MINISTRY OF EDUCATION
AND TRAINING

VIETNAM ACADEMY OF
SCIENCE
AND TECHNOLOGY

GRADUATE UNIVERSITY OF SCIENCE AND TECHNOLOGY

Pham Thanh Dat

RESEARCH ON METHODS FOR SHORT-TERM FORECASTING AND DECISION SUPPORT IN HIV SURVEILLANCE AND TREATMENT USING DEEP LEARNING

SUMMARY OF DISSERTATION ON INFORMATION SYSTEM
Code: 9 48 01 04

Hà Nội – 2026

This dissertation was completed at the Graduate University of Science and Technology, Vietnam Academy of Science and Technology

Scientific Supervisors:

1. Assoc. Prof. Dr. Nguyen Truong Thang Institute of Information Technology, Vietnam Academy of Science and Technology
2. Assoc. Prof. Dr. Nguyen Viet Anh Institute of Information Technology, Vietnam Academy of Science and Technology

Reviewer 1: _____

Reviewer 2: _____

Reviewer 3: _____

The dissertation will be defended before the Academy-level Doctoral Dissertation Evaluation Committee at the Graduate University of Science and Technology, Vietnam Academy of Science and Technology at …… (time), on …… (day) …… (month), …… (year).

INTRODUCTION

Despite substantial progress in HIV prevention and treatment, HIV remains a major public health challenge. Unlike acute infectious diseases, HIV is a chronic disease that progresses over time and requires continuous surveillance, treatment, and intervention through multiple levels of decision-making, ranging from epidemiological surveillance and treatment management to intervention evaluation. In practice, HIV systems generate large volumes of heterogeneous data, creating a growing demand for intelligent analytical and forecasting methods to support timely and effective decision-making at different management levels.

This dissertation investigates deep learning methods for three major tasks in HIV surveillance and treatment, including forecasting future HIV infections, predicting disease progression risk, and estimating causal effects for decision support. The objective is to develop deep learning models capable of comprehensively exploiting HIV data to improve surveillance, prognosis, and intervention. The research methodology combines theoretical analysis with experimental evaluation using HIV surveillance data, clinical data, and survey data.

The main contributions of this dissertation include: (1) a spatio-temporal graph-based forecasting model for capturing spatial and temporal relationships among HIV cases; (2) a Mamba-based survival analysis model; and (3) a causal inference framework that integrates deep learning and causal estimation framework with adaptive coordination through reinforcement learning. These contributions provide a data-driven foundation for decision support in HIV management.

Chapter 1. THEORETICAL BACKGROUND

This chapter presents the theoretical foundations of modern machine learning approaches for three main research areas: short-term forecasting, survival analysis, and causal effect estimation. It first reviews deep learning models for time-series and spatio-temporal data, including RNNs, LSTMs, GRUs, CNNs, GCNs, GATs, ST-GNNs, Transformers, and advanced architectures such as TITAN and Mamba for long or irregular time series. The chapter then introduces the mathematical foundations of survival analysis and representative deep learning models, including DeepSurv, DeepHit, and SurvTRACE, for handling censored data and nonlinear relationships between patient characteristics and event risk. Causal inference is reviewed through directed acyclic graphs (DAGs), the potential outcomes framework, and models such as TARNet, DragonNet, and DR Learner for estimating individualized treatment effects, together with fundamental concepts of causal graph construction and uncertainty in medical forecasting. Finally, evaluation metrics are presented to assess model accuracy, discrimination, calibration, and reliability across the three research tasks.

Chương 2. FORECASTING HIV INFECTIONS USING AN HIV CASE GRAPH

1. Introduction

This chapter proposes a spatiotemporal graph model based on the connectivity among HIV cases, integrating a graph pooling mechanism to transfer propagation knowledge learned at the city level to short-term forecasting at the district level, thereby improving prediction accuracy.

2. Proposed Method

2.1. City-level spatio-temporal representation of HIV-positive case graphs

The HIV-positive case graph is represented as $G = (V, E)$, where each node v_i is associated with a feature vector $\mathbf{x}_i \in \mathbb{R}^{d_x}$, and each edge e_{ij} is associated with an edge feature vector $\mathbf{f}_{ij} \in \mathbb{R}^{d_f}$ describing the epidemiological relationship between two HIV-positive cases. The epidemiological association between two cases is quantified by following.

$$\rho_{ij} = \xi_{ij} \times \tau_{ij} \times \kappa_{ij}. \tag{2.1}$$

The spatial correlation is defined as: $\xi_{ij} = \begin{cases} 1, & \text{same area,} \\ 1 - \dfrac{d_{ij}}{\max(d_{ij})}, & \text{otherwise,} \end{cases}$

where $d_{ij} = \sqrt{(\text{long}_i - \text{long}_j)^2 + (\text{lat}_i - \text{lat}_j)^2}$.

The temporal correlation is defined as: $\tau_{ij} = \begin{cases} 1, & r_i = r_j, \\ 1 - \dfrac{|r_i - r_j|}{\max(r_{ij})}, & \text{otherwise.} \end{cases}$

The transmission-route correlation is defined as $\kappa_{ij} = \begin{cases} 1, & \text{same route,} \\ \dfrac{L_i + L_j}{L}, & \text{otherwise.} \end{cases}$

An edge is created only when $\rho_{ij} > \beta$, thereby controlling the graph density while preserving only epidemiologically meaningful relationships.

2.2. Forecasting framework

The proposed framework (fig 2.1) constructs a city-level spatio-temporal graph from HIV-positive cases and uses an attention-based GNN to learn node representations. Regional subgraphs are extracted and pooled to retain informative cases, followed by global mean pooling and an MLP for regional forecasting.

2.3. Embedding node and edge features into a shared representation space

The proposed model employs an MLP to project node and edge features into a common embedding space with dimension d_{em}:

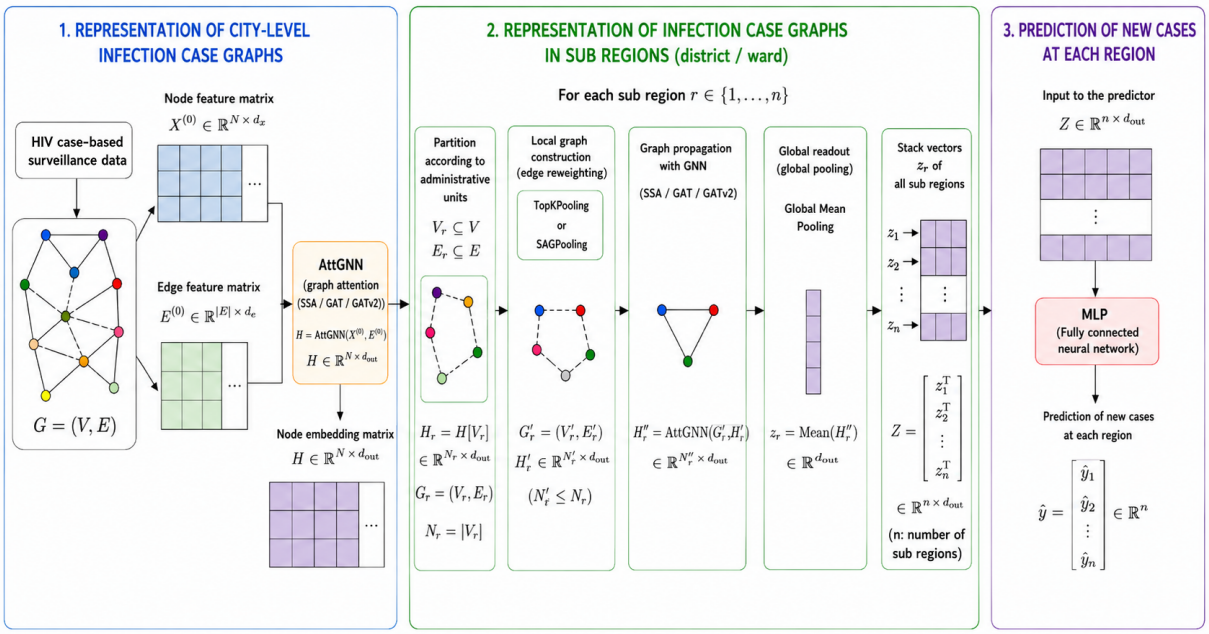

Figure 2.1: Overall architecture of the proposed framework.

