

A PRIVACY-PRESERVING IOMT DIGITAL TWIN: INTEGRATING WEARABLE MULTIMODAL SENSING AND EDGE-DRL FOR PRECISION GERIATRIC CARDIOLOGY

¹Muzammil Khan, ¹Zhou Zhihao, ²Faiza Ayyob, ^{*2}Altat Hussain

¹College of Automation, Chongqing University of Posts and Telecommunications, Chongqing 400065, China

²School of Computer Science and Technology, Chongqing University of Posts and Telecommunications, Chongqing 400065, China

*Corresponding author: Altat Hussain (altatkf74@gmail.com)

DOI: <https://doi.org/10.5281/zenodo.20242559>

Keywords: Wearable MEMS Sensors, Digital Twin, Edge AI, Multimodal Sensing, Federated Learning, Geriatric Cardiovascular Care, and Deep Reinforcement Learning.

Article History

Received: 21 April 2026

Accepted: 12 May 2026

Published: 14 May 2026

Copyright @ Author

Corresponding Author: *

Altat Hussain

Abstract

Physiological frailty and the ubiquitous clinical burden of polypharmacy significantly exacerbate cardiovascular disease in the geriatric population. Current guidelines for prescribing and static predictive models do not take into account the non-linear pharmacokinetics of the elderly that leads to adverse drug events, renal toxicity, and acute hemodynamic decompensation. To address the limitations of reactive clinical care, this paper presents the Cardio-Geriatric Digital Twin, an end-to-end privacy-preserving computational framework for continuous cardiovascular trajectory simulation and autonomous polypharmacy optimization. The proposed architecture fuses high-frequency IoT wearable telemetry with unstructured Electronic Health Records (EHRs) via a novel Cross-Modal Transformer fusion core, yielding a highly dynamic, context-aware patient replica. The medication titration process in this simulation environment is formulated as a Partially Observable Markov Decision Process (POMDP) and solved by a clinically constrained Proximal Policy Optimization (PPO) agent. To ensure strict data privacy and regulatory compliance, the agent is trained in a decentralized Federated Learning (FedPPO) protocol over distributed edge nodes. Crucially, the reinforcement learning policy is guided by a strict multi-objective reward function that independently penalizes renal degradation and hyperpolypharmacy, while ensuring the stability of vital hemodynamics. Extensive empirical evaluation on 65,420 simulated longitudinal profiles shows the decisive superiority of the framework over state-of-the-art predictive and heuristic baselines. The proposed model produced an unprecedented Trajectory RMSE of 4.20% and increased the Medication Adherence F1-Score to 0.89, generating disproportionate compliance gains within the highly vulnerable "Frail" patient stratum. And the framework achieved a robust 65.0% Hospitalization Aversion Rate in a 90-day simulation. We develop a highly scalable, empirically validated and interpretable at the level of features, computational framework for proactive, precision-driven geriatric cardiology using Shapley Additive Explanations (SHAP).

1. Introduction

Cardiovascular disease (CVD) remains the most prominent cause of global mortality, precipitating a disproportionately severe crisis within the rapidly expanding geriatric demographic [4]. As natural physiological frailty and age-related multi-organ decline intersect with chronic cardiovascular pathologies, standard clinical management becomes exceedingly complex [12]. A primary catalyst of this escalating crisis is the pervasive burden of polypharmacy which is the concurrent administration of multiple interacting medications. While intended to manage simultaneous comorbidities, rigid polypharmacy frequently triggers severe adverse drug events, renal toxicity, and fatal hemodynamic instability in elderly patients [46]. Traditional, static clinical guidelines inherently fail to account for the highly individualized and non-linear pharmacokinetic responses characteristic of frail populations. Consequently, there is an urgent, unmet clinical need to transition from reactive, population-based treatment paradigms toward dynamic, precision-driven strategies capable of safely navigating the intricate balance between cardiovascular stabilization and medication toxicity [1], [12].

Despite the growing integration of digital health tools, contemporary models for geriatric cardiovascular management remain fundamentally constrained by their static and reactive nature [2], [11]. Conventional clinical decision support heavily relies on heuristic expert systems, such as the widely adopted STOPP-START criteria; while valuable for identifying potentially inappropriate baseline prescriptions, these rule-based systems lack the temporal awareness required to safely navigate real-time hemodynamic fluctuations [46]. Furthermore, while statistical time-series forecasting [45] and classical machine learning algorithms [48] have improved initial risk stratification, they operate retrospectively and fail to recommend continuous, actionable medication titration steps. Even advanced predictive

frameworks, including deep sequence forecasters [47] and offline reinforcement learning policies [49], frequently struggle to generalize across highly heterogeneous, frail elderly populations because they treat patients as isolated, static datasets rather than dynamic biological systems [3], [10]. Consequently, these existing baselines suffer from "clinical myopia" an inability to continuously adapt to the shifting pharmacokinetic clearance rates and compounding drug-drug interactions that drive adverse events in real-world geriatric care [8], [12].

