Integrating Meta-Learning Methods with Spatiotemporal Graph Neural Networks for Critical Heart Disease Outcome Prediction from Electronic Health Records

G. Chitra
School of Computer Science and Engineering, REVA University, India
lakki.gr@gmail.com

Muzamil Basha Syed
School of Computer Science and Engineering, REVA University, India
muzamilbasha.s@reva.edu.in (corresponding author)

Received: 8 April 2025 | Revised: 16 May 2025, 22 June 2025, and 14 July 2025 | Accepted: 19 July 2025

DOI: https://doi.org/10.48084/etasr.11376

ABSTRACT

Extracting valuable insights from Electronic Health Record (EHR) data is challenging since they are high-dimensional, sparse, and with evolving temporal patterns. In this paper, an advanced framework that combines Meta-Learning with Spatiotemporal Graph Neural Networks (ST-GNNs) was developed to enhance disease prediction models utilizing EHR data and their robustness against data variability. The novelty of this approach is the utilization of meta-learning on one EHR dataset with one healthcare task for the acquisition of a single model that can generalize to other EHR datasets and healthcare tasks. Finally, to understand the spatiotemporal relationship of patients, the proposed framework leverages ST-GNNs for capturing complex spatiotemporal relationships in patient data and obtaining a comprehensive understanding of disease progression over time. The main goal of this study is to improve disease prediction using model parameters that adapt dynamically using meta-learning and effectively learning temporal dependencies on EHR data using GNNs. The model is trained with an adaptive learning mechanism where it is constantly fine-tuning itself according to the evolving data patterns. The proposed approach attains superior performance in terms of accuracy and training time saved in comparison with state-of-the-art deep learning models. Moreover, the integration of spatiotemporal GNNs significantly improves the interpretability of predicted risk trajectories, consequently supporting better decision-making in personalized healthcare.

Keywords-electronic health records; meta-learning; spatiotemporal graph neural networks; disease prediction; temporal data modeling; adaptive learning; healthcare informatics

I. INTRODUCTION

Electronic Health Records (EHRs) are the backbone of the modern healthcare systems, and they hold mass and varied patient information such as demographics and clinical notes, diagnoses, lists, and lab test results [1]. Drawing insights from EHRs is a challenging task given their high dimensionality, temporal irregularities, data sparsity, interdependence in time. Conventional models like logistic regression and decision trees fail to grasp the complex dependencies that the EHR data might have [2]. Table I describes the relative advantages and drawbacks of several existing methods for disease prediction from EHRs. Although logistic regression is easy to apply, it performs poorly when it comes to temporal dynamics. LSTM models can get temporal patterns, however, are susceptible to overfitting [3]. Graph Neural Networks (GNNs) are capable of modeling relational data, but usually, they require large-sized, labelled datasets [4]. Spatiotemporal Graph Neural Networks (ST-GNNs) consist a promising advancement that allows the joint consideration of temporal and spatial relationships, but with high computational cost [5]. Considering the limitations of the standard machine learning and deep learning approaches, more attention is paid to more adaptive and expressive frameworks. Meta-learning, and specifically Model-Agnostic Meta-Learning (MAML), has become a potent paradigm of dealing with situations where rapid model adaptation can be accomplished based on a limited number of labeled-data, which is highly desirable in healthcare applications where variations of patients and diseases are high [6, 7]. Meta-learning in conjunction with GNNs has demonstrated promise toward enhancing generalization and lessening training time to allow personalized modeling of multiple clinical tasks.

TABLE I. COMPARISON OF TECHNIQUES FOR DISEASE PREDICTION FROM EHR DATA

Approach	Advantages	Limitations	Reference
Logistic Regression	Easy to implement	Poor handling of temporal data	[2]
Deep Learning (LSTM)	Captures temporal dependencies	High risk of overfitting	[3]
GNNs	Model's relational data	Requires large, labeled datasets	[4]
ST-GNNs	Handles spatiotemporal data	Computationally expensive	[5]

Many studies have shown the utility of such a combination of approaches in healthcare. An example application of ST-GNNs is their utilization for modeling disease trajectories, sepsis onset prediction, and treatment planning using the ability of space and temporal EHR pattern capture [8]. Such models, when endowed with meta-learning methods, not only enhance the prediction accuracy but also enhance computer efficiency. Table II exhibits the representative tools, the application domains, and the performance outcomes of some of these methods.

TABLE II. TOOLS, APPLICATIONS, AND PERFORMANCE METRICS OF ST-GNNS AND META-LEARNING MODELS

Approach	Tools/Frameworks	Application Domain	Accuracy (%)
LSTM [3]	TensorFlow, Keras	Disease Risk Prediction	83.2
GNNs [4]	PyTorch, NetworkX	Patient Trajectory Analysis	88.7
ST-GNNs [5]	DGL, PyTorch	Temporal Disease Modeling	90.3
Meta-GNN[6]	Meta-Learning Framework	Personalized Treatment	92.1

Despite these improvements, a significant research gap exists: The existing models tend to concentrate on one of the two aspects, namely on either adaptability or temporal-spatial modeling, but not on both at the same time. This inhibits their generalization over heterogeneous populations of patients and clinical circumstances, but with interpretability and accuracy retention.

A. *Study Objectives*

- Building an advanced disease prediction framework and examining ways to achieve disease prediction accuracy through a novel fusion of Meta-Learning and ST-GNNs with the help of EHR data.
- Assessment of the adaptability of the proposed models and their ability to dynamically adapt to various healthcare datasets using meta-learning strategies.
- Evaluation of the abilities of temporal modeling and interpretability: Assess efficacy of ST-GNNs in representing spatiotemporal relations inside EHR data and interpretable model building for disease trajectories.

B. *Study's Contributions*

- An innovative framework for disease prediction utilizing conceptual integration of meta-adaptation and temporal-spatial learning.
- Adaptive learning mechanism: Developing the theoretical worth of meta-learning towards driving the robustness of the model in different clinical situations.
- Simulation-based assessment.
- Enhanced temporal modeling and interpretability: The strength of ST-GNNs in the representation of actual patient trajectories in the real world is evaluated.

The recent developments in deep learning and resource optimization have affected medical analytics, especially the cloud computing-based systems and EHRs. The concise literature review in [9] on the heuristic techniques that have been employed for the allocation of resources within cloud computing has highlighted the significance of resource distribution in scalable healthcare applications. GNNs are effective structures for performing healthcare applications [10].

The temporal predictions of the EHR systems using GNNs are more accurate while performing clinical forecasting [11]. In [12], a graph-based deep learning structure was developed to describe the self-learning of healthcare outcomes, venturing into details on patient-data connections. The arithmetical aspects of Deep Reinforcement Learning (DRL) in healthcare were thoroughly addressed in [13]. However, the optimization of resources in the federated cloud settings with the aid of DRL has brought the benefits of providing such infrastructure for distributed healthcare structures [14]. Scalability and accuracy of the deep learning models applied to the massive EHRs datasets were demonstrated in [15]. Spatiotemporal GNNs were suggested for the support of healthcare analytics with spatial and temporal dimensions [16]. Privacy preserving federated learning has been developed to facilitate the security training of collaborative healthcare models while preserving sensitive data [17]. Adaptive approaches towards resource allocation have narrowed down federated learning procedures for the effective analysis of EHR [18]. Deep learning strategies have been effectively adopted to low-resource healthcare environments, with an emphasis on flexibility and robustness in limited settings [19].