$\text{FeatureEMBED}(v_i) = \text{MLP}(x_i, d_{em}), \text{FeatureEMBED}(e_{ij}) = \text{MLP}(f_{ij}, d_{em})$. Here, $x_i \in \mathbb{R}^{d_x}$ denotes the epidemiological feature vector of node v_i, whereas $f_{ij} \in \mathbb{R}^{d_f}$ represents the spatio-temporal features of edge e_{ij}.

2.4. Attention-based graph neural network for the city-level graph

The graph self-attention mechanism enables the model to learn the propagation influence among HIV-positive cases through adaptive attention weights rather than assuming equal contributions from neighboring nodes. Given the input feature matrix $D = [d_1, \ldots, d_n]$, the attention mechanism is computed as

$$SCORE_{att} = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right), \qquad Y = SCORE_{att} \cdot V, \tag{2.2}$$

where $Q = DW_q^T, K = DW_k^T, V = DW_v^T$. The multi-head attention mechanism is defined as $\text{Multihead}(Q, K, V) = \text{CONCAT}(Y^1, \ldots, Y^C)W_o$. To incorporate spatio-temporal information, edge features are integrated into the attention weights as

$$\alpha_{c,ij} = \frac{(q_{c,i}, k_{c,j} + f_{c,ij})}{\sum_{u \in \mathcal{N}(i)} (q_{c,i}, k_{c,u} + f_{c,iu})}, \tag{2.3}$$

where $(q, k) = \exp\left(\frac{q^\top k}{\sqrt{d_k}}\right)$. The node representation is updated as follows.

$$h_i' = \Big\|_{c=1}^{C} \left[\sum_{j \in \mathcal{N}(i)} \alpha_{c,ij}(v_{c,j} + e_{c,ij}) \right], \qquad v_{c,j} = W_{c,v} h_j + b_{c,v}. \tag{2.4}$$

3. Graph Attention Network (GAT) with edge features.

In the GAT model, the attention coefficient between the target node v_i and one of its neighboring nodes v_j (including the node itself) is computed from both node features and edge features. Let $h_i = W_i x_i, h_j = W_j x_j, h_{ij} = W_{ij} e_{ij}$. The attention coefficient is then computed as follows.

$$\alpha_{ij} = \frac{\exp\left(\text{LeakyReLU}\left(\mathcal{W}_{a,i}^{\top} h_i + \mathcal{W}_{a,j}^{\top} h_j + \mathcal{W}_{a,ij}^{\top} h_{ij}\right)\right)}{\sum\limits_{k \in \mathcal{N}(i) \cup \{i\}} \exp\left(\text{LeakyReLU}\left(\mathcal{W}_{a,i}^{\top} h_i + \mathcal{W}_{a,k}^{\top} h_k + \mathcal{W}_{a,ik}^{\top} h_{ik}\right)\right)}, \tag{2.5}$$

and the multi-head output representation is calculated as follows.

$$h_i' = \frac{1}{C} \sum_{c=1}^{C} \sum_{j \in \mathcal{N}(i) \cup \{i\}} \alpha_{ij}^{(c)} h_j. \tag{2.6}$$

4. Graph Attention Network v2 (GATv2).

GATv2 improves the flexibility of the attention mechanism by applying the nonlinear transformation before projection through the attention weights. The attention coefficient is defined as follows.

$$\alpha_{ij} = \frac{\exp\left(\mathcal{W}_a^{\top} \text{LeakyReLU}(h_i + h_j + h_{ij})\right)}{\sum\limits_{k \in \mathcal{N}(i) \cup \{i\}} \exp\left(\mathcal{W}_a^{\top} \text{LeakyReLU}(h_i + h_k + h_{ik})\right)}. \tag{2.7}$$

2.5. Construction of regional subgraphs from the city-level HIV-positive case graph

For each region $r \in \{1, 2, \ldots, n\}$, a regional subgraph $G_r = (V_r, E_r)$ is constructed, where $V_r = \{v_{r1}, \ldots, v_{rL_r}\}$ denotes the set of HIV-positive cases recorded in region r, and $L_r = |V_r|$.

Node representations in each regional subgraph are inherited from the city-level graph to preserve inter-regional information, while the edge set E_r is reconstructed using the same thresholding criterion, with an edge created when the epidemiological correlation coefficient satisfies $\rho_{ij} > \beta$.

2.6. Graph pooling for regional subgraphs

The task of forecasting new HIV infections is formulated as a graph-level regression problem. Graph pooling (node-drop pooling) is employed to simplify each regional graph, suppress noisy nodes, and retain the HIV-positive cases

that contribute most significantly to epidemic propagation. The pooling procedure consists of three stages: (i) node scoring, (ii) top-k node selection, and (iii) graph reconstruction.

1. TopKPooling. For a regional graph $G_r = (X_r, A_r)$, the node scores are computed as $S_r = \frac{X_r p}{\|p\|}, IDX_r = \text{RANK}(S_r, k), \widetilde{S}_r = \text{sigmoid}(S_r(IDX_r))$. The node features and adjacency matrix are updated as follows.

$$X'_r = X_r[IDX_r, :] \odot (\widetilde{S}_r \mathbf{1}_{d_x}^\top), A'_r = A_r[IDX_r, IDX_r]. \tag{2.8}$$

2. SAGPool. The node scores are computed using a graph convolutional network (GCN) as $S_r = \sigma\left(\widetilde{D}_r^{-\frac{1}{2}} \widetilde{A}_r \widetilde{D}_r^{-\frac{1}{2}} X_r W_{\text{SAGPool}}\right)$, after which the top $\lfloor kL_r \rfloor$ nodes are retained to construct the pooled graph G'_r.

2.7. Representation Update and Graph-level Aggregation

After graph pooling, an attention-based GNN aggregates local information from the retained informative cases to capture regional transmission dynamics. To obtain a fixed-length representation for each regional subgraph G_r, global mean pooling is employed as: $\text{output}(G_r) = \frac{1}{L_r} \sum_{i=1}^{L_r} x_i, \forall r \in \{1, 2, \ldots, n\}$.

2.8. Prediction Layer and Model Output

The graph-level representations of all regions are concatenated into a matrix $\text{output}(G) \in \mathbb{R}^{n \times d_{em}}$. This matrix is subsequently fed into a multilayer perceptron (MLP) to perform graph-level regression and predict the number of newly reported HIV-positive cases for each region: $\text{PredictOutput} = \text{MLP}(\text{output}(G), 1)$.

3. Experiments and Results

3.1. Dataset

The study used retrospective, fully anonymized HIV surveillance data from 40,611 cases across 23 districts in Ho Chi Minh City between 2008 and 2019. The data were represented as graphs, where each case was modeled as a node with encoded epidemiological features, and edges were defined based on spatial proximity and feature similarity. Graphs were constructed at three-month intervals to support short-term HIV forecasting.

3.2. Baseline Models

The proposed model was experimentally compared with ARIMA, LSTM, Bi-LSTM, CNN, Stacked-LSTM, CNN-LSTM, DCRNN, MPNN-LSTM, GAT-LSTM, GCN-LSTM, and CGN-GRU.

3.3. Experimental Results

The study evaluated combinations of GSA, GAT, and GATv2 with TopKPooling and SAGPool, using GIN as a baseline; the results are presented in Table 2.1.

The results indicate that the combination of GSA for learning the city-level

Table 2.1: Selection of the optimal combination of GNN architecture and graph pooling method

	GSA		GATv2		GAT		GIN	
	MSE	R^2	**MSE**	R^2	**MSE**	R^2	**MSE**	R^2
TopKPooling + GAT + Global Mean Pooling	**21.64**	**0.558**	23.29	0.525	35.45	0.240	41.43	0.112
SAGPool + GAT + Global Mean Pooling	inf	inf	inf	inf	46.02	0.061	inf	inf
Global Mean Pooling	25.90	0.444	26.48	0.443	24.98	0.464	42.81	0.081

graph representation together with TopKPooling and GAT for regional subgraphs achieves the best forecasting performance, yielding the lowest mean squared error (MSE) and the highest coefficient of determination (R^2).