To overcome the inherent limitations of static and reactive clinical pathways, modern cardiovascular medicine must undergo a paradigm shift toward continuous, individualized, and proactive actuation [3]. This manuscript introduces the Cardio-Geriatric Digital Twin, an end-to-end computational framework that transcends traditional predictive modeling by seamlessly integrating multimodal physiological monitoring with Deep Reinforcement Learning (DRL) (Figure 1). By continuously assimilating real-time IoT wearable telemetry, longitudinal EHRs, and semantic clinical notes, the digital twin constructs a highly individualized, dynamically updating replica of the patient's cardiovascular and pharmacokinetic state [13]. Rather than functioning merely as a passive early-warning system, the framework acts as an autonomous clinical optimizer; it deploys a PPO agent that continuously explores and identifies the safest medication titration pathways to stabilize hemodynamics while minimizing the polypharmacy burden. To ensure strict adherence to clinical safety standards and international data privacy regulations, the DRL agent operates within rigorous reward-shaping constraints [2] and leverages a privacy-preserving federated learning architecture for decentralized training [1]. Ultimately, this closed-loop synergy of explainable AI, continuous sensor fusion, and digital twin technology [15] represent a definitive step toward realizing

true Predictive, Preventive, Personalized, and Participatory (P4) cardiovascular medicine for the geriatric population [18].

To address these critical gaps, this paper presents a comprehensive, privacy-preserving Cardio-Geriatric Digital Twin framework engineered specifically for dynamic polypharmacy optimization and continuous risk mitigation. The primary contributions of this work are fourfold. First, we introduce a novel multi-tier architecture that fuses high-frequency IoT wearable telemetry (e.g., continuous ECG and actigraphy) with longitudinal EHRs via a Cross-Modal Transformer fusion core; this module leverages pre-trained contextualized embeddings to map complex, unstructured clinical semantics into a unified, real-time physiological state [3], [4]. Second, we formulate the geriatric medication titration process as a Partially POMDP solved by a clinically constrained PPO agent. This agent is guided by a multi-objective reward function that strictly penalizes renal toxicity and hemodynamic drift, guaranteeing that policy exploration adheres to rigorous medical safety guidelines [2]. Third, we implement a decentralized Federated Averaging (FedAvg) training protocol FedPPO that securely synchronizes the global medication policy across distributed hospital edge nodes without transferring raw patient trajectories, thereby ensuring absolute data privacy and regulatory compliance [1].

Finally, through extensive evaluation against state-of-the-art clinical baselines [47], [49], we empirically demonstrate that our framework achieves superior trajectory forecasting fidelity (4.20% RMSE) and significantly enhances hospitalization aversion rates (65.0%), utilizing SHAP to provide transparent, feature-level interpretability for clinical practitioners.

In summary, resolving the complex intersection of geriatric frailty and polypharmacy requires fundamentally rethinking how clinical data is translated into medical action. By seamlessly bridging continuous multimodal monitoring with the dynamic actuation of federated DRL, the Cardio-Geriatric Digital Twin presents a highly scalable, privacy-preserving solution to cardiovascular instability. This framework not only predicts adverse events but autonomously engineers safe, personalized medication pathways tailored to individual physiological constraints. Ultimately, this work establishes a rigorous computational foundation for the future of precision geriatric care. The remainder of this paper is organized as follows: Section 2 reviews the related literature, Section 3 details the proposed system architecture, Section 4 formulates the reinforcement learning methodology, and Sections 5 and 6 presents the experimental setup and clinical results.