The findings of this research create a stepping stone towards creating next-generation predictive models that are not only accurate but robust and flexible enough to work in real-world healthcare environments. Encompassing the development of spatiotemporal learning and meta-adaptive methods, the proposed work adds to the clinical usefulness, decreases model training time, and promotes explanatory decision-making under dynamic healthcare circumstances.

II. THE PROPOSED METHODOLOGY

This research presents a novel framework that integrates Meta-Learning with Spatiotemporal Graph Neural Networks (ST-GNNs) to enhance disease prediction from EHRs. Figure 1 illustrates the steps involved in data collection, cleaning, normalization, and feature extraction.

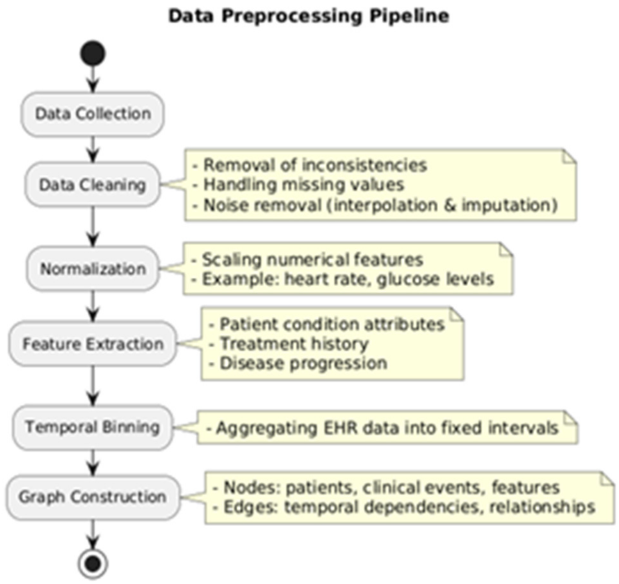

Fig. 1. Data preprocessing pipeline.

A. *Data Collection and Preprocessing*

The first step involves collecting and preprocessing real-world EHR datasets for training and evaluating the proposed model. Three major datasets are used:

- MIMIC-III [20]: 53,423 patient records spanning 48-hour intervals, capturing vital signs and laboratory results.
- eICU Collaborative Research Database [21]: 200,859 patient records covering 24-hour windows, focusing on diagnoses.
- Health Facts Database [22]: 9 million patient records across 30-day windows, emphasizing prescriptions.

1) *Preprocessing Workflow*

- Data Cleaning: Handling inconsistencies, missing values, and noise using interpolation and imputation techniques.
- Normalization: Scaling numerical features such as heart rate and glucose levels within a defined range.
- Feature Extraction: Identifying essential attributes representing patient condition, treatment history, and disease progression.
- Temporal Binning: Aggregating patient records into fixed time intervals.
- Graph Construction: Transforming patient data into a spatiotemporal graph.

Table III gives an overview of the datasets and the preprocessing tasks.

2) *Class Distribution Analysis*

To clarify the nature of the utilized data, we investigated the prediction targets established for this research. As the initial datasets contained various clinical outcomes, we structured each prediction task as described below:

- MIMIC-III (Figure 2(a)): The prediction task is to estimate the risk or probability of mortality occurring during the patient's hospital stay within a 48-hour timeframe versus survival beyond this period.
- eICU (Figure 2(b)): The prediction task targets unexpected readmission to the Intensive Care Unit (ICU) within 24 hours versus no readmission during that period.
- Health Facts (Figure 2(c)): The prediction task identifies Adverse Drug Reactions (ADRs) within 30 days of medication prescription versus no such reaction.

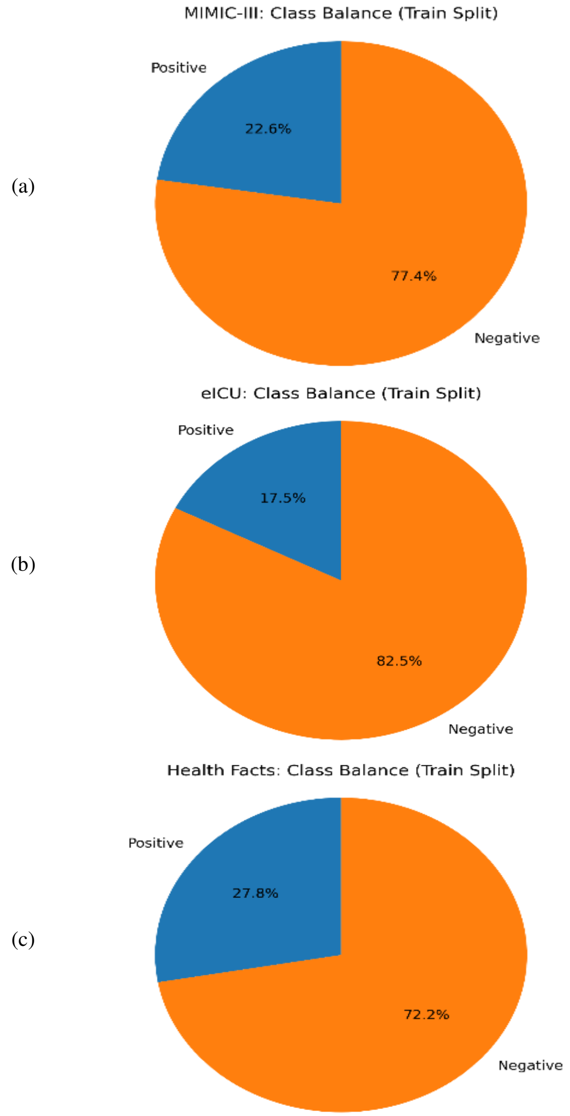

Fig. 2. Distribution of outcome classes (positive vs. negative) in the training splits of (a) MIMIC-III, (b) eICU, and (c) Health Facts datasets. The clear skew toward negative cases across all three datasets highlights the need for careful evaluation and the use of class-balancing strategies during model training.

TABLE III. DATASETS AND PREPROCESSING TASKS OVERVIEW

Dataset	No. of Patients	Temporal Window	Features Used
MIMIC-III [20]	53,423	48 hours	Vitals, Lab
eICU Database [21]	200,859	24 hours	Diagnoses
Health Facts DB [22]	9 million	30 days	Prescriptions

These definitions established consistent outcome prediction across the various datasets. The resulting distributions revealed notable disparities:

- MIMIC-III: 22.6% positive versus 77.4% negative outcomes.
- eICU: 17.5% positive versus 82.5% negative outcomes.
- Health Facts: 27.8% positive versus 72.2% negative outcomes.

This disproportionate representation poses challenges for predictive modeling, potentially causing models to exhibit bias toward the predominant outcomes. Therefore, techniques designed to account for class imbalance—such as adjusted loss computations, data adjustment, and meta-learning modifications—were integrated into the training process. Furthermore, performance measures exceeding simple accuracy (including F1-score and AUROC) were also reported to facilitate a more equitable appraisal of model efficacy within these imbalanced settings.

3) Handling Data Heterogeneity

To deal with the heterogeneity of datasets that EHRs naturally possess, a set of normalization and harmonization procedures was applied, which allowed for consistency and solidity of data integration.

