therefore remains a hypothesis rather than a demonstrated decomposition of the end-to-end accuracy.

A third limitation concerns statistical-conclusion validity. All experiments use a single 80:20 train and test split on a single dataset, with no cross-validation, repeated runs, confidence intervals, or significance testing, so the reported point estimates carry an unquantified variance. The timing results in Section 4.6 measure only the labelling stage of the pipeline; the end-to-end latency of the full framework is not measured.

Finally, regarding external validity, the evaluation uses a single centre (the MIMIC-IV database from the Beth Israel Deaconess Medical Center) with no external-site validation, and the LIME-based explanations are integrated by design but have not yet been assessed in a prospective clinician study. The most important direction for future work is therefore to re-ground the prediction target on the observed 30-day readmission event derived from the MIMIC-IV admission timestamps, withholding the derived risk label L and the raw scores I and V from the classifier inputs, and to re-evaluate the framework using precision, recall, F1-score, PR-AUC, and a confusion matrix on that target. This should be accompanied by an ablation study that removes one module at a time to quantify each module's contribution, and by k-fold cross-validation with confidence intervals and a significance test. External validation on a second clinical site and a prospective study of the LIME explanations with clinicians are the longer-term extensions.

CHAPTER 5

CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT

5.1 Conclusion

This thesis has proposed and evaluated a multimodal deep-learning framework for hospital readmission risk prediction. The framework integrates three modalities of clinical evidence (unstructured clinical notes, time-series vital signs, and structured ICD diagnosis codes) and produces an interpretable readmission risk classification together with a feature-level explanation. Five novel modules were introduced to address specific limitations of existing approaches: NR-ClinicalBERT for semantic contextual analysis with non-linear ridge feature learning; Co-AVAE for cross-model integration analysis through co-attention augmentation of the variational autoencoder; LDS-FL for risk labelling with a Logarithmic Dual Sigmoid membership function; WCRF for sequential pattern analysis of vital signs through Wald-distributed temporal intervals; and HSM-BiGRU for the final classification with Hessian regularization and SMU activation. LIME-based explainability was integrated for transparent and clinically trustworthy predictions.

Empirical evaluation on the MIMIC-IV dataset (v3.1), with 299,712 patient records, showed consistent improvements over backbone baselines across all five sub-tasks. NR-ClinicalBERT reached an MRR of 0.93 and MAP of 0.91, beating ClinicalBERT, RoBERTa, DistilBERT and BERT. Co-AVAE reached an MSE of 0.011, almost an order of magnitude lower than the conventional VAE, and reduced perceptual loss and MMD substantially relative to four autoencoder baselines. LDS-FL reached a defuzzification time of 235 ms, which brings the labelling latency down to roughly 42% of that of conventional FL. WCRF improved the log-likelihood from 1.142 to 1.845 and reduced the sequence error rate from 0.19 to 0.08 relative to the conventional CRF. HSM-BiGRU predicted readmission risk with an accuracy of 98.67%, an RMSE of 0.154, an MAE of 0.118, a Cohen's Kappa of 0.961 and a ROC-AUC of 0.9867, beating the BiGRU, BiLSTM, LSTM and RNN baselines. It is also compared qualitatively against six recent works on hospital readmission prediction in Table 4.6.

5.2 Practical Implications

The proposed framework can be deployed in hospitals and healthcare centres to make faster and more accurate decisions about patient care and discharge planning. Because the framework runs on data already routinely captured in modern EHR systems (clinical notes, vital signs and ICD codes), integrating it into existing clinical workflows places no additional data-collection burden on clinicians or patients. The interpretable output lets care teams identify patients who need extra support before discharge, plan post-discharge follow-up appointments early, and allocate

transitional-care resources to those who most need them. Hospitals operating under value-based reimbursement schemes that penalise excess 30-day readmissions will benefit both clinically and economically from such early-warning tools.