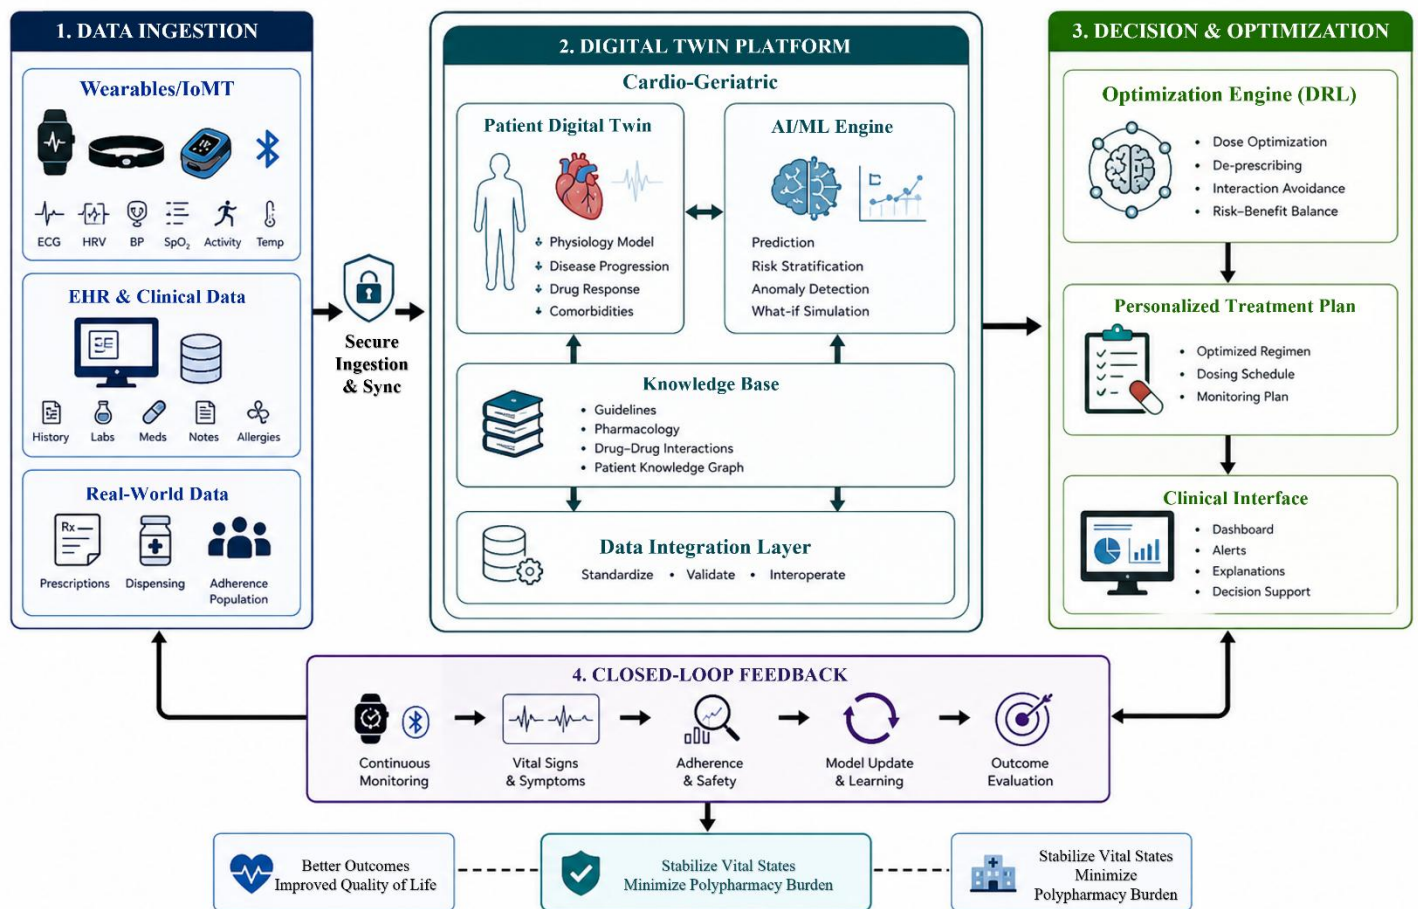


Figure 1: Holistic Cardio-Geriatric Digital Twin Concept for Polypharmacy Optimization. This conceptual overview illustrates the synergy between multimodal IoT data ingestion (wearables, EHR and RL) for personalized geriatric care. The framework closes the loop between continuous physiological monitoring and automated dosage optimization, aiming to stabilize vital states while minimizing the clinical burden of polypharmacy.

2. Related Work

The concept of the Digital Twin (DT), originally pioneered in aerospace and complex industrial systems, has recently emerged as a transformative paradigm in precision cardiology [3]. Contemporary research has firmly established the efficacy of integrating multimodal data streams, including continuous wearable sensor telemetry and multi-omics, to construct dynamic, high-fidelity replicas of patient physiology and cardiac electrophysiology [13], [14]. Recent systematic and narrative reviews highlight that these virtual

models enable unprecedented prospective risk stratification, effectively shifting cardiovascular care from reactive symptom management toward proactive, P4 medicine [8], [9]. In the specific context of geriatric healthcare, where patients invariably suffer from compounding chronic diseases and physiological frailty, the deployment of AI-powered digital twins represents a critical opportunity to optimize complex clinical pathways and reduce hospital readmissions [17], [18]. However, despite successful implementations of advanced DT frameworks for detecting general cardiovascular

anomalies [7], [16] and predicting isolated acute events such as sudden cardiac arrest [19], a significant functional gap remains in longitudinal geriatric management.