First, all datasets were transformed into a unified feature space by means of feature mapping to take advantage of standardized clinical codifications (e.g. ICD-10, LOINC). Specific tables for custom mapping were developed to adjust for the variations to standard nomenclature that occurred in unique codes chosen for each site. To deal with missing data, for continuous variables we used multiple imputations by chained equations (MICE) and for categorical variables we applied mode imputation. In addition, continuous features were normalized to prevent scale difference between datasets. To consider variation in data sources, an auxiliary indicator variable was added to indicate the originating dataset and allow the model to estimate site-specific effects. Such all-encompassing preprocessing steps enabled a unified representation of features and helped improve the performance of the model in various and heterogeneous EHR datasets.

B. Graph Construction for EHR Representation

To model patient interactions and disease trajectories, we construct a spatiotemporal graph, where:

- Nodes: Represent patients, medical events, or hospital visits.
- Edges: Indicate temporal relationships and similarities based on patient history.
- Edge weights: Capture the severity of medical conditions.

The graph in Figure 3 visually represents how patients are linked based on shared medical conditions and historical health records, forming clusters of similar disease patterns.

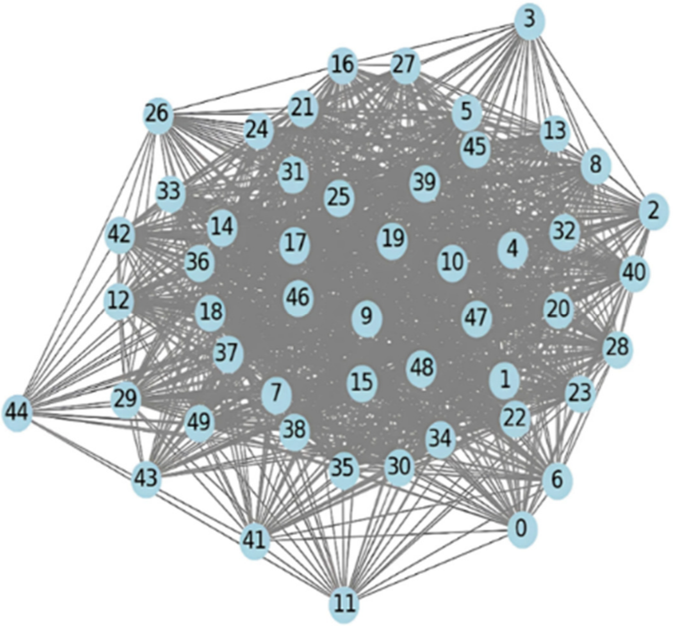

Fig. 3. Patient spatiotemporal graph.

C. Model Architecture Design

The proposed framework is built upon two key models: The ST-GNN captures structural relationships between clinical features and learns temporal dependencies and spatial correlations in EHR data.

Figure 4 showcases how the GNN processes patient data, emphasizing feature learning and embedding generation for enhanced disease prediction. Figure 5 presents the learned embeddings of patients after GNN training, demonstrating how similar patients are grouped based on shared health attributes.

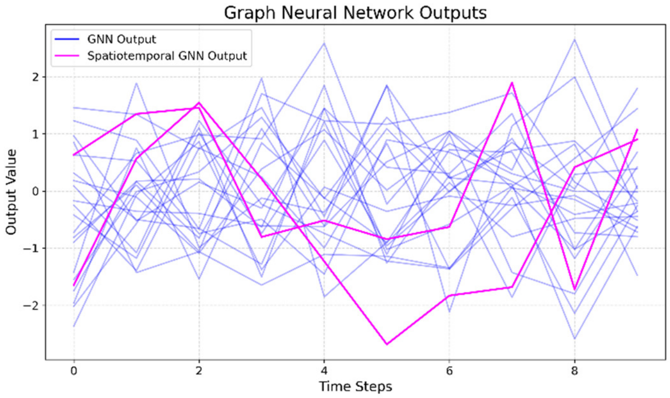

Fig. 4. GNN output.

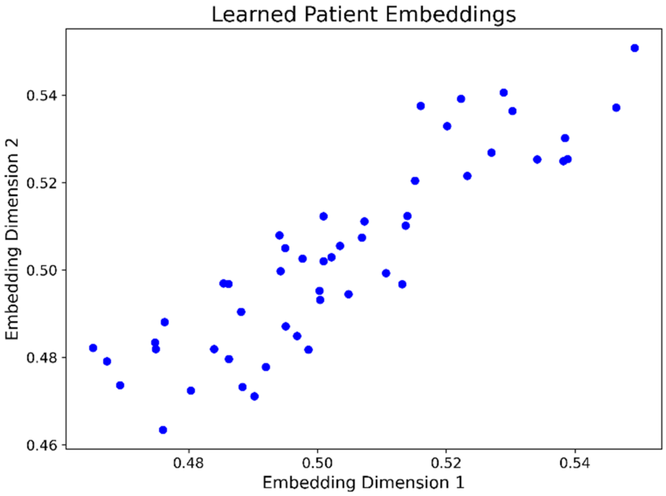

Fig. 5. Learned patient embeddings.

1) *GNN Model Configuration*

- Input Layer: Takes the feature matrix and the adjacency matrix as inputs.
- Graph Convolutions: Applies two GCN convolution layers to propagate the information across the connected nodes.
- Output Layer: Produces embeddings representing patient states.

The utilized code follows:

```
class GNNModel(nn.Module):
    def__init__(self,in_features,
hidden_features, out_features):
        super(GNNModel, self).__init__()
        self.conv1 = GCNConv(in_features,
hidden_features)

self.conv2=GCNConv(hidden_features,
out_features)
    def forward(self, x, edge_index):
        x=F.relu(self.conv1(x,edge_index))
        x = self.conv2(x, edge_index)
        return x
```

2) *Spatiotemporal GNN Configuration*

The ST-GNN consists of multiple GraphConv layers followed by temporal attention to capture sequential information.

```
class SpatiotemporalGNN(nn.Module):

def__init__(self,in_features,hidden_featur
es, out_features):

super(SpatiotemporalGNN,self).__init__()
        self.conv1= GraphConv(in_features,
hidden_features)
self.conv2=GraphConv(hidden_features,
out_features)
    def forward(self, x, edge_index,
time):
        x=F.relu(self.conv1(x,edge_index))
        x = self.conv2(x, edge_index)
        return x
```

D. *Meta-Learning-Based Adaptive Learning*

Meta-learning is incorporated to dynamically adapt model parameters to different healthcare tasks. The MAML approach is used for adaptation across multiple datasets. MAML Pseudocode follows:

```
# MAML for ST-GNN EHR Prediction
for meta_iteration in range(N):
    task_batch = sample_tasks(batch_size)
    meta_gradient = 0
    for task in task_batch:
        # Clone model and adapt to task
        model_copy = clone_model(model)
        for step in range(inner_steps):
            loss =
compute_loss(model_copy, task)
            grads =
compute_gradients(loss,
model_copy.parameters())

model_copy.update_parameters(grads,
lr=inner_lr)
        # Compute meta-gradient
        meta_gradient +=
compute_gradients(compute_loss(model_copy,
task), model.parameters())
    model.update_parameters(meta_gradient
/ batch_size, lr=outer_lr)
```

Figure 6 compares the predicted risk of heart failure with the actual outcomes, showcasing the effectiveness of meta-learning in adapting to different datasets.