Selection of the TopKPooling parameter k. To evaluate the effect of the node retention ratio in TopKPooling, experiments were conducted with k values ranging from 0.1 to 0.9. The results show that the TopKPooling parameter has a substantial impact on the performance of both the GSA and GATv2 models. For GSA, the optimal performance was achieved at $k = 0.3$, yielding an MSE of 21.64 and an R^2 of 0.558. In contrast, GATv2 achieved its best performance at $k = 0.4$, with an MSE of 23.29 and an R^2 of 0.525. Accordingly, these values were selected for all subsequent experiments.

Analysis of the impact of the post-pooling GNN architecture. Replacing the GAT layer with GCN after the pooling stage led to a significant decline in predictive performance, highlighting the importance of the attention mechanism for learning the spatiotemporal dynamics of HIV transmission.

Comparison with existing forecasting methods. After determining the optimal model configuration, the proposed method was compared with a variety of representative forecasting approaches. As shown in Table 2.2, the proposed model consistently achieves the lowest mean squared error (MSE) and the highest coefficient of determination (R^2), demonstrating the effectiveness of learning HIV transmission patterns through a city-level HIV-positive case graph.

Table 2.2: Comparison between the proposed model and baseline methods

Method	**Proposed**	**MPNN LSTM**	**GAT LSTM**	**DC RNN**	**GCN GRU**	**GCN LSTM**	**LSTM**	**Bi-LSTM**	**Stacked LSTM**	**CNN**	**CNN-LSTM**	**ARIMA**
MSE	**21.64**	35.11	50.15	52.78	52.32	60.53	48.51	45.99	47.06	46.27	52.22	3944.64
R^2	**0.558**	0.244	-0.008	-0.136	-0.126	-0.303	0.010	0.010	-0.013	0.003	-0.120	-0.768

Analysis of epidemiological features. Ablation results showed that location, occupation, and transmission route improved forecasting performance, whereas age group had little effect. Removing the risk-group feature reduced MSE from 21.64 to 19.52, suggesting a potential negative effect on model performance.

Forecasting performance over different prediction horizons. The results in Table 2.3 show that the proposed model achieves its best performance for short-term forecasting (15 days and 1 month).

Table 2.3: Forecasting performance under different prediction horizons

	15 Days	1 Month	2 Months	3 Months	1 Month (1-Month Input)	1 Month (2-Month Input)
MSE	11.00	19.52	38.83	41.48	22.43	22.64
R^2	0.784	0.602	0.369	0.128	0.557	0.553

Chương 3. RISK PREDICTION THROUGH SURVIVAL ANALYSIS IN HIV TREATMENT

1. Introduction

This chapter proposes a personalized survival analysis model based on Mamba, integrating dynamic memory and Bayesian components to learn representations of disease states, capture long-term dependencies, and quantify uncertainty, thereby improving prediction quality.

2. Proposed Method

2.1. Overview of the Proposed Architecture

Figure 3.1 illustrates an overview of the proposed architecture, which integrates specialized encoders, temporal memory, and an SSM model to learn representations of HIV patients, while simultaneously predicting survival and treatment outcomes within a conditional Bayesian framework, incorporating SHAP to enhance interpretability.

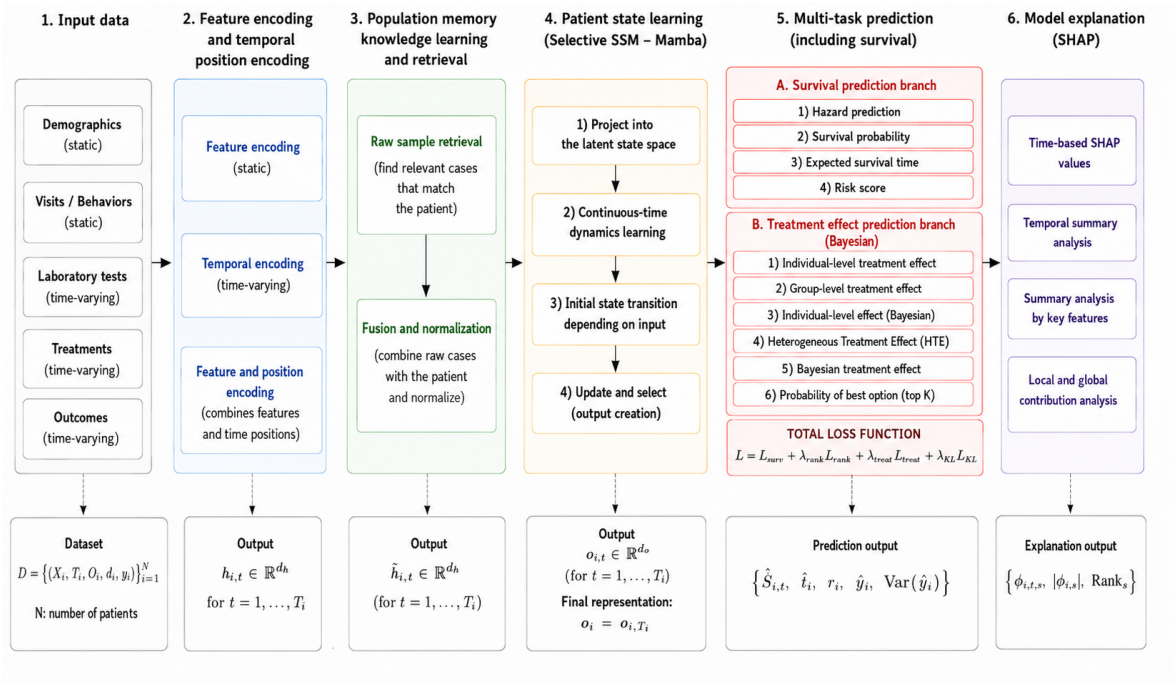

Figure 3.1: Overview of the proposed model architecture

2.2. Clinical Feature Encoding and Treatment Representation

Clinical features and treatment information are embedded separately and then combined into a unified input representation, enabling the model to simultaneously capture disease status and treatment effects over time.

2.3. Memory-based Population Knowledge Learning

After the feature encoder produces the patient representation $\mathbf{h}_{i,t}$, a memory-based population knowledge learning module is introduced to capture common disease progression patterns from the entire cohort. This module consists of three main components.

Global prototype memory bank. The model maintains a set of K learnable prototypes: $\mathbf{M} = [\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_K]^\top \in \mathbb{R}^{K \times d}$. The prototypes are randomly initialized and updated through backpropagation during training to represent common disease progression patterns within the population.

Time-aware cross-attention knowledge retrieval. The follow-up time information is first mapped into a time embedding: $\mathbf{e}_{i,t} = f_{\text{time}}(\tau_{i,t})$. The query vector is constructed from the patient representation and temporal information: $\mathbf{q}_{i,t} = \mathbf{W}_Q(\mathbf{h}_{i,t} + \mathbf{e}_{i,t})$. The key and value matrices are generated from the memory bank: $\mathbf{K} = \mathbf{M}\mathbf{W}_K, \mathbf{V} = \mathbf{M}\mathbf{W}_V$. The attention weights are computed using the scaled dot-product attention mechanism as follows. $\boldsymbol{\alpha}_{i,t} = \text{Softmax}\left(\frac{\mathbf{q}_{i,t}\mathbf{K}^\top}{\sqrt{d}}\right)$. The retrieved population knowledge representation is: $\mathbf{r}_{i,t} = \boldsymbol{\alpha}_{i,t}\mathbf{V}$. The attention mechanism enables the model to aggregate information from multiple prototypes rather than relying on a single prototype.

Fusion of population knowledge and individual representation. The patient representation and the retrieved knowledge are fused through an MLP with a residual connection: $\tilde{\mathbf{h}}_{i,t} = \mathbf{h}_{i,t} + f_{\text{fusion}}\left([\mathbf{h}_{i,t}; \mathbf{r}_{i,t}]\right)$. where $f_{\text{fusion}}(\cdot)$ is an MLP consisting of linear layers, activation functions, normalization, and dropout. The resulting representation $\tilde{\mathbf{h}}_{i,t}$ preserves patient-specific characteristics while incorporating knowledge learned from the entire population. This enhanced representation serves as the input to the disease-state learning module for temporal disease modeling and survival prediction.