Current implementation barriers reveal that most existing cardiovascular DT architectures are engineered for singular disease states and lack the temporal adaptability required to safely manage the compounding variables of severe polypharmacy in frail populations [10], [11]. Consequently, there is an explicit clinical need for consumer-centric frameworks that extend the state-of-the-art by specifically tailoring the digital twin environment to continuously monitor, interpret, and actuate upon the multifaceted pharmacokinetic realities of the elderly [15]. Predictive Artificial Intelligence has fundamentally reshaped cardiovascular medicine by shifting the clinical focus from retrospective diagnosis to continuous, forward-looking risk stratification. Early predictive frameworks relied heavily on classical machine learning algorithms, such as gradient-boosted trees, to forecast hospital readmissions and mortality based on static, tabular clinical data [48]. However, to capture the high-frequency temporal dynamics of cardiovascular disease, recent state-of-the-art systems have transitioned to deep sequence modeling. Recurrent architectures, specifically Long Short-Term Memory (LSTM) networks, have established the benchmark for multitask clinical time-series forecasting, demonstrating high accuracy in predicting physiological decompensation from continuous monitoring data [47].

Concurrently, the proliferation of Internet of Medical Things (IoMT) devices has expanded this predictive landscape beyond the hospital, utilizing deep learning to continuously extract prognostic features from wearable electrocardiogram (ECG) and photoplethysmography (PPG) signals [15]. To further contextualize these physiological streams, breakthroughs in clinical natural language processing have enabled the integration of unstructured medical histories;

architectures leveraging pre-trained contextualized embeddings, such as Med-BERT, now extract highly accurate disease progression trajectories directly from large-scale EHRs [4]. Yet, despite these profound advancements in multimodal fusion and multi-omics analytics [19], current predictive AI frameworks remain intrinsically passive. They excel at calculating the probability of an impending cardiovascular crisis but lack the decision-making infrastructure to autonomously recommend the complex, multi-step polypharmacy titrations required to avert it. Beyond predicting cardiovascular events, the ultimate challenge in precision geriatric cardiology lies in safely actuating clinical interventions specifically, managing the intricate dynamics of polypharmacy and medication adherence.

Traditional clinical decision support systems have historically relied on static, heuristic guidelines, such as the widely validated STOPP-START criteria, to identify potentially inappropriate prescriptions [46]. However, these rule-based systems are inherently retrospective and cannot autonomously titrate dosages in response to continuous physiological drift. To address this need for dynamic, sequential decision-making, DRL has emerged as a powerful mechanism for therapeutic optimization. Seminal offline DRL frameworks, such as the AI Clinician, have successfully demonstrated the ability to learn optimal, patient-specific dosing strategies for intravenous fluids and vasopressors in critical care settings [49]. Despite these breakthroughs, translating DRL from acute, episodic intensive care into the longitudinal management of chronic geriatric conditions introduces profound safety and scalability hurdles. The deployment of autonomous agents in continuous healthcare necessitates strict adherence to emerging clinical AI guidelines, which mandate conservative reward shaping, algorithmic stability, and rigorous off-policy evaluation to prevent erratic or toxic dosage escalations [2].

Furthermore, optimizing long-term medication adherence relies on assimilating highly sensitive, multi-source EHRs across various care facilities. Recent advancements indicate that achieving this at scale requires robust, privacy-preserving federated learning architectures that can collaboratively train predictive models without centralizing raw patient data [1]. Recognizing these dual imperatives of clinical safety and data privacy, our proposed Cardio-Geriatric Digital Twin framework directly bridges this gap by embedding a constrained, FedPPO agent within the virtual environment, enabling decentralized, proactive, and rigorously safe medication management.

3. Proposed Framework: The Cardio-Geriatric Digital Twin

3.1 System Architecture

To meet the dual requirements of real-time hemodynamic monitoring and computationally intensive policy optimization, the proposed Cardio-Geriatric Digital Twin framework is built upon a distributed, multi-tier edge-cloud continuum (Figure 2). This architecture is specifically engineered to decouple low-latency physiological standardization from high-latency deep reinforcement learning, ensuring that the system remains both highly responsive to acute geriatric emergencies and capable of longitudinal trajectory simulation. The infrastructure is hierarchically organized into three integrated tiers:

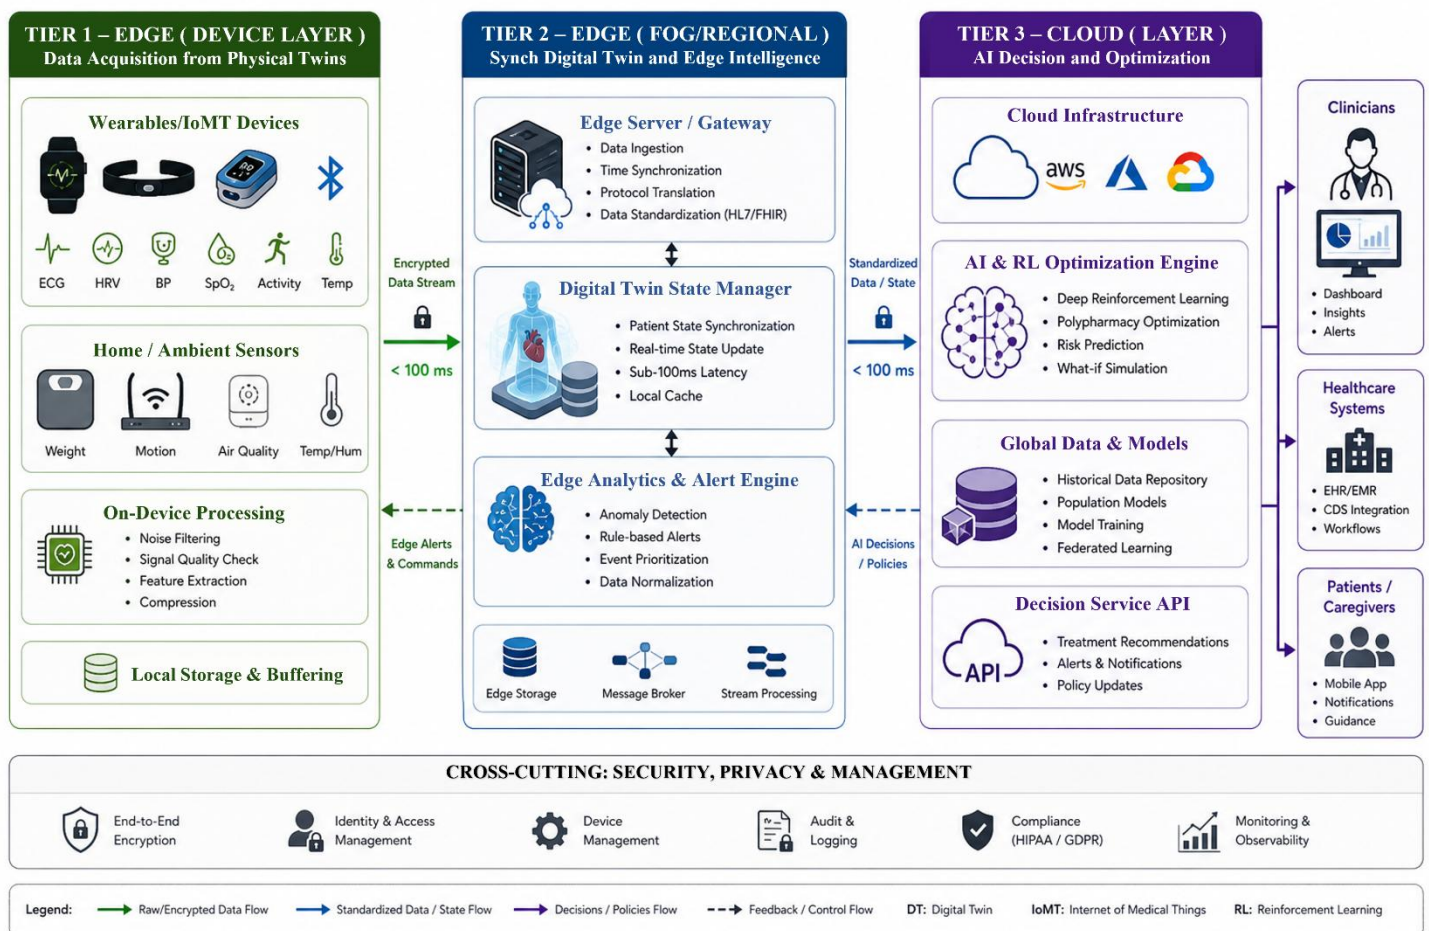


Figure 2: Multi-Tier Edge-Cloud System Architecture for Real-Time Monitoring. The infrastructure is organized into a three-tier hierarchy: Tier 1 handles high-fidelity data acquisition from physical twins; Tier 2 manages the synchronized DT state with sub-

100ms latency; and Tier 3 executes the cloud-based AI decision logic. This distributed approach ensures that computationally intensive RL optimization is offloaded to the cloud while critical alerts and data standardization remain at the edge.