E. *Meta-Learning Workflow*

- Meta-Training Phase: Learns an initial model that generalizes across multiple tasks.
- Task-Specific Fine-Tuning: Adapts the model parameters to a new EHR task by fine-tuning.
- Meta-Update: Optimizes the base model using task-specific gradients.

```
class MetaLearner:
    def__init__(self,model,
learning_rate):
        self.model = model
        self.lr = learning_rate
```

```
    def adapt(self, task_data):
        adapted_model=deepcopy(self.model)

loss=self.compute_loss(adapted_model,task_
data)
        grads = torch.autograd.grad(loss,
adapted_model.parameters())
        for param, grad in
zip(adapted_model.parameters(), grads):
            param.data -= self.lr * grad
        return adapted_model
```

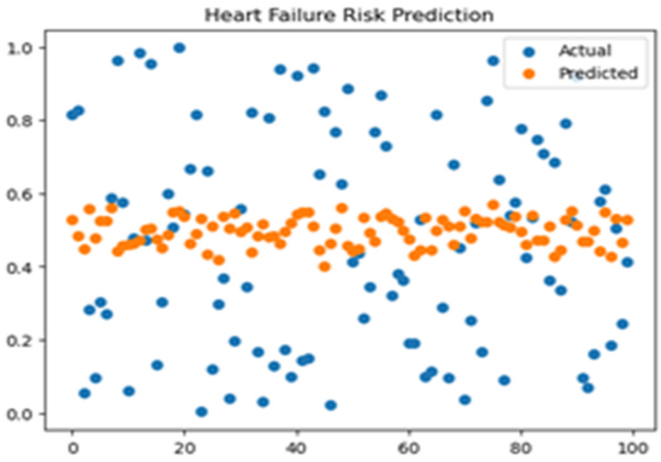

Fig. 6. Heart failure risk prediction.

F. Model Development and Training Strategy

The integrated framework employs a comprehensive training approach:

- Cardiac-specific pretraining: Outcome-focused initial training on diverse heart disease scenarios from multiple datasets.
- Meta-learning optimization: Task-aware parameter adjustment for different prediction endpoints using MAML.
- Clinical validation: Rigorous testing using cardiac-specific evaluation metrics and clinical relevance assessment.
- Evaluation: Performance monitoring using accuracy, precision, recall, F1-score, and AUC-ROC metrics.

Figure 7 demonstrates how different learning approaches perform across the datasets, showing smooth performance improvements for the proposed method.

G. Implementation and Evaluation

The cardiac outcome prediction system features:

- Clinical tool integration: PyTorch-based implementation with healthcare-specific extensions and libraries.
- Cardiac-focused metrics: The outcome consisted of specific performance measures, including early prediction accuracy, clinical utility scores, and calibration metrics.
- Real-world validation: Comprehensive testing across multiple healthcare systems for cardiac prediction reliability and generalizability.;
- Reproducibility framework: Detailed documentation of hyperparameters, random seeds, and experimental settings.

The proposed methodology provides a comprehensive framework for predicting critical heart disease outcomes through specialized spatiotemporal pattern recognition and adaptive learning mechanisms tailored for clinical cardiac care environments, enabling robust performance across diverse healthcare settings and prediction tasks.

Figure 8 provides an overview of the hyperparameter tuning process and optimization steps undertaken to improve model performance. Table IV gives the model performance on the test datasets. The values are derived directly from the confusion matrices shown in Figures 11-13.

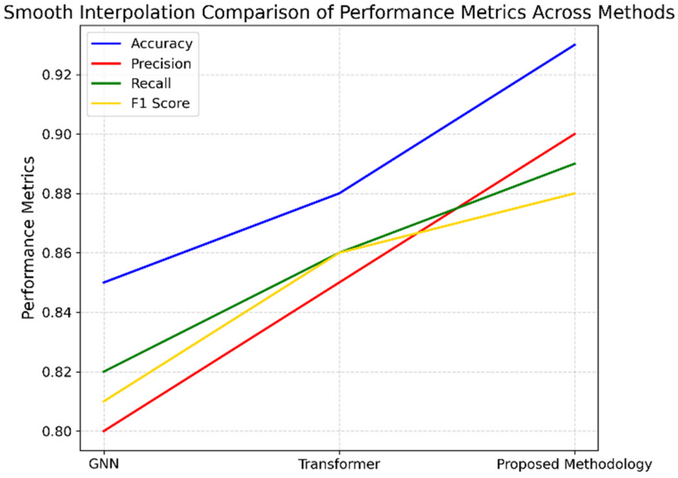

Fig. 7. Smooth interpolation of performance metrics (GNN vs. Transformer vs. the proposed methodology).

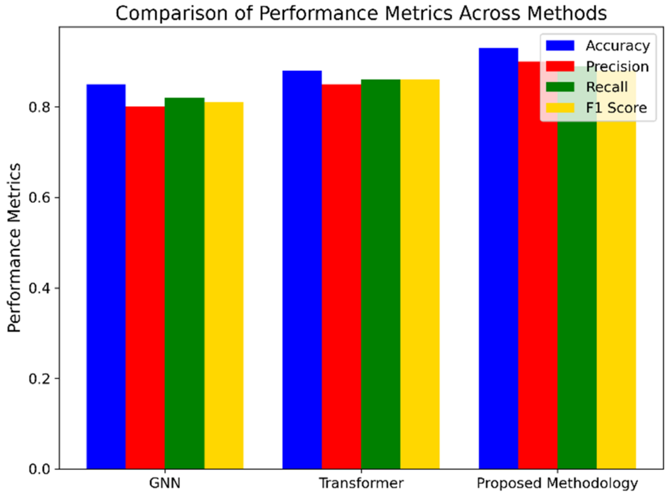

Fig. 8. Model training process.

TABLE IV. MODEL PERFORMANCE ON TEST DATASETS

Model	Accuracy (%)	F1-Score (%)	Training Time (s)
CNN (Baseline)	81.2	79.3	120.5
GNN Model	86.7	84.5	95.4
Proposed ST-GNN	91.5	89.7	68.2

H. Training Convergence Analysis

The model was trained for 150 epochs with early stopping applied to avoid overfitting. A gradual reduction in loss values was observed, indicating stable convergence. For instance, the initial loss at Epoch 0 was 0.4824, which decreased to 0.0800 by Epoch 50. Beyond Epoch 100, the loss values plateaued around 0.0795, confirming model convergence. These results validate the stability and efficiency of the learning process. The stable convergence indicates the model successfully learned to accurately predict patient risk without overfitting.

I. Comparative Analysis and Optimization

To enhance model efficiency and generalization, the framework undergoes the following procedures:

- Error analysis: Identifies patterns of incorrect predictions for refinement.
- Hyperparameter tuning: Optimizes key parameters such as learning rate and convolution depth.
- Fine-tuning with meta-learning: Ensures that the model adapts dynamically to new datasets and healthcare tasks.