2.4. Latent Disease-State Learning Using a Selective State Space Model

The memory-enhanced patient representation, denoted by $\tilde{\mathbf{h}}_{i,t}$, is fed into a Selective State Space Model (SSSM) based on the Mamba architecture to model disease progression over time with linear computational complexity with respect to sequence length. The input representation is first projected into the state space: $\mathbf{u}_{i,t} = \mathbf{W}_{\text{in}}\tilde{\mathbf{h}}_{i,t}$, after which the disease state is updated according to the selective state space model as $\mathbf{s}_{i,t} = \bar{\mathbf{A}}_{i,t}\mathbf{s}_{i,t-1} + \bar{\mathbf{B}}_{i,t}\mathbf{u}_{i,t}$, where the input-dependent state transition matrices are learned from the current representation as $\bar{\mathbf{A}}_{i,t} = \exp(\Delta_{i,t}\mathbf{A}), \bar{\mathbf{B}}_{i,t} = \Delta_{i,t}\mathbf{B}$. The model output is generated through a selective gating mechanism as $\mathbf{o}_{i,t} = \mathbf{W}_{\text{out}}\left(\mathbf{g}_{i,t} \odot \mathbf{y}_{i,t}\right)$. The representation at the final follow-up visit: $\mathbf{o}_i = \mathbf{o}_{i,T_i}$, is used as the overall patient representation and serves as the shared input to both the survival prediction branch and the treatment outcome prediction branch.

2.5. Multi-task Learning for Survival and Treatment Outcome Prediction

The disease-state representation $\mathbf{o}_i$ is used as the shared input for two prediction branches: survival prediction and treatment outcome prediction.

Survival prediction branch. The disease-state representation is used to predict the hazard function $\hat{\mathbf{h}}_i = f_{\text{hazard}}(\mathbf{o}_i)$,from which the survival probability is computed as follows.

$$\hat{S}_{i,t} = \prod_{k=1}^{t} (1 - \hat{h}_{i,k}). \tag{3.1}$$

Treatment outcome prediction branch. The treatment effect is modeled as the combination of population-level, subgroup-level, and individual-level effects as follows.

$$e_i = e_{\text{pop}} + e_{\text{sub},i} + e_{\text{ind},i}, \tag{3.2}$$

which is combined with the disease-state representation to predict treatment outcomes as follows.

$$\hat{y}_i = f_{\text{out}} \left([\mathbf{o}_i; \mathbf{e}_i; \text{emb}(d_i)] \right). \tag{3.3}$$

Multi-task learning. The two prediction branches are jointly optimized using the following loss function as follows.

$$\mathcal{L} = \mathcal{L}_{\text{surv}} + \lambda_{\text{rank}} \mathcal{L}_{\text{rank}} + \lambda_{\text{treat}} \left(\mathcal{L}_{\text{treat}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}} \right). \tag{3.4}$$

The multi-task learning strategy enables the two prediction tasks to share a common disease-state representation while optimizing each task-specific branch independently to improve predictive performance.

2.6. Interpretability of Predictions Using SHAP

After model training, SHAP is employed to quantify the contribution of each input feature to the predicted survival outcomes and treatment outcomes without modifying the model architecture. The SHAP value for the i-th feature is computed based on the Shapley value

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} \left[f(S \cup \{i\}) - f(S) \right]. \tag{3.5}$$

For longitudinal data, SHAP values are computed at each follow-up visit and aggregated over time to evaluate the average contribution of each feature:

$$\overline{|\phi_i|} = \frac{1}{T} \sum_{t=1}^{T} |\phi_{t,i}|. \tag{3.6}$$

The resulting SHAP values are used to identify risk factors, protective fac-

tors, and generate explanation reports that facilitate the interpretation of model predictions.

2.7. Model Training Strategy

The model is trained using a curriculum learning strategy with progressively increasing complexity. The model is first trained only for survival prediction, followed by the introduction of the ranking loss, and finally the treatment outcome prediction task is incorporated. Model parameters are optimized using AdamW through backpropagation, together with *ReduceLROnPlateau* and *early stopping* on the validation set to improve training stability and reduce overfitting.

3. Experiments and Results

3.1. Data Description

The experimental design and evaluation results of the proposed model were conducted on two publicly available AIDS Clinical Trials Group (ACTG) clinical trial datasets, namely ACTG 388 and ACTG 175. These two datasets represent two complementary scenarios in survival prediction: ACTG 388 is a longitudinal dataset that enables the progression of each patient to be tracked through repeated measurements of CD4 count and viral load at different time points. In contrast, ACTG 175 is a static dataset in which each patient is represented only by baseline characteristics collected at the beginning of the study.

3.2. Baseline Models for Comparison

The study selected baseline models representing different approaches to survival analysis, including the classical statistical model CoxPH, Cox-based deep learning models such as DeepSurv, the discrete-time model DeepHit, and the Transformer-based model SurvTRACE. For time-series data in ACTG 388, RNN-based models, including GRU-Survival and LSTM-Survival, were additionally included to exploit longitudinal information.

3.3. Experiments on ACTG 388 (Longitudinal Cohort)

Overall Performance on ACTG 388

According to Table 3.1, the proposed model outperformed GRU-Survival, achieving a higher C-index and lower IBS and Calibration MSE values. These results indicate substantial improvements in risk discrimination, predictive accuracy, and probability calibration.

Table 3.1: Key performance metrics on ACTG 388

Metric	Proposed Model	Best Baseline (GRU)	Improvement
C-index	0.824	0.762	+8.1%
Integrated Brier Score (IBS)	0.087	0.112	-22.3%
Calibration MSE	0.012	0.021	-42.9%
Time-dependent AUC (24 months)	0.801	0.745	+7.5%

Experiments on ACTG 388 showed that the proposed model effectively

leveraged longitudinal data, increasing the C-index from 0.770 to 0.824 and improving early prediction at 3–6 months. Compared with standard care, the model achieved a treatment adherence rate of 74.1% versus 61.8%, reduced the time to viral suppression from 22.7 to 16.3 weeks, and increased the 24-month CD4 gain from +124 to +168 cells/μL.

Performance Comparison with Baseline Models

Experimental results on the ACTG 388 dataset demonstrate that the proposed model achieves the best performance across all evaluation metrics. Specifically, it attains a C-index of 0.824, substantially outperforming competing methods, including GRU-Survival (0.762), LSTM Survival (0.756), SurvTRACE (0.739), DeepHit (0.708), DeepSurv (0.672), and Cox PH (0.634). The proposed model also achieves the lowest Integrated Brier Score (IBS) of 0.087 and the lowest Calibration MSE of 0.012, indicating superior survival probability estimation and calibration. Furthermore, it reaches an AUC of 0.801 at the 24-month time horizon, outperforming all baseline models. These results demonstrate that integrating longitudinal representation learning with a prototype memory module and a selective state-space model enables more effective modeling of temporal clinical dynamics, thereby substantially improving survival prediction performance on longitudinal HIV data.

Statistical Significance Assessment Using DeLong's Test

The DeLong test showed that the proposed model achieved a statistically significant improvement in AUC compared with Cox PH, DeepSurv, DeepHit, SurvTRACE, LSTM-Survival, and GRU-Survival.

Characteristics of the Study Population Stratified by Baseline CD4 Count

Table 3.2 shows clear survival stratification according to baseline CD4 count. The group with CD4 < 50 cells/μL had the highest risk (HR = 2.84; 95% CI: 2.12–3.81; $p < 0.001$) and the poorest prognosis compared with the group with CD4 > 200 cells/μL. This finding confirms the central prognostic role of CD4 count in HIV disease progression.

Table 3.2: Survival outcomes according to baseline CD4 stratification (ACTG 388)

CD4 Stratum	N	Proportion	6 Months	95% CI	12 Months	95% CI	24 Months	95% CI	Median Time
< 50 cells/μL	127	24.90%	0.732	(0.651–0.813)	0.583	(0.498–0.668)	0.441	(0.355–0.527)	14.2 months
50–200 cells/μL	226	44.20%	0.823	(0.769–0.877)	0.721	(0.661–0.781)	0.634	(0.571–0.697)	21.7 months
> 200 cells/μL	158	30.90%	0.886	(0.834–0.938)	0.823	(0.758–0.888)	0.759	(0.684–0.834)	Not reached

Ablation Study

Table 3.3 shows that all temporal modeling components improved predictive performance, with the Selective SSM making the largest contribution. When all components were fully integrated, the model achieved the best overall performance (C-index = 0.824, IBS = 0.087, Calibration MSE = 0.012), confirming

its effectiveness in exploiting the temporal structure of the ACTG 388 dataset.