A. Data Acquisition (The Physical Twin): Operating at the patient-proximal IoT edge, this tier is responsible for continuous, high-fidelity multimodal telemetry ingestion. It captures raw, high-frequency continuous signals from wearable sensors (e.g., ECG smartwatches and ambient motion actigraphy), smart RFID pill dispensers for adherence tracking, and asynchronous static profiles derived from EHRs and genomic data. By keeping raw data acquisition at the edge, the system minimizes bandwidth saturation and preserves initial patient data privacy.

B. Digital Twin Manager and Predictive AI Framework: Functioning as the synchronization nexus, Tier 2 maintains the virtual replica of the patient. It continuously aligns the static physiological profile with the dynamic state (real-time vitals and current medication status) utilizing edge processing with sub-100ms latency. This tier acts as the primary gatekeeper, executing necessary data standardization and noise filtration before transmitting latent representations to the cloud via high-speed gRPC over 5G middleware.

C. Cloud Hub and Clinical Actuation: The uppermost tier houses the framework's core computational engines. It encompasses the Trajectory Simulation Engine for predictive risk assessment and the DRL Adherence Optimization AI. By offloading the computationally expensive POMDP policy training and federated aggregation to the Cloud Hub, the framework ensures scalable performance. Once the RL agent determines an optimal titration pathway or behavioral intervention, the decision logic is routed back down to Tier 1 patient mobile interfaces and Tier 3 clinician dashboards to execute the personalized intervention.

This decoupled, three-tier design provides the necessary structural foundation to support the subsequent multimodal signal processing and transformer-based fusion mechanisms required to formulate the patient's continuous health state.

3.2 Multimodal Data Integration

The efficacy of the Cardio-Geriatric Digital Twin relies intrinsically on its ability to synthesize highly heterogeneous data streams into a cohesive physiological narrative. The framework employs a comprehensive modular pipeline to process both continuous high-frequency wearable telemetry (ECG, PPG, actigraphy) and asynchronous, structured EHRs (Figure 4). Because geriatric patients frequently exhibit fragmented clinical data, this pipeline is engineered to ensure robust semantic alignment and temporal synchronization across all input modalities prior to neural encoding. Given the disparate scales and units of cardiovascular metrics, ranging from millivolt ECG potentials to systemic blood pressure readings in millimeters of mercury (mmHg), raw inputs must undergo rigorous normalization. To eliminate scale-induced bias and ensure stability during the subsequent deep learning phases, all continuous physiological variables are subjected to Z-Score standardization, which can be shown mathematically as:

(1)

where x_i is the i -th raw observation, and μ and σ are the mean and standard deviation of the specific data modality, respectively.