5.3 Social Impact

Unplanned readmissions disproportionately affect elderly patients, patients with multiple chronic conditions, and patients with limited access to community-based primary care. By flagging high-risk patients before discharge, the proposed framework can support more equitable allocation of post-acute resources, reduce avoidable suffering for patients and their families, and ease the financial strain that repeated hospitalisations place on households. At the system level, accurate readmission prediction reduces avoidable occupancy of hospital beds, frees capacity for new admissions, shortens waiting lists, and improves the overall use of finite healthcare resources. The integrated explainability layer also ensures that the clinician, not the algorithm, remains the final decision-maker. This preserves professional accountability and the patient-clinician relationship.

5.4 Future Scope

Several directions for future work stand out. First, the present framework uses three modalities; future versions can incorporate laboratory reports, medication histories, radiology findings, wearable-sensor data and genomic information for a fuller patient profile. Second, the model can be extended to predict specific causes of readmission, rather than the binary or three-level outcome considered here, which would support condition-specific interventions. Third, prospective validation on external datasets from multiple institutions and geographies will help assess generalisability and surface possible site-specific biases. Fourth, federated learning [31] could be explored as a way to train across institutions without centralised data sharing, which would address the privacy concerns intrinsic to clinical data. Fifth, integration with real-time streaming infrastructure would allow continuous risk reassessment as new observations arrive during a patient's stay, rather than a single prediction at discharge. Finally, formal user studies with clinicians could quantify the effect of the LIME-based explanations on clinical decision-making, trust, and downstream patient outcomes.

REFERENCES

- [1] M. Alabdulhafith, H. Saleh, H. Elmannai, Z. H. Ali, S. El-Sappagh, J. Hu, and N. El-Rashidy, "A clinical decision support system for edge/cloud ICU readmission model based on particle swarm optimization, ensemble machine learning, and explainable artificial intelligence," *IEEE Access*, vol. 11, pp. 100604–100621, 2023, doi: 10.1109/ACCESS.2023.3312343.
- [2] S. Albahli, "Predictive analytics for diabetic patient care: leveraging AI to forecast readmission and hospital stays," *Computer Modeling in Engineering & Sciences*, vol. 143, no. 1, pp. 1–34, 2025, doi: 10.32604/cmescs.2025.058821.
- [3] H. Al-Jehani, A. Hafez, U. Elhadi, W. Issawi, and A. Anazi, "Neurosurgical readmissions as a key performance indicator: a single-center analysis of causes, patterns, and clinical relevance," *Interdisciplinary Neurosurgery: Advanced Techniques and Case Management*, vol. 43, pp. 1–6, 2026, doi: 10.1016/j.inat.2026.102242.
- [4] T. Almeida, P. Moreno, and C. Barata, "Prediction of 30-day hospital readmission with clinical notes and EHR information," in *Iberian Conference on Pattern Recognition and Image Analysis*, 2025, pp. 1–13.
- [5] S. G. Alonso, I. H. Montano, I. de la T. Diez, J.-A. Roman-Gallego, M.-L. Perez-Delgado, and M. A. Conde, "SchizoReadPredict: a tool for prediction readmission risk of schizophrenia patients in a Spanish region," *Procedia Computer Science*, vol. 263, pp. 719–725, 2025, doi: 10.1016/j.procs.2025.07.086.
- [6] K. Biswas, S. Kumar, S. Banerjee, and A. K. Pandey, "SMU: smooth activation function for deep networks using smoothing maximum technique," *arXiv preprint arXiv:2111.04682*, 2021.
- [7] M. Bonomo, M. G. Hermesen, S. Kaskovich, M. J. Hemmrich, J. C. Rojas, K. A. Carey, L. R. Venable, M. M. Churpek, and V. G. Press, "Using machine learning to predict likelihood and cause of readmission after hospitalization for chronic obstructive pulmonary disease exacerbation," *International Journal of Chronic Obstructive Pulmonary Disease*, vol. 17, pp. 2701–2709, 2022, doi: 10.2147/COPD.S379700.
- [8] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001, doi: 10.1023/A:1010933404324.
- [9] M. E. Charlson, P. Pompei, K. L. Ales, and C. R. MacKenzie, "A new method of classifying prognostic comorbidity in longitudinal studies: development and validation," *Journal of Chronic Diseases*, vol. 40, no. 5, pp. 373–383, 1987, doi: 10.1016/0021-9681(87)90171-8.
- [10] T. Chen and C. Guestrin, "XGBoost: a scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.
- [11] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1724–1734, doi: 10.3115/v1/D14-1179.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proc. 2019 Conf. of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, 2019, pp. 4171–4186, doi: 10.18653/v1/N19-1423.