Table 3.3: Ablation study: Effects of temporal components on ACTG 388

Configuration	**C-index**	**Change**	**IBS**	**Calibration MSE**
Baseline (without temporal modeling)	0.741	Baseline	0.134	0.029
+ Selective SSM	0.789	+0.048	0.103	0.018
+ Temporal Memory	0.768	+0.027	0.119	0.023
+ Bayesian Treatment	0.756	+0.015	0.127	0.026
SSM + Memory	0.813	+0.072	0.091	0.014
SSM + Bayesian	0.802	+0.061	0.096	0.016
Memory + Bayesian	0.784	+0.043	0.108	0.020
Proposed Model (Full)	**0.824**	**+0.083**	**0.087**	**0.012**

Predictive Performance in the Competing-Risks Setting

For the three competing events in ACTG 388, the proposed model achieved higher and more stable time-dependent AUC and C-index values than the baseline models at all follow-up time points from 3 to 24 months.

Table 3.4: Time-dependent AUC-ROC values in the competing-risks setting on ACTG 388

Time (Months)	**Cox PH**	**DeepSurv**	**DeepHit**	**SurvTRACE**	**LSTM-Surv**	**GRU-Surv**	**Proposed Model**
3	0.621	0.659	0.702	0.734	0.748	0.753	0.812
6	0.635	0.672	0.715	0.748	0.763	0.769	0.826
12	0.642	0.681	0.723	0.756	0.772	0.778	0.834
18	0.628	0.668	0.708	0.741	0.759	0.765	0.821
24	0.612	0.649	0.684	0.718	0.739	0.745	0.801

Temporal SHAP Analysis

Table 3.5 shows that CD4 count and viral load (log RNA) were the two most important features, with clear temporal variation in their contributions.

Table 3.5: Global feature importance averaged over time on ACTG 388

Feature	**Importance Score**	**Temporal Variability**	**Critical Period**	**Clinical Interpretation**
CD4 Count	0.342	High (σ=0.089)	Months 0–6	Baseline immune status determines prognosis
Log RNA	0.298	Very high (σ=0.134)	Months 0–3 and 12–18	Viral dynamics and risk of virologic rebound
Treatment Type	0.187	Moderate (σ=0.045)	Months 0–12	Effectiveness of the initial treatment response
Treatment Duration	0.156	High (σ=0.067)	Throughout follow-up	Disease progression over time
CD4%	0.143	Moderate (σ=0.052)	Months 3–9	Immune recovery
Treatment Interruption	0.089	Low (σ=0.023)	Months 6–24	Long-term treatment tolerability
Weight	0.067	Low (σ=0.018)	Months 0–6	Baseline health status
Age	0.045	Very low (σ=0.009)	Stable	Demographic factor

Table 3.6 shows that the model effectively explained the trajectory of a patient who achieved sustained viral suppression. The early increase in CD4 count and decrease in RNA level were clearly reflected in the SHAP values, indicating a protective effect over time.

Table 3.6: Example patient trajectory: Complete response (ID: 12884) on ACTG 388

Time Point	**CD4**	**CD4 SHAP**	**Log RNA**	**RNA SHAP**
Baseline (Month 0)	18 cells/μL	-0.78	5.14	+0.89
Month 3	164 cells/μL	-0.31	3.18	+0.34
Month 6	145 cells/μL	-0.18	0.95	-0.67
Month 12	193 cells/μL	+0.24	1.76	-0.23
Month 24	186 cells/μL	+0.31	2.30	+0.15

Clinical Validation

Table 3.7 shows a high level of agreement (70–77%) between the model recommendations and clinical practice. Cases in which clinical treatment was consistent with the model recommendations also achieved better clinical out-

comes, particularly during the first 12 months.

Table 3.7: Agreement with treatment recommendations over time on ACTG 388

Time Interval	Agreement Rate	95% CI	CD4 Benefit	Viral Suppression	Satisfaction Rate
0–6 months	73.2%	(68.9–77.5%)	+89 cells/μL	68.3%	82.1%
6–12 months	76.8%	(72.7–80.9%)	+127 cells/μL	71.5%	85.7%
12–24 months	71.4%	(66.8–76.0%)	+93 cells/μL	64.9%	79.3%
Overall	74.1%	(70.8–77.4%)	+103 cells/μL	68.2%	82.4%

Table 3.8 shows that the model-identified high-outcome group achieved substantially better clinical outcomes than the standard-care group, including faster viral suppression, greater CD4 recovery, and lower rates of treatment discontinuation, virologic rebound, and drug resistance.

Table 3.8: Comparison of outcomes between the model-identified favorable-trajectory group and the standard-care group on ACTG 388

Outcome	**Model-Identified Favorable Group**	**Standard Care**	**Difference**	**95% CI**	**P-value**	**Effect Size**
Time to VL < 50	16.3 weeks	22.7 weeks	-6.4 weeks	(-8.9, -3.9)	< 0.001	d = 0.72
CD4 increase after 24 months	+168 cells/μL	+124 cells/μL	+44 cells/μL	(28, 60)	< 0.001	d = 0.54
Treatment discontinuation	14.2%	22.8%	-8.6%	(-12.4, -4.8)	< 0.001	OR = 0.57
Virologic rebound	18.9%	28.4%	-9.5%	(-14.1, -4.9)	< 0.001	OR = 0.59
Drug resistance	7.3%	12.1%	-4.8%	(-8.2, -1.4)	0.006	OR = 0.57

Scalability Analysis

Table 3.9: Scalability analysis on ACTG 388

Sequence Length	Inference Time	Memory Usage	Accuracy Loss
≤ 5	18 ms	2.1 GB	None
≥ 21	145 ms	14.6 GB	$< 0.7\%$

The performance degradation of less than 0.7% when processing long sequences confirms the linear computational complexity of the Selective State

Space Model with respect to sequence length. This property ensures the model's scalability for long-term longitudinal healthcare data.

3.4. Experimental Results on the ACTG 175 Dataset

Overall Performance Comparison among Models

Table 3.10 shows that the proposed model achieved the best overall performance, with the highest C-index, the lowest IBS, and the smallest calibration error. These results demonstrate its superior ability to discriminate risk and calibrate survival probabilities compared with the competing methods.

Table 3.10: Performance comparison of survival analysis models on ACTG175.

Model	**C-index**	**95% CI**	**IBS**	**95% CI**	**Calib. MSE**	**95% CI**
Cox PH	0.61	(0.587–0.633)	0.187	(0.178–0.196)	0.048	(0.044–0.052)
DeepSurv	0.65	(0.628–0.672)	0.165	(0.157–0.173)	0.037	(0.034–0.040)
DeepHit	0.69	(0.669–0.711)	0.142	(0.135–0.149)	0.034	(0.031–0.037)
SurvTRACE	0.72	(0.700–0.740)	0.128	(0.122–0.134)	0.029	(0.027–0.031)
Proposed Model	**0.77**	**(0.752–0.788)**	**0.106**	**(0.101–0.111)**	**0.016**	**(0.014–0.018)**

Time-Dependent AUC-ROC Analysis

Table 3.11 shows that the proposed model achieved the highest time-dependent AUC-ROC at every follow-up time point, demonstrating stable and superior risk discrimination compared with the competing methods in both short-term and long-term prediction.

Table 3.11: Time-dependent AUC-ROC values on ACTG175.

Time (Days)	**Cox PH**	**DeepSurv**	**DeepHit**	**SurvTRACE**	**Proposed Model**
180	0.578	0.620	0.672	0.701	0.758
360	0.603	0.633	0.685	0.715	0.769
540	0.614	0.652	0.691	0.723	0.776
720	0.621	0.658	0.705	0.731	0.783
900	0.608	0.640	0.687	0.718	0.774

Kaplan–Meier Survival Analysis by Baseline CD4 Group

The Kaplan–Meier analysis (Table 3.12) shows clear risk stratification according to baseline CD4 count, with the high-CD4 group exhibiting substantially higher survival probabilities. The log-rank test ($p < 0.001$) confirms that the difference is statistically significant, indicating that CD4 count is an important prognostic factor.