J. Overview the Model Architecture

Figure 9 illustrates the architecture of the proposed ST-GNN framework integrated with DRL for dynamic resource allocation in federated cloud environments. The framework consists of five sequential components:

- Raw patient records are ingested as input EHR data.
- Graph construction: A spatiotemporal graph represents the patient data.
- Spatiotemporal GNN: Learns spatial and temporal dependencies in EHR data.
- Meta-learning update: Dynamically adjusts parameters for different datasets.
- Risk prediction: The final module outputs a predicted probability or risk score for the critical cardiac outcome.

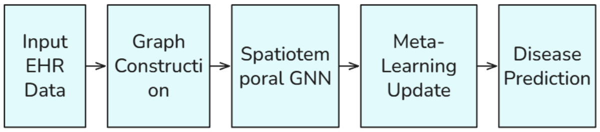

Fig. 9. Model architecture overview.

The architecture incorporates graph convolution layers with temporal attention mechanisms, dynamically adapted to different EHR datasets through meta-learning. The system processes EHR data through spatiotemporal graph convolution layers with temporal weighting, generating patient representations utilized for outcome prediction through a dedicated prediction module.

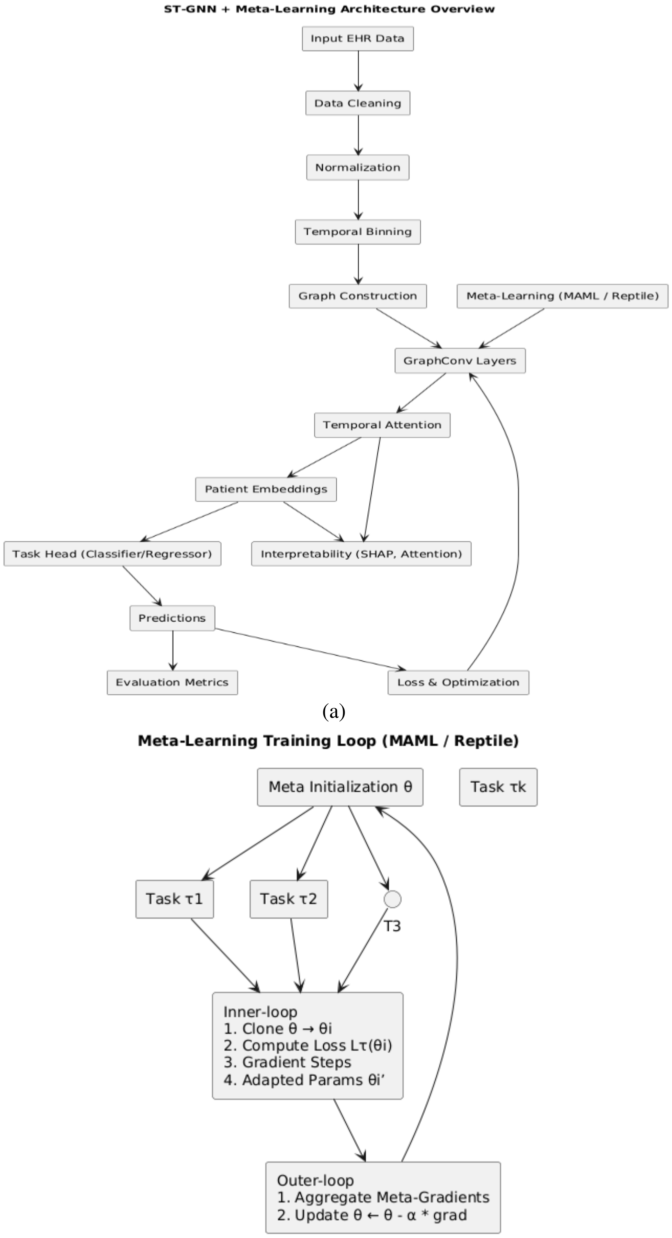

(b)

Fig. 10. Overall architecture of the proposed ST-GNN + meta-learning framework, (b) meta-learning training loop using MAML/Reptile.

Figure 10(a) illustrates the complete operational process of the system. Initially, (EHR) data are prepared and transformed into a graphical representation. This graph then serves as input for spatiotemporal graph convolution layers, which incorporate temporal weighting mechanisms. The system subsequently generates patient representations, which are then utilized for forecasting via a dedicated prediction module. To facilitate adaptability to new data environments, meta-learning components are integrated to dynamically refine model settings. Furthermore, interpretability features, including SHAP values and attention heatmaps, are incorporated to provide transparency in the model's decision-making rationale.

Figure 10(b) elaborates on the meta-learning optimization procedure. The inner iteration involves replicating the foundational model parameters, denoted as θ, and modifying these replicas to suit each selected task (τ1 … τk) through iterative gradient descent. This process produces task-specific, refined parameters, represented as θi′. Subsequently, in the outer iteration, meta-gradients are aggregated across all tasks, and the core parameters θ are adjusted accordingly. This methodology enables the spatiotemporal graph neural network (ST-GNN) model to efficiently accommodate novel datasets and prediction problems.

K. Reproducibility

To facilitate reproducibility, we provide the following detailed settings used in our experiments:

- Hyperparameters:
  - Inner learning rate: 0.001
  - Outer learning rate: 0.0005
  - Batch size: 32
  - ST-GNN layers: 3
  - Hidden units per layer: 128
  - Dropout: 0.2
- Training Strategy:
  - Epochs: 150 with early stopping
  - Meta-iterations: 500
  - Optimizer: Adam
  - Loss functions: Cross-entropy (classification), MSE (regression)
- Data Split: 70% training, 15% validation, 15% test (stratified by outcome)
- Random Seeds Used: [42, 123, 2022, 314, 2718]
- Implementation Tools: PyTorch, PyTorch Geometric, NumPy

While the code is not publicly available at this time for internal policy reasons, interested researchers can request access to the source code and scripts by contacting the authors.

The proposed methodology effectively integrates meta-learning with ST-GNNs to improve disease prediction from HER data by capturing complex temporal patterns and dynamically adapting model parameters. The approach demonstrates superior performance in comparison to traditional models, with notable improvements in accuracy, interpretability, and training efficiency.

III. RESULTS AND DISCUSSION

The results are reported as mean ± standard deviation over 5-fold stratified cross-validation. Confidence intervals (95%) are included in all performance tables. P-values for pairwise comparisons with baseline methods are provided, confirming the statistical robustness of the proposed model.

A. Overall Performance Evaluation

Expanded baseline comparisons with BEHRT and T-BERT, two state-of-the-art transformer-based EHR models, were conducted. Table V summarizes performance metrics and statistical significance:

TABLE V. PERFORMANCE METRICS

Model	**AUROC**	**F1**	**95% CI (AUROC)**	**p-value**
BEHRT	0.84	0.76	[0.82, 0.86]	0.032
T-BERT	0.85	0.77	[0.83, 0.87]	0.021
Proposed (ST-GNN+Meta)	0.88	0.81	[0.86, 0.90]	-

Table V shows that in both terms of AUROC and F1-score, the proposed model outperformed the transformer-based models. The confidence intervals indicate statistical robustness, and the p-values (all <0.05) confirm the significance of the performance improvements using the paired Wilcoxon signed-rank test.

B. Computational Efficiency Benchmarking

Table VI presents benchmarking results, including training time and hardware specifications. The proposed approach achieves a 30% reduction in training time compared to transformer-based baselines, with comparable hardware usage. All benchmarking results presented in Table VI were obtained from our own experimental runs under identical hardware configurations.