- [13] J. Donzé, D. Aujesky, D. Williams, and J. L. Schnipper, "Potentially avoidable 30-day hospital readmissions in medical patients: derivation and validation of a prediction model," *JAMA Internal Medicine*, vol. 173, no. 8, pp. 632–638, 2013, doi: 10.1001/jamainternmed.2013.3023.
- [14] C. L. Dugan, M. V. Kudlinski, S. Ganesh, G. Rosenberg, T. Harano, S. Wightman, S. Atay, A. W. Kim, and B. V. Udelsman, "Social contributions as risk factors for readmissions after lung transplantation: clinical and financial implications," *JHLT Open*, vol. 9, pp. 1–9, 2025, doi: 10.1016/j.jhlto.2025.100300.
- [15] Z. He, H. Li, S. Tian, J. Wen, G. Yuan, and A. Li, "A comparative study of structured and narrative EHR data for 30-day readmission risk assessment," *Electronics*, vol. 14, pp. 1–18, 2025, doi: 10.3390/electronics14204033.
- [16] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.
- [17] K. Huang, J. Altosaar, and R. Ranganath, "ClinicalBERT: modeling clinical notes and predicting hospital readmission," *arXiv preprint arXiv:1904.05342*, 2019.
- [18] Y. Huang, T. Feng, X. Song, K. Luo, H. Chen, Y. Xie, Y. Xu, and X. Chen, "Explainable machine learning model for 1-year readmission risk prediction in AECOPD patients: integrating relief feature selection with sample augmentation," *BMC Infectious Diseases*, vol. 25, pp. 1–14, 2025, doi: 10.1186/s12879-025-12010-1.
- [19] S. F. Jencks, M. V. Williams, and E. A. Coleman, "Rehospitalizations among patients in the Medicare fee-for-service program," *New England Journal of Medicine*, vol. 360, no. 14, pp. 1418–1428, 2009, doi: 10.1056/NEJMs0803563.
- [20] A. Johnson, L. Bulgarelli, T. Pollard, B. Gow, B. Moody, S. Horng, L. A. Celi, and R. Mark, "MIMIC-IV (version 3.1)," *PhysioNet*, 2024. [Online]. Available: <https://physionet.org/content/mimiciv/3.1/>
- [21] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2014, arXiv:1312.6114.
- [22] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2017, arXiv:1609.02907.
- [23] J. Lafferty, A. McCallum, and F. Pereira, "Conditional random fields: probabilistic models for segmenting and labeling sequence data," in *Proc. 18th Int. Conf. on Machine Learning (ICML)*, 2001, pp. 282–289.
- [24] J. Li, L. Liang, Y. S. Cai, Y. Zuo, J. Su, L. Feng, H. Wang, and Z. Tong, "Tracking COPD exacerbation patterns and forecasting readmission risks utilizing electronic medical records," *International Journal of Medical Informatics*, vol. 189, pp. 1–7, 2024, doi: 10.1016/j.ijmedinf.2024.105505.
- [25] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: a robustly optimized BERT pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [26] Y. Liu, L. Du, L. Li, L. Xiong, H. Luo, E. Kwaku, X. Mei, Y. Y. Cui, Y. Zhou, L. Zeng, S. Li, K. Wang, J. Zheng, Z. Liu, H. Hu, and R. Yue, "Development and validation of a machine learning-based readmission risk prediction model for non-ST elevation

myocardial infarction patients after percutaneous coronary intervention," *Scientific Reports*, vol. 14, pp. 1–13, 2024, doi: 10.1038/s41598-024-64048-x.