Table 3.12: Kaplan–Meier results according to baseline CD4 stratification (ACTG175).

Group	**N**	**1-Year Survival (95% CI)**	**2-Year Survival (95% CI)**	**Median Survival Time (Days)**
CD4 $>$ 350	927	0.924 (0.906–0.942)	0.864 (0.841–0.887)	Not reached
CD4 $\leq$ 350	1212	0.847 (0.826–0.868)	0.684 (0.658–0.710)	798

Log-rank test: $\chi^2 = 23.8, p = 0.00000791$

Treatment-Specific Prediction Analysis across Patient Subgroups.

The study also analyzed treatment-specific outcomes, showing that the model predicted distinct and clinically consistent outcomes across treatment regimens, with ZDV+ddI yielding the best prognosis and ZDV monotherapy the worst, particularly in patients with low CD4 counts; this pattern was consistent across patient groups, demonstrating the model's ability to personalize predictions based on immune status and treatment history.

Calibration Analysis and Practical Implications.

Calibration analysis showed that the proposed model achieved the lowest calibration error (0.016), outperforming the comparison methods and demonstrating accurate and reliable survival probability estimation.

Treatment Recommendation Analysis.

The results demonstrate the model's potential to support personalized treatment recommendations, with moderate agreement with clinical practice (68.0%), particularly among high-risk groups. Patients whose treatment choices aligned with the model's recommendations achieved better treatment outcomes.

Feature Importance Analysis.

Feature importance analysis showed that CD4 Count was the most important feature across all treatment groups (0.22–0.28), followed by Prior ZDV and the CD4/CD8 ratio. Other demographic and clinical factors had lower importance. These results indicate that the model focuses on key determinants of survival prediction while adapting to differences across treatment regimens.

Statistical Significance Testing between Models.

Statistical significance testing showed that the proposed model achieved a significantly higher C-index than Cox PH, DeepSurv, DeepHit, and SurvTRACE, with improvements ranging from 0.05 to 0.16. All pairwise comparisons yielded p-values < 0.05, confirming that the observed performance gains are statistically significant rather than attributable to random variation.

Chương 4. CAUSAL EFFECT ESTIMATION AND OUTCOME PREDICTION IN HIV SURVEILLANCE AND TREATMENT

1. Introduction

This chapter analyzes the challenges of causal inference using observational HIV data and proposes a modular causal inference framework that decouples representation learning from treatment effect estimation. This approach improves the reliability of individualized treatment effect estimation and supports clinical decision-making in HIV care.

2. Proposed Method

2.1. Definition of the Individual Causal Inference Problem

The problem is formulated within the potential outcomes framework for observational HIV data, in which each individual is characterized by a set of baseline covariates X, a discrete intervention variable T, and a binary outcome Y. The model aims to learn the conditional expectation function of the potential outcome, $\mu(x,t) = \mathbb{E}[Y(t) \mid X = x]$, from which causal effects can be derived at both the individual level and the conditional level according to baseline characteristics. Identification of these quantities relies on fundamental assumptions, including consistency, ignorability, positivity, and SUTVA, while also requiring consideration of the challenges associated with their application to real-world observational data.

2.2. Construction and Evaluation of Causal Graphs for Intervention Inference

Two candidate causal graphs are considered: an expert knowledge graph G_{EX} and a data-driven graph G_{REX} constructed using SHAP-bootstrap, ANM–HSIC, and cycle removal. Sensitivity analysis is used to select the final graph G_{final}, which guides intervention effect estimation based on causal rather than purely statistical associations.

2.3. Overall System Architecture

Two causal graphs are constructed from expert knowledge and data-driven inference using the REX model, after which the optimal structure is selected through sensitivity analysis. The causal variable set is represented as a feature matrix ($X \in \mathbb{R}^{n \times p}$), which helps reduce bias caused by irrelevant or post-intervention variables.

The matrix (X) is provided to two components: (i) TITAN for nonlinear representation learning and (ii) DRLearner for causal effect estimation according to the doubly robust principle, enabling the calculation of (ITE) and (ATE).

A reinforcement learning agent learns a policy ($\pi(\cdot)$) to dynamically com-

bine the outputs of the two components for each sample, thereby balancing representation capability and estimation reliability. The CAUSALRLSTACK framework therefore provides a flexible and accurate approach to individual causal effect estimation.

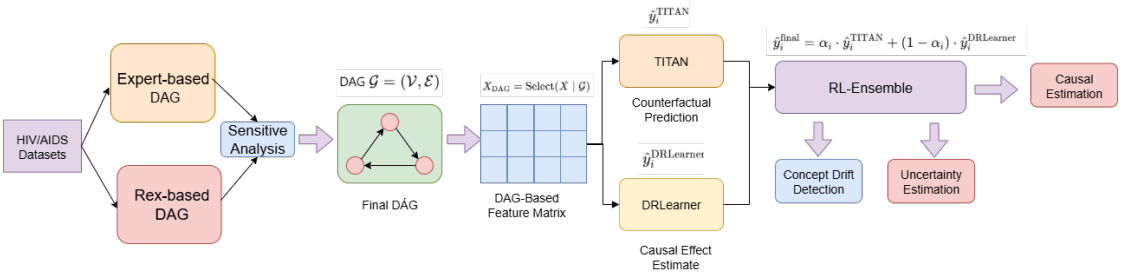

Figure 4.1: Overall architecture of the hybrid causal prediction model.

Figure 4.1 illustrates the overall architecture of the CAUSALRLSTACK framework. The process begins with HIV/AIDS data, followed by causal graph construction to identify features that are meaningful for inference. TITAN and DRLearner operate in parallel to exploit two different aspects of the problem: representation learning and causal effect estimation. Finally, the reinforcement learning module acts as a coordinator, flexibly combining the two sources of information and generating the final causal prediction together with its associated uncertainty.

2.4. Deep Causal Representation Learning with the TITAN Architecture

The TITAN module is designed to learn a latent representation $z_i \in \mathbb{R}^d$ from the input $x_i \in \mathbb{R}^p$, supporting both prediction and counterfactual inference: $z_i = \Phi_{\text{TITAN}}(x_i; \theta_\Phi)$. First, the features are embedded through one-dimensional convolution and Transformer layers as follow.

$$h_i^{(0)} = \text{Conv1D}(x_i), \quad \tilde{h}_i^{(\ell)} = \text{MHA}(h_i^{(\ell-1)}), \quad h_i^{(\ell)} = \text{LN}(h_i^{(\ell-1)} + \tilde{h}_i^{(\ell)}),$$

$$h_i^{(\ell)} = \text{LN}(h_i^{(\ell)} + \text{FFN}(h_i^{(\ell)})).$$

To improve stability under distributional changes, an external memory is integrated:

$$g_i = \sigma(W_g h_i^{(L)} + b_g), \quad M_i = g_i \odot M_{i-1} + (1 - g_i) \odot h_i^{(L)}, \quad z_i = [h_i^{(L)}, |, M_i].$$

The representation z_i is used to estimate the potential outcome under each intervention level t as follow.

$$\hat{\mu}_i(t) = \sigma(W_o z_i + b_o), \quad \mu(x_i, t) = \mathbb{E}[Y_i(t) \mid X_i]. \tag{4.1}$$

The model is trained using the binary cross-entropy loss as follow.

$$\mathcal{L}_{\text{TITAN}} = -\frac{1}{n}\sum_{i=1}^{n}\left[y_i \log \hat{\mu}_i(T_i) + (1-y_i)\log\left(1-\hat{\mu}_i(T_i)\right)\right]. \tag{4.2}$$

After training, the model generates potential outcomes for every $t \in \mathcal{T}$ and estimates the individual treatment effect: $\widehat{\text{ITE}}_i = \hat{\mu}_i(1) - \hat{\mu}_i(0)$.

2.5. Personalized Intervention Effect Estimation with DRLearner

DRLearner is integrated as a parallel branch alongside TITAN to estimate causal effects according to the *doubly robust* principle. This method combines two components: an outcome regression model and a propensity score model, ensuring consistency if at least one of the two models is correctly specified.