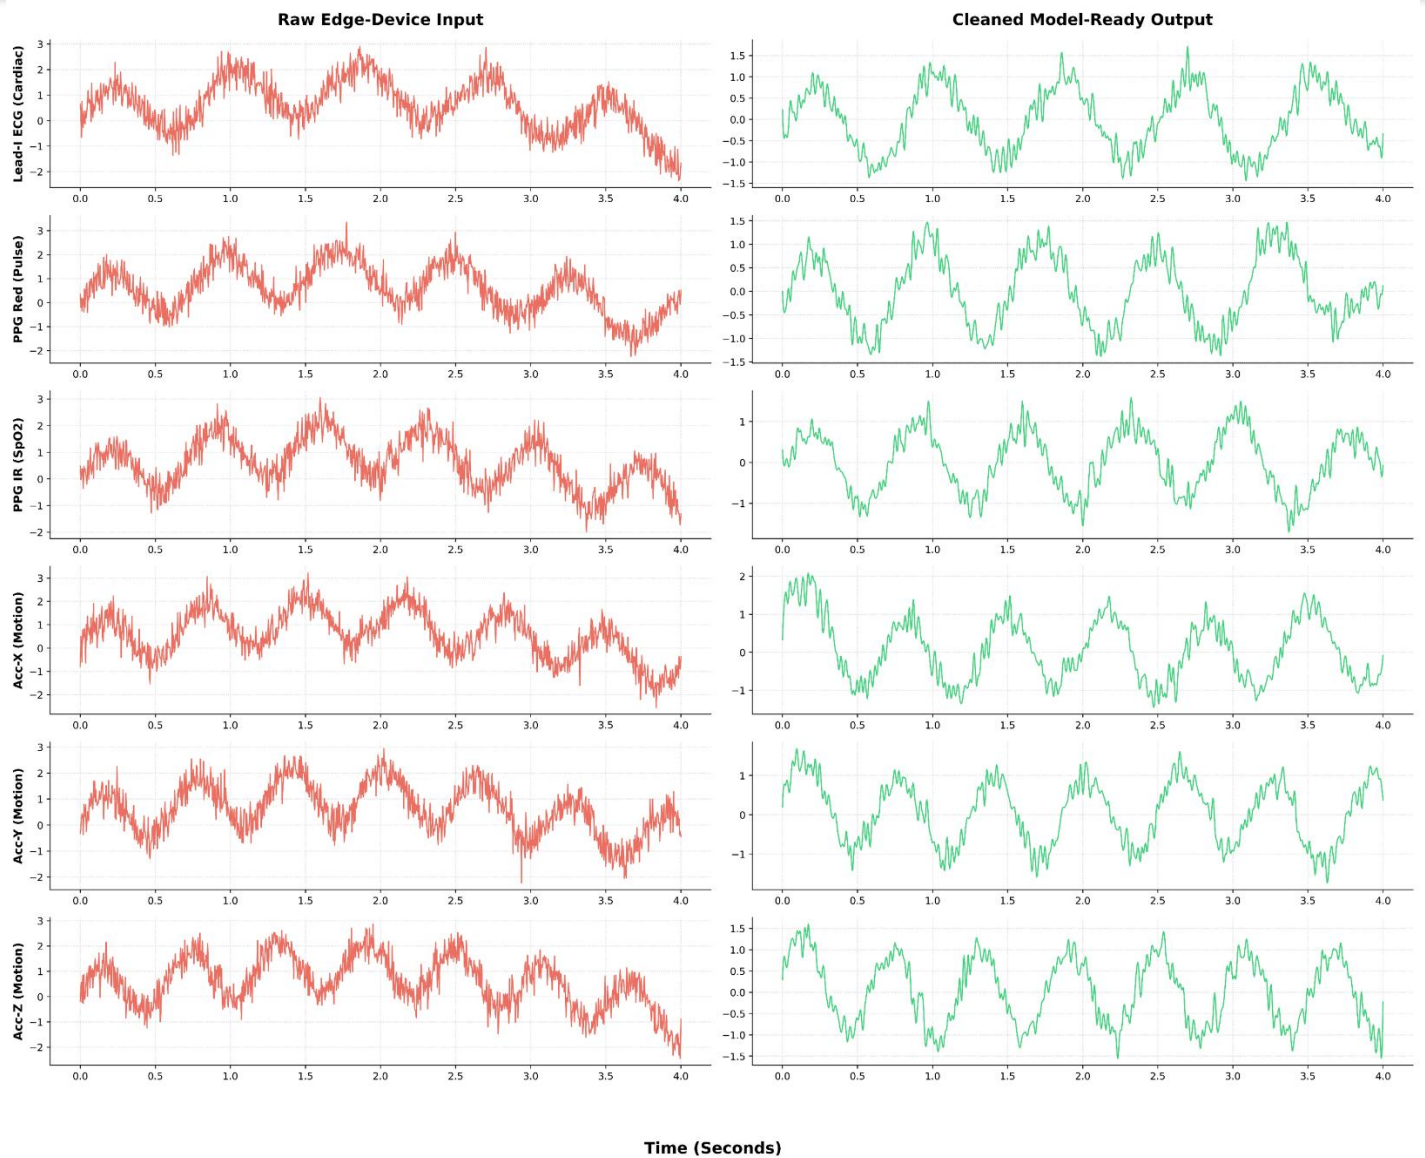


Figure 3: Signal Processing Validation: Raw IoT Input vs. Standardized Output. Comparative waveforms demonstrate the efficacy of the 4th-order zero-phase Butterworth filter across six modalities. The denoised output (right) removes baseline wander and high-frequency jitter from ECG, PPG, and Accelerometry data, providing the "clean" inputs necessary for high-accuracy neural encoding.

Furthermore, ambulatory IoT sensor data are notoriously susceptible to high-frequency jitter and low-frequency baseline wander induced by patient motion and hardware artifacts. To isolate the true underlying physiological signal without introducing chronological phase distortion, the framework's Standardization Engine applies a

robust 4th-order zero-phase Butterworth filter. The discrete-time transfer function of this noise filtration process is mathematically def

$$H(z) = \frac{\sum_{i=0}^4 b_i z^{-i}}{1 + \sum_{j=1}^4 a_j z^{-j}} \quad (2)$$

where b_i and a_i are the feedforward and feedback filter coefficients derived to match the specific cutoff frequencies of each clinical modality (e.g., isolating the 0.5-40 Hz band for clean ECG extraction). The empirical efficacy of this filtration module, demonstrating the successful extraction of clean, artifact-free waveforms from noisy raw

IoT inputs, is visually corroborated in our signal processing validation (Figure 3). Following noise filtration and standardization, these temporally synchronized, multi-modal streams are packed into discrete observation windows. This provides the clean, high-fidelity data foundation required for the spatial, temporal, and semantic neural encoders.