- [27] R. Loutati, A. Ben-Yehuda, S. Rosenberg, and Y. Rottenberg, "Multimodal machine learning for prediction of 30-day readmission risk in elderly population," *The American Journal of Medicine*, vol. 137, no. 7, pp. 617–628, 2024, doi: 10.1016/j.amjmed.2024.04.002.
- [28] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 29, 2016, pp. 289–297.
- [29] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017, pp. 4765–4774.
- [30] C. Mancheno-Velasco, M. Narvaez-Camargo, O. M. Lozano-Rojas, and M. Sanchez-Garcia, "Readmission and dropout in outpatient centers: an analysis of real-world data in patients with dual-diagnosis," *International Journal of Mental Health and Addiction*, vol. 23, no. 6, pp. 4423–4443, 2025, doi: 10.1007/s11469-024-01360-7.
- [31] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1273–1282.
- [32] M. Pham, T. T. Mai, M. Crane, M. Ebiele, R. Brennan, M. E. Ward, U. Geary, N. McDonald, and M. Bezbradica, "Forecasting patient early readmission from Irish hospital discharge records using conventional machine learning models," *Diagnostics*, vol. 14, pp. 1–19, 2024, doi: 10.3390/diagnostics14212405.
- [33] N. Rasheed, U. Khaliq, and A. Ahmed, "Predicting patient ICU readmission using recurrent neural networks with long short-term memory," *IEEE Access*, vol. 13, pp. 153375–153385, 2025, doi: 10.1109/ACCESS.2025.3601117.
- [34] M. T. Ribeiro, S. Singh, and C. Guestrin, "\"Why should I trust you?\": explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
- [35] Royal College of Physicians, *National Early Warning Score (NEWS) 2: Standardising the Assessment of Acute-Illness Severity in the NHS*. London, U.K.: RCP, 2017.
- [36] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.
- [37] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997, doi: 10.1109/78.650093.
- [38] X. Song, J. Shi, C. Zhu, F. Xian, Z. Dong, and J. Li, "XGBoost machine learning algorithm for predicting unplanned readmission in elderly patients with coronary heart disease," *Geriatric Nursing*, vol. 66, pp. 1–8, 2025, doi: 10.1016/j.gerinurse.2025.103609.
- [39] S. Tang, A. Tariq, J. A. Dunnmon, U. Sharma, P. Elugunti, D. L. Rubin, B. N. Patel, and I. Banerjee, "Predicting 30-day all-cause hospital readmission using multimodal

- spatiotemporal graph neural networks," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 4, pp. 2071–2082, 2024, doi: 10.1109/JBHI.2023.3236888.
- [40] C. van Walraven, I. A. Dhalla, C. Bell, E. Etchells, I. G. Stiell, K. Zarnke, P. C. Austin, and A. J. Forster, "Derivation and validation of an index to predict early death or unplanned readmission after discharge from hospital to the community," *Canadian Medical Association Journal*, vol. 182, no. 6, pp. 551–557, 2010, doi: 10.1503/cmaj.091117.
- [41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017, pp. 5998–6008.
- [42] A. F. Woudneh, "Prediction of 30-day unplanned hospital readmission among elderly patients using patient and ward level factors," *Scientific Reports*, vol. 15, pp. 1–11, 2025, doi: 10.1038/s41598-025-28060-z.
- [43] Y. Zhai, Y. Li, L. He, X. Li, and W. Luo, "Leveraging environmental information for enhanced prediction of cardiac readmissions," *iScience*, vol. 28, pp. 1–20, 2025, doi: 10.1016/j.isci.2025.114104.
- [44] T. Zhu, J. Jiang, Y. Hu, and W. Zhang, "Individualized prediction of psychiatric readmissions for patients with major depressive disorder: a 10-year retrospective cohort study," *Translational Psychiatry*, vol. 12, pp. 1–18, 2022, doi: 10.1038/s41398-022-01937-7.