For each individual i, two outcome regression functions are estimated:

$$\mu_1(x_i) = \mathbb{E}[Y_i \mid X_i = x_i, T_i = 1], \quad \mu_0(x_i) = \mathbb{E}[Y_i \mid X_i = x_i, T_i = 0], \tag{4.3}$$

together with the propensity score:

$$\hat{e}(x_i) = \mathbb{P}(T_i = 1 \mid X_i = x_i). \tag{4.4}$$

The individual treatment effect (ITE) is calculated using the doubly robust formula:

$$\hat{\tau}_i = \left(\frac{t_i - \hat{e}(x_i)}{\hat{e}(x_i)(1-\hat{e}(x_i))}\right)(y_i - \hat{\mu}_{t_i}(x_i)) + \hat{\mu}_1(x_i) - \hat{\mu}_0(x_i), \tag{4.5}$$

and the average treatment effect is given by:

$$\text{ATE} = \frac{1}{n}\sum_{i=1}^{n}\hat{\tau}_i. \tag{4.6}$$

For multi-arm interventions $T \in \{0, 1, \ldots, K\}$, the effect between two treatment arms a, b is extended as:

$$\hat{\tau}_i(a,b) = \frac{\mathbf{1}[T_i = a]}{\hat{e}_a(X_i)}(Y_i - \hat{\mu}_a(X_i)) - \frac{\mathbf{1}[T_i = b]}{\hat{e}_b(X_i)}(Y_i - \hat{\mu}_b(X_i)) + \hat{\mu}_a(X_i) - \hat{\mu}_b(X_i). \tag{4.7}$$

The model is implemented using Random Forest for outcome regression and Logistic Regression for propensity score estimation, combined with 5-fold cross-validation to ensure out-of-sample generalizability.

2.6. Adaptive Coordination Using Reinforcement Learning for Counterfactual Estimation

The model uses reinforcement learning (RL) to dynamically coordinate two parallel inference branches: TITAN for deep representation learning and DR-Learner for doubly robust estimation. Instead of using a fixed mixing weight, the model learns an optimal individual-level combination strategy through a Markov decision process (MDP).

State and action modeling. For each individual i, the agent state is constructed from the two counterfactual predictions and the uncertainty signal:

$$s_i = \left[\hat{y}_i^{\text{TITAN}}, \hat{y}_i^{\text{DRLearner}}, u_i\right]. \tag{4.8}$$

The action corresponds to a continuous mixing weight:

$$a_i \equiv \alpha_i \in [0, 1], \tag{4.9}$$

generated by the policy:

$$\alpha_i = \pi_\theta(a_i \mid s_i). \tag{4.10}$$

Prediction mixing mechanism. The final prediction is a convex combination of the outputs from the two branches:

$$\hat{y}_i = \alpha_i \hat{y}_i^{\text{TITAN}} + (1 - \alpha_i)\hat{y}_i^{\text{DRLearner}}. \tag{4.11}$$

Reward design. The reward is defined as the negative cross-entropy loss:

$$r_i = -\Big[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)\Big]. \tag{4.12}$$

Policy update. The policy is optimized using an Actor–Critic mechanism with policy gradient:

$$\theta \leftarrow \theta + \eta \nabla_\theta \log \pi_\theta(a_i \mid s_i) Q(s_i, a_i). \tag{4.13}$$

Through this mechanism, the mixing weight is learned adaptively for each individual, enabling the system to optimally exploit the relative reliability of the two inference branches and improve the quality of counterfactual predictions in heterogeneous data settings.

2.7. Uncertainty Estimation and Adaptation to Distributional Changes

The study additionally incorporates an uncertainty estimation mechanism based on Monte Carlo Dropout to assess the reliability of individual-level causal effect estimates through the decomposition of aleatoric and epistemic uncertainty. At the same time, the system integrates concept drift detection to trigger model updates when the data distribution changes, thereby maintaining stability and robustness in real-world environments.

3. Experiments and Results

3.1. Dataset Description

The experiments and evaluation used two publicly available HIV datasets from the Kaggle platform to assess the causal inference framework in two different settings.

The first dataset, EDHS–HIV/AIDS, originates from the Ethiopia Demographic and Health Survey and reflects population-level demographic characteristics, behaviors, and HIV-related awareness.

The second dataset, AIDS Virus Infection Prediction, is based on the ACTG 175 clinical trial dataset and represents the HIV treatment trial setting, including demographic, clinical, and immunological characteristics.

3.2. Simulation Study

1. Data Generating Process. The study constructed a semi-realistic causal system simulating an HIV setting with $n = 5000$. The system included observed confounders, an unobserved confounder $X_7 \sim \mathcal{N}(0, 1)$, an instrumental variable Z, a mediator M, and age-dependent heterogeneous treatment effects as follows.

$$\tau_i = -1.5 + 0.08(X_{0i} - 50). \tag{4.14}$$

Treatment assignment and outcomes were modeled as follows:

$$\text{logit } \Pr(T = 1 \mid X, Z) = \ldots, \tag{4.15}$$

$$M = 2 - 1.5T + 0.5X_3 + \varepsilon_M, \tag{4.16}$$

$$\text{logit } \Pr(Y = 1) = \cdots + T \cdot \tau + \varepsilon_Y. \tag{4.17}$$

The observed dataset excluded the variable X_7 to simulate unmeasured confounding. In addition, a reference dataset containing the true individual treatment effects (ITE) was constructed for evaluation.

2. Causal Graphs: Ground-Truth and Estimated Structures

Figure 4.2 presents the ground-truth causal graph, in which the confounders, instrumental variable, and mediator are explicitly defined according to the data-generating mechanism. In contrast, Figure 4.3 presents the graph estimated from observational data, which has a denser structure and contains deviations caused by correlated relationships. This difference highlights the challenges of observational data and the importance of robust causal effect estimation methods.

3. Simulation Study Results. Table 4.1 shows that the proposed model achieved the best performance across the causal evaluation metrics, with the lowest PEHE, ATE error, and bias, together with the highest R^2or ITE estimation. Table 4.2 shows that CAUSALRLSTACK not only outperforms competing methods in causal inference but also achieves superior predictive performance.

4.Robustness Testing. The robustness experiments showed that the model

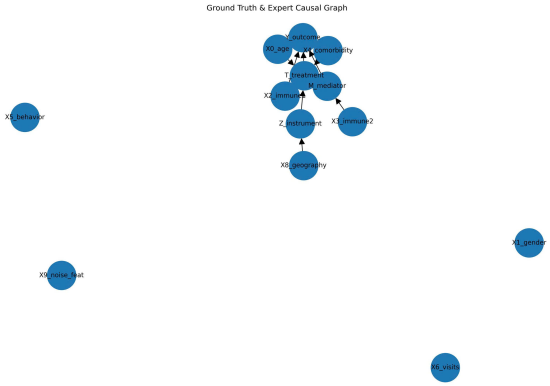

Figure 4.2: Ground-truth causal graph

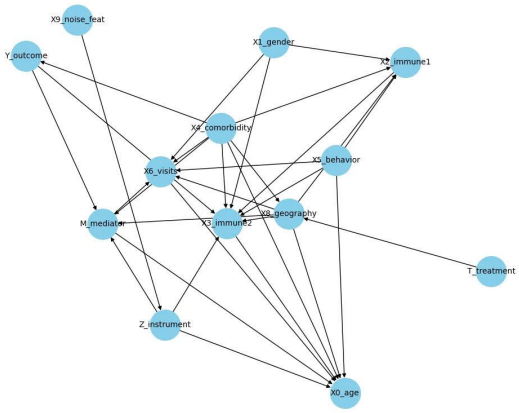

Figure 4.3: Estimated causal graph

Table 4.1: Performance comparison of causal inference models on simulated data

Method	**PEHE** ↓	**ATE Error** ↓	**ATE Bias**	**ITE** R^2 ↑	**Estimated ATE**
CAUSALRLSTACK	**0.0980**	**0.0198**	**-0.0089**	**0.5834**	**0.2261**
CAT Network	0.1089	0.0234	0.0187	0.5445	0.2537
Orthogonal Random Forests	0.1156	0.0267	-0.0145	0.5198	0.2205
Double/Debiased ML	0.1203	0.0289	0.0223	0.4987	0.2573
Causal Forest	0.1267	0.0321	-0.0198	0.4743	0.2152
X-Learner	0.1334	0.0367	0.0298	0.5456	0.2648
CEVAE	0.1423	0.0412	-0.0267	0.5123	0.2083

maintained stable performance when the sample size was reduced, the confounding strength increased, and the outcome model was misspecified. In particular, owing to its doubly robust property, the ATE bias remained low even under adverse scenarios, whereas methods relying on a single robustness mechanism exhibited substantial performance degradation.