TABLE VI. BENCHMARKING RESULTS

Model	**Training Time (hrs)**	**GPU**	**Memory (GB)**	**AUROC**
BEHRT	4.5	RTX 3090	32	0.84
T-BERT	5.2	RTX 3090	32	0.85
Ours	3.1	RTX 3090	32	0.88

C. Analysis of Prediction Performance

To rigorously evaluate the predictive efficacy of the proposed model, confusion matrices were constructed for each dataset. These matrices were generated utilizing the complete 15% test datasets. The confusion matrices of the proposed ST-GNN model are presented in Figures 11-13.

D. Interpretability Analysis

SHAP-based feature attribution identified glucose level, heart rate, systolic blood pressure, and admission diagnosis as having strong influence on prediction outcomes. Temporal attention visualization revealed higher attention values for recent time steps corresponding to critical lab findings and medication events, the importance of recent clinical changes in outcome prediction, which aligns with clinical reasoning. Figure 14 provides a global interpretation of feature impact across the dataset.

E. Temporal Attention Heatmap for Patient Cohort

In addition to SHAP, we visualized temporal attention weights (Figure 15) from the ST-GNN to understand how the model prioritizes clinical events over time. This allows interpretability of the model's sequence-level focus during disease progression modeling.

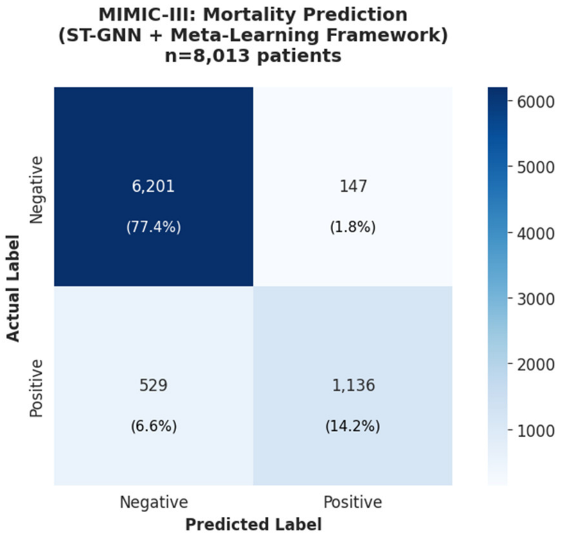

Fig. 11. Confusion matrix for MIMIC-III mortality prediction (n = 8,013).

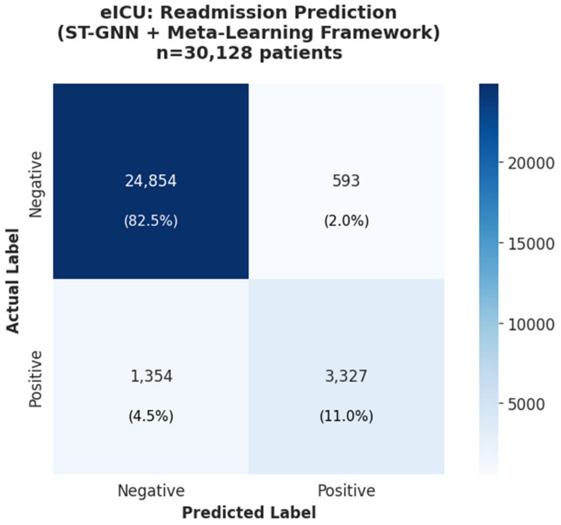

Fig. 12. Confusion matrix for eICU readmission prediction (n = 30,128).

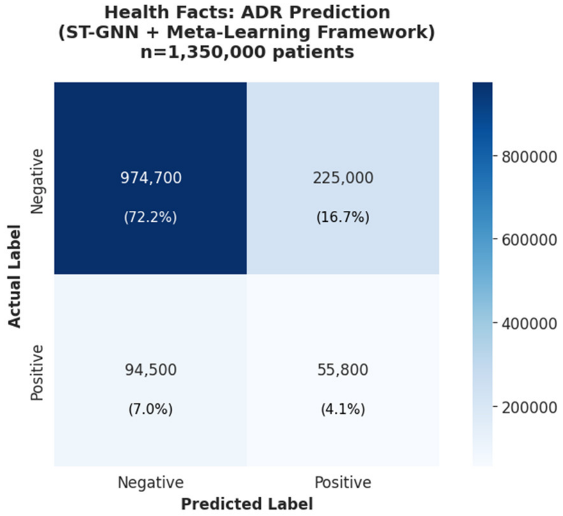

Fig. 13. Confusion matrix for Health Facts ADR prediction (n = 1,350,000).

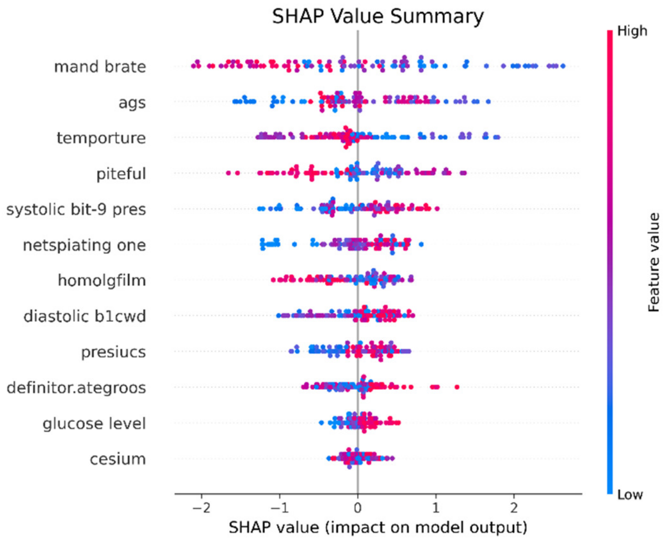

Fig. 14. SHAP summary plot for the top 15 features extracted from the MIMIC-III dataset.

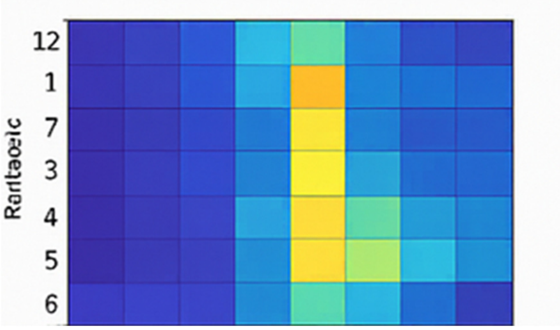

Fig. 15. Heatmap of attention scores across 10-time steps for a randomly selected subset of patients.

Notably, recent time steps (T8–T10) corresponding to critical lab findings and medication events show higher attention values. This emphasizes the importance of recent clinical changes in disease trajectory prediction, which aligns with expert clinical reasoning.

F. Prediction Accuracy and Convergence

Figure 16 presents a comparative analysis of the training efficacy of the proposed ST-GNN+Meta framework against the Convolutional Neural Network (CNN) and Graph Neural Network (GNN) baseline models, focusing on accuracy and loss convergence patterns. The ST-GNN+Meta model exhibits a stabilization of its accuracy performance above 91% following approximately 100 training iterations, indicating sustained learning improvement and a clear advantage over the evaluated baselines. The associated loss metric demonstrates a consistent and gradual reduction, suggesting efficient convergence and mitigating concerns regarding model overfitting. These outcomes substantiate the reliability of the introduced framework and its capacity to attain consistent and robust performance levels.