3.3. Experimental Results on the HIV Datasets

Construction and Selection of Causal Graphs for the Two Datasets. For both datasets, the causal structure was constructed using a two-step procedure: (i) establishing a DAG based on epidemiological and clinical expert knowledge and (ii) inferring a DAG from the data using the REX discovery framework. The two structures were compared in terms of biological plausibility, stability of ATE/ITE estimation, and sensitivity analysis.

The expert-defined DAG served as the reference structure, whereas the data-driven DAG was refined to remove implausible edges and retain empirically supported relationships. The final graph for each dataset, $G_{\text{final}}^{\text{dataset1}}$ and $G_{\text{final}}^{\text{dataset2}}$,

Table 4.2: Comparison of predictive performance on simulated data

Method	Accuracy	F1-score	AUC-ROC
Proposed Model	**0.8734**	**0.862**	**0.924**
CAT Network	0.8267	0.814	0.889
Orthogonal Random Forests	0.8189	0.807	0.883
Double/Debiased ML	0.8123	0.798	0.876
Causal Forest	0.8045	0.789	0.867
X-Learner	0.7967	0.776	0.854
CEVAE	0.7834	0.761	0.841

was used to identify the adjustment set, including confounders, mediators, and instrumental variables, thereby providing the foundation for the entire causal effect estimation procedure in the thesis.

Comparison with Existing Causal Inference Methods. The comparison methods included *Double/Debiased Machine Learning (DML)*, *Orthogonal Random Forests (ORF)*, *Causal Forests*, *X-Learner*, *CEVAE (Causal Effect Variational Autoencoder)*, and *Causality-Aware Transformer (CAT)*.

Experimental Results.

1. Main Results. Table 4.3 shows that the proposed model outperformed all baseline methods on both datasets.

Table 4.3: Comparison of the proposed model with existing causal inference methods

	Dataset 1			Dataset 2		
Method	**Accuracy**	**F1-Score**	**AUC-ROC**	**Accuracy**	**F1-Score**	**AUC-ROC**
Proposed Model	**0.861**	**0.845**	**0.897**	**0.855**	**0.839**	**0.892**
Double/Debiased Machine Learning	0.774	0.749	0.830	0.768	0.744	0.827
Orthogonal Random Forests	0.781	0.758	0.834	0.774	0.753	0.831
CEVAE	0.743	0.720	0.798	0.740	0.753	0.795
Causality-Aware Transformer (CAT) Networks	0.784	0.763	0.838	0.778	0.758	0.835
Causal Forest	0.768	0.747	0.824	0.763	0.742	0.821
X-Learner	0.756	0.738	0.812	0.751	0.733	0.808

Experimental results show that the full model integrating TITAN, DRLearner, and reinforcement learning (RL) achieves the best performance across both datasets and all evaluation metrics compared with ablated variants in which one or more components are removed.

2. Average Treatment Effect Estimation (ATE). Beyond predictive performance, the study further assessed causal inference capability through Average Treatment Effect (ATE) estimation. The proposed model achieved ATE estimates of 0.247 and 0.243 on the two datasets, respectively, demonstrating stable treatment effect estimation across datasets while maintaining consistency with the trends observed in competing methods.

3. Ablation Study Table 4.4 shows that the full model achieved the highest performance on both datasets. The results indicate that removing key components such as RL-ensemble, TITAN, the causal graph, uncertainty estimation, or concept drift detection resulted in a clear reduction in performance.

Table 4.4: Ablation study: Performance effects of removing individual system components

	Dataset 1			**Dataset 2**		
Technique	**Accuracy**	**F1**	**ATE**	**Accuracy**	**F1**	**ATE**
Proposed Model	**0.861**	**0.845**	**0.247**	**0.855**	**0.839**	**0.243**
Without RL Ensemble	0.789	0.767	0.235	0.788	0.765	0.231
Without TITAN	0.762	0.729	0.230	0.755	0.722	0.227
Without DRLearner	0.779	0.753	0.232	0.773	0.748	0.229
Without Causal Graph	0.815	0.792	0.239	0.807	0.787	0.235
Without Uncertainty Estimation	0.850	0.831	0.244	0.843	0.825	0.241
Without Concept Drift Detection	0.857	0.838	0.246	0.849	0.831	0.242

4.Uncertainty Analysis. Monte Carlo Dropout evaluation showed that the proposed model achieved the lowest predictive uncertainty on both datasets (0.093 and 0.092, respectively), indicating higher confidence and reliability than the alternative model configurations.

GENERAL CONCLUSION

This thesis develops deep learning methods for HIV prediction and decision support through three main contributions: (i) an attention-based graph neural network with graph pooling for epidemiological forecasting, (ii) a Selective State Space Model for longitudinal survival analysis, and (iii) a hybrid causal inference framework integrating deep representation learning, DR-Learner, and reinforcement learning. The proposed methods can be extended to incorporate additional data sources and adapted to other infectious diseases.

LIST OF PUBLICATIONS BY THE AUTHOR

CT1. **Phạm Thành Đạt**, Khai Quang Tran, Viet Anh Nguyen, "CAUSAL-RLSTACK: Adaptive balancing of deep representation and causal effect estimation with application to HIV-related health data," BioData Mining, vol. 18, p. 77, 2025, ISSN 1756-0381, https://doi.org/10.1186/s13040-025-00492-3 (SCIE Q1)

CT2. **Phạm Thành Đạt**, Duong Van Nguyen, Thanh Tan Tran, Viet Anh Nguyen, "Casenet: A novel HIV forecasting model using attention-based spatio-temporal graph network leveraging city-wide infection correlations," Neural Computing and Applications, vol. 38, p. 119, 2026, ISSN 0941-0643, https://doi.org/10.1007/s00521-025-11804-3 (SCOPUS Q1)

CT3. **Phạm Thành Đạt**, Duong Nguyen Van, Thanh Tran Tan, Viet Anh Nguyen, "Spatio-temporal graph learning with epidemiological factors for HIV epidemic short-term prediction," Journal of Computer Science and Cybernetics, 2024, ISSN 1813-9663, https://doi.org/10.15625/1813-9663/21185

CT4. **Phạm Thành Đạt**, Duong Van Nguyen, Thanh Tan Tran, Thang Truong Nguyen, "ViT-Survive: Temporal patch-based survival analysis of HIV treatment dropout using longitudinal data," in *Advances in Information and Communication Technology*, Proc. ICTA 2025, LNNS 1833, Springer, 2026, pp. 1–10, https://doi.org/10.1007/978-3-032-18162-6_-14

CT5. **Phạm Thành Đạt**, Nguyen Van Duong, Tran Tan Thanh, Vu Tuan Anh, Nguyen Truong Thang, "Improving survival prediction of HIV treatment dropout with group-wise self-attention network," Proceedings of the 18th National Conference on Fundamental and Applied Information Technology Research (FAIR 2025), Dec. 2025, https://doi.org/10.15625/vap.2025.0322

CT6. **Phạm Thành Đạt**, Nguyen Van Duong, and Nguyen Viet Anh, "Harnessing spatio-temporal learning with 2-layer attention-driven graph neural networks for HIV epidemic forecasting in Ho Chi Minh City," Proceedings of the 27th National Conference on Information and Communication Technology (ICTA 2024), Nha Trang, Vietnam, Oct. 2024

CT7. **Phạm Thành Đạt**, Nguyen Van Duong, Nguyen Tuan Cuong, and Ha Thi Hong Van, "Apply probabilistic approach with unsupervised learning for patient matching in health information exchange," Proceedings of the 16th National Conference on Fundamental and Applied Information Technology Research (FAIR 2023), Da Nang, Vietnam, Sep. 2023, https://doi.org/10.15625/vap.2023.0031

CT8. **Phạm Thành Đạt**, Tran Quang Khai, and Nguyen Viet Anh, "A Unified Memory-Guided Framework for Disease State Learning in Personalized HIV Survival Analysis", submitted to Discover Artificial Intelligent, 2026.