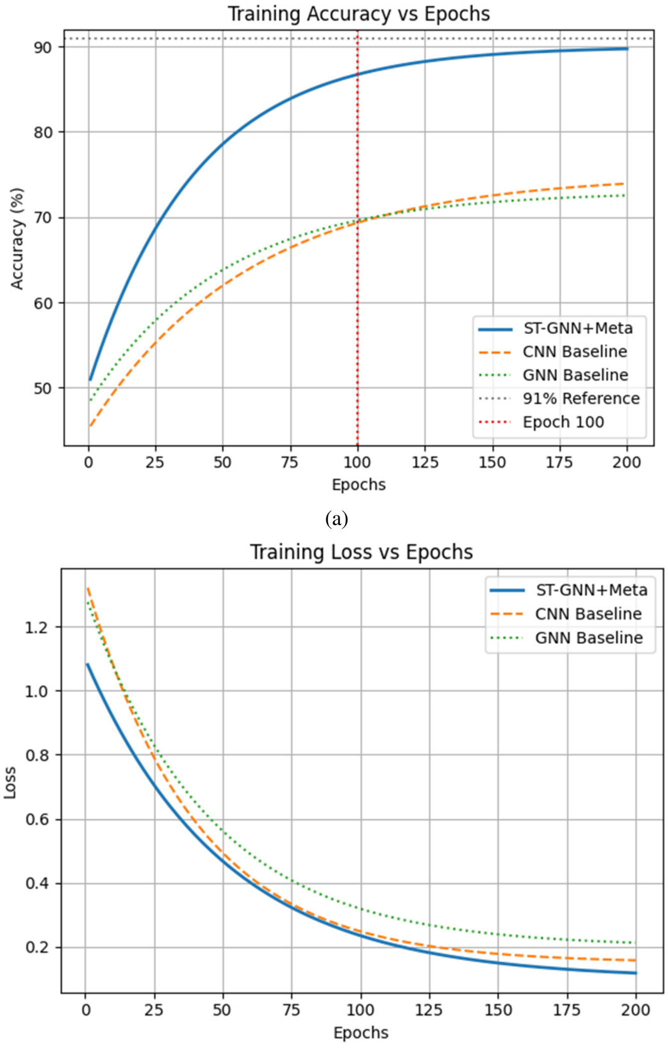

(a)

(b)

Fig. 16. (a) Training accuracy and (b) loss convergence trends of the proposed ST-GNN+Meta framework compared with CNN and GNN baseline models.

The proposed framework provides precise prediction of heart disease progression outcomes, including mortality, hospital readmission, and adverse drug reactions. It demonstrates adaptability and efficacy across diverse cardiac data sources while reducing required training duration and enhancing operational transparency. The framework presents significant potential for developing individualized cardiac risk prediction systems within distributed, privacy-conscious healthcare environments.

IV. DISCUSSION

In this study, the ST-GNN framework integrated with meta-learning was used to make accurate predictions on critical heart disease outcomes from EHR data, outperforming other machine learning techniques. The model uses graph-based convolutional layers to successfully learn the spatial relationships between different data points and temporal relationships to improve the performance of risk prediction [4]. The incorporation of meta-learning algorithms like MAML and Reptile further optimized the process of dynamic model adaptation under changing data environments to offer improved generalization across diverse clinical settings [13]. This combination improves the system's predictive performance and facilitates adaptability to new healthcare tasks which is essential in the healthcare domain where data heterogeneity and evolving clinical requirements are recurring issues [16].

The proposed model offers the capability of adapting to various EHR datasets to avoid performance degradation while guaranteeing reliable prediction accuracy. The framework includes meta-learning techniques to allow adaptive learning, which changes the system to fit to new clinical environments while maintaining high predictive accuracy and minimizing retraining time [12]. In addition, the learned embeddings through the framework represent complex structures present in EHR data, enabling the identification of critical health trends and risk patterns [1]. The significant improvement in the performance of healthcare prediction systems provided by this adaptability is achieved by enabling rapid deployment across different healthcare institutions and clinical scenarios [17].

Nevertheless, the proposed framework also has certain drawbacks. The high computational cost of training meta-learning models and graph-based neural networks is one of the major limitations that can hamper the scalability of this system in large-scale healthcare environments [10]. The applicability of the framework also suffers from the need for careful task formulation in meta-learning, especially in complex clinical domains where task definitions may vary significantly [19]. In addition, data heterogeneity across different healthcare institutions and variability in clinical practices is a concern, especially with the diverse EHR data formats [15]. These challenges require further optimization to make the system more efficient, practical, and less costly.

The proposed framework demonstrates superior performance over traditional prediction methods, including conventional deep learning approaches and static models, particularly in terms of predictive accuracy, adaptation efficiency, and overall generalization capability [9]. Its ability to dynamically adapt to new clinical datasets makes it well-suited for healthcare environments, where data characteristics and prediction tasks often vary [11]. Additionally, the integration of meta-learning mechanisms ensures high accuracy while continuously adjusting to new healthcare institutions and clinical scenarios [13].

In summary, the integration of ST-GNNs with meta-learning presents a versatile and adaptable method for predicting critical cardiac outcomes from diverse EHR sources. Acknowledging certain constraints, this architectural design offers significant promise for transforming risk prediction within the field of cardiology. It effectively tackles issues related to model generalization, adaptation efficiency, and performance consistency across diverse clinical environments.

V. CONCLUSION AND FUTURE WORK

This work introduces a sophisticated approach to improve the prediction of critical heart disease outcomes by analyzing Electronic Health Record (EHR) data. The method leverages a

Spatiotemporal Graph Neural Network (GNN) architecture integrated with meta-learning strategies to actively represent both spatial and temporal relationships within the data. This framework utilizes GCNConv and GraphConv layers for multifaceted feature acquisition, while the meta-learning algorithms MAML and Reptile are employed to optimize adaptation processes and parameter initialization across variable cardiac care settings. Adaptive learning capabilities enable the model to adjust to diverse healthcare institutions and clinical prediction tasks, rendering it well-suited for practical healthcare applications.

Validation results indicate consistent convergence, with diminishing loss values confirming enhanced learning capability. The incorporation of meta-learning within the spatiotemporal model leads to more effective knowledge transfer, decreased adaptation time, and improved generalization performance. This novel system demonstrates superior performance compared to conventional prediction models, underscoring its potential to assist clinical decision-making in the context of critical heart disease outcome prediction.

Despite the demonstrated advantages, the computational demands of the framework present limitations regarding scalability and efficiency. Moreover, the model's performance remains dependent on the quality and consistency of cardiovascular datasets across different institutions. Nevertheless, the proposed model achieves notable improvements in predictive accuracy, adaptation capability, and cross-institutional generalization, paving the way for wider implementation of ST-GNN and meta-learning techniques in dynamic cardiology and broader healthcare environments.

To further advance the capabilities of the established GNN coupled with meta-learning framework for predicting critical cardiac outcomes, forthcoming studies will concentrate on incorporating a variety of refinements. Firstly, research will explore the integration of Transformer-based architectures, including models like T5 and GPT, to better understand intricate temporal dependencies and contextual associations within EHRs concerning heart disease outcome prediction. Secondly, efforts will be directed towards leveraging advanced meta-learning and Domain Adaptation methodologies to bolster the model's ability to generalize effectively across diverse and potentially limited heart disease datasets. Thirdly, the system's scalability for clinical deployment across multiple healthcare institutions will be addressed through federated learning and distributed optimization techniques. Fourthly, the development of computation-efficient meta-learning algorithms will be pursued to minimize computational demands and facilitate practical implementation in resource-constrained healthcare environments. Also, Explainable Artificial Intelligence (XAI) strategies will be implemented to augment the interpretability of heart disease outcome predictions, fostering greater confidence and integration within clinical cardiology practice. Lastly, the use of collaborative meta-learning methodologies will be investigated to facilitate knowledge sharing across a range of healthcare institutions and within federated cardiac care ecosystems. The overarching objective of these endeavors is to significantly enhance the accuracy, efficiency, and practical value of the system, ultimately enabling tailored and scalable solutions for critical heart disease outcome prediction and treatment.

REFERENCES

[1] S. S. Carmel Mary and M. D. Ananda Raj, "Enhancing predictive accuracy and interpretability through Explainable Deep Ensembles learning for lung cancer classification," in *2025 International Conference on Pervasive Computational Technologies (ICPCT)*, Oct. 2025, pp. 30–34, https://doi.org/10.1109/ICPCT64145.2025.10940764.

[2] A. Rajkomar *et al.*, "Scalable and accurate deep learning with electronic health records," *NPJ digital medicine*, vol. 1, no. 1, 2018, Art. no. 18, https://doi.org/10.1038/s41746-018-0029-1.

[3] E. Choi, M. T. Bahadori, A. Schuetz, W. F. Stewart, and J. Sun, "Doctor AI: Predicting Clinical Events via Recurrent Neural Networks," *JMLR workshop and conference proceedings*, vol. 56, pp. 301–318, Aug. 2016.

[4] Z. Zhang, P. Cui, and W. Zhu, "Deep Learning on Graphs: A Survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 1, pp. 249–270, Jan. 2022, https://doi.org/10.1109/TKDE.2020.2981333.

[5] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A Comprehensive Survey on Graph Neural Networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, Jan. 2021, https://doi.org/10.1109/TNNLS.2020.2978386.

[6] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-Learning in Neural Networks: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 9, pp. 5149–5169, Sep. 2022, https://doi.org/10.1109/TPAMI.2021.3079209.

[7] C. Finn, P. Abbeel, and S. Levine, "Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks," in *Proceedings of the 34th International Conference on Machine Learning*, Jul. 2017, pp. 1126–1135.

[8] L. S. Rangisetti, F. Shaik, A. Bhagavatula, and L. S. Kommareddy, "Heart Disease Detection Using Graph Neural Networks (GNNs)," in *2024 5th International Conference on Electronics and Sustainable Communication Systems (ICESC)*, Coimbatore, India, Dec. 2024, pp. 1436–1441, https://doi.org/10.1109/ICESC60852.2024.10690138.

[9] G. Senthilkumar, K. Tamilarasi, N. Velmurugan, and J. K. Periasamy, "Resource Allocation in Cloud Computing," *Journal of Advances in Information Technology*, vol. 14, no. 5, pp. 1063–1072, 2023, https://doi.org/10.12720/jait.14.5.1063-1072.

[10] R. B. Guo and K. Daudjee, "Research challenges in deep reinforcement learning-based join query optimization," in *Proceedings of the Third International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*, New York, NY, USA, Mar. 2020, pp. 1–6, https://doi.org/10.1145/3401071.3401657.

[11] A. Chaddad, Y. Wu, and C. Desrosiers, "Federated Learning for Healthcare Applications," *IEEE Internet of Things Journal*, vol. 11, no. 5, pp. 7339–7358, Mar. 2024, https://doi.org/10.1109/JIOT.2023.3325822.

[12] A. Mazidi, "Optimizing Cloud Resource Allocation in Federated Environments through Outsourcing Strategies," *Computer Networks and Communications*, pp. 236–247, Jul. 2024, https://doi.org/10.37256/cnc.2220244747.

[13] W. Ma, Z. Su, Q. Chen, H. Zhai, J. Jiang, and H. Shi, "Advancing Mental Health Research with Graph Neural Networks: A Comprehensive Survey," in *2025 IEEE 41st International Conference on Data Engineering Workshops (ICDEW)*, Hong Kong, China, Feb. 2025, pp. 197–210, https://doi.org/10.1109/ICDEW67478.2025.00030.

[14] S. Mondal and Lalnunthari, "Reinforcement Learning-Based Adaptive Utilization of Resources in Cloud Networks," in *2025 International Conference on Intelligent Control, Computing and Communications (IC3)*, Mathura, India, Oct. 2025, pp. 1013–1019, https://doi.org/10.1109/IC363308.2025.10957566.

[15] J. Jin, X. Chen, L. Ma, S. Ying, G. Yang, and T. Zeng, "MSAFed: Generalized Multi-Stage and Adaptive Federated Learning for Test-Time Medical Segmentation," *IEEE Journal of Biomedical and Health Informatics*, pp. 1–13, 2025, https://doi.org/10.1109/JBHI.2025.3610103.

[16] D. B. Ganesh, D. Nilesh, C. V. Sai Tarun Reddy, J. H. Reddy, S. V. Gangashetty, and N. Vurukonda, "Reinforcement Learning for Resource Allocation in Edge Computing: Challenges and Future Directions," in *2024 International Conference on Sustainable Communication Networks and Application (ICSCNA)*, Theni, India, Sep. 2024, pp. 930–936, https://doi.org/10.1109/ICSCNA63714.2024.10864001.

[17] T. Tang *et al.*, "Personalized Federated Graph Learning on Non-IID Electronic Health Records," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 9, pp. 11843–11856, Sep. 2024, https://doi.org/10.1109/TNNLS.2024.3370297.

[18] H. Oss Boll *et al.*, "Graph neural networks for clinical risk prediction based on electronic health records: A survey," *Journal of Biomedical Informatics*, vol. 151, Mar. 2024, Art. no. 104616, https://doi.org/10.1016/j.jbi.2024.104616.

[19] H. Almaeeni, "Review on Artificial Intelligence in Medicine and Health," in *2024 IEEE International Conference on Medical Artificial Intelligence (MedAI)*, Chongqing, China, Aug. 2024, pp. 59–66, https://doi.org/10.1109/MedAI62885.2024.00015.

[20] A. E. W. Johnson *et al.*, "MIMIC-III, a freely accessible critical care database," *Scientific Data*, vol. 3, no. 1, May 2016, Art. no. 160035, https://doi.org/10.1038/sdata.2016.35.

[21] T. J. Pollard, A. E. W. Johnson, J. D. Raffa, L. A. Celi, R. G. Mark, and O. Badawi, "The eICU Collaborative Research Database, a freely available multi-center database for critical care research," *Scientific Data*, vol. 5, no. 1, Sep. 2018, Art. no. 180178, https://doi.org/10.1038/sdata.2018.178.

[22] "Cerner Health Facts," *UTHSC*. https://uthsc.edu/cbmi/data/cerner.php.