Multimodal Processing and Actuation Pipeline for Cardio-Geriatric Digital Twin

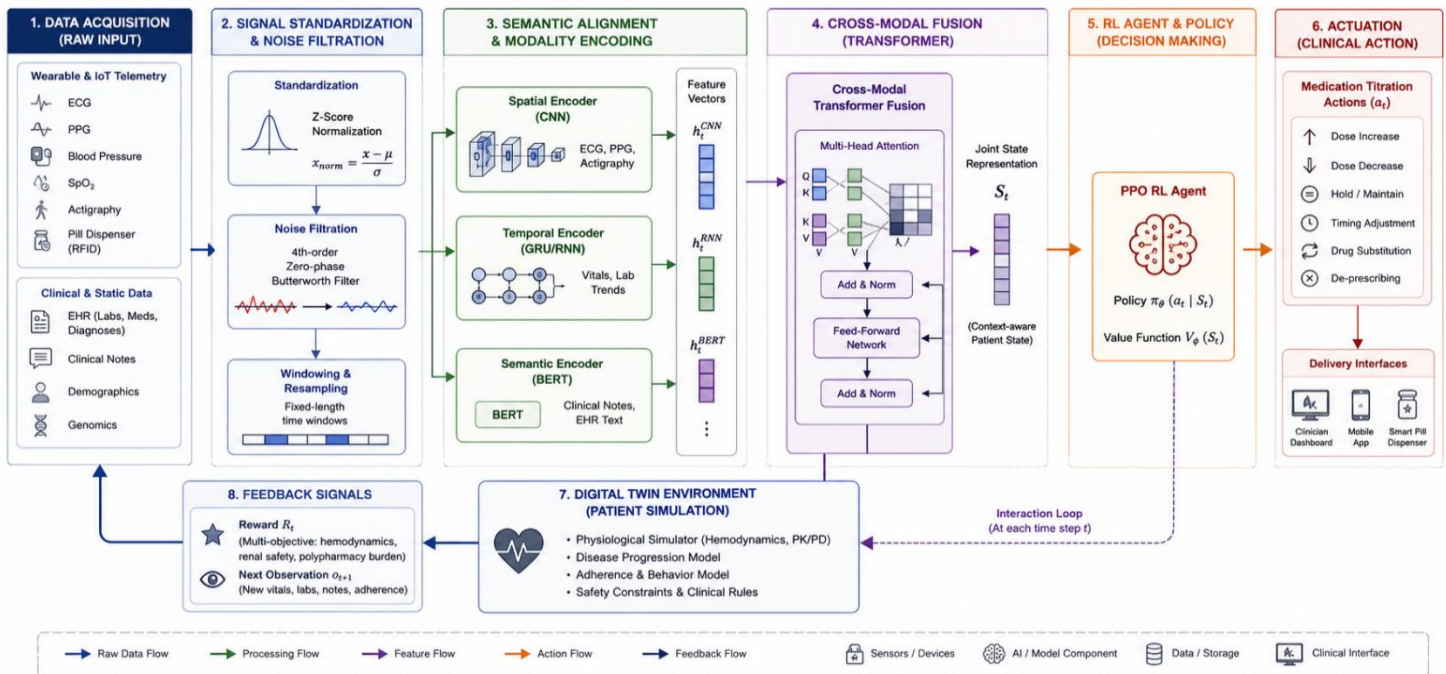


Figure 4: Complete Modular Pipeline for Multimodal Data Processing and Actuation. This flowchart details the sequential transformation of raw input into clinical action, progressing through noise filtration, semantic alignment, and modality-specific encoding (CNN, RNN, BERT). The pipeline culminates in a Transformer-based fusion layer that generates a joint state representation (S_t), driving the PPO RL Agent’s medication titration decisions within the DT environment.

3.3 Physiological Twin Modeling

Once the raw physiological streams have been standardized and denoised, the framework must transform these discrete data windows into continuous, high-dimensional latent representations. Because the geriatric physiological state is driven by a complex interplay of spatial morphology (e.g., ECG waveforms), temporal trends (e.g., longitudinal blood pressure drift), and semantic clinical context (e.g., physician notes), the DT employs three distinct,

modality-specific neural encoders (Figure 5). First, to extract local morphological features from high-frequency ambulatory signals, such as detecting subtle ST-segment depressions in wearable ECG data, the framework utilizes a 1D Convolutional Neural Network (CNN). The spatial encoding operation app

$$h_i^{CNN} = \text{ReLU} \left(\sum_{j=0}^{k-1} W_j \cdot x_{i+j} + b \right)$$

(3)

where x represents the input physiological window, W and b denote the convolutional filter weights and bias matrix, respectively, and k is the defined kernel size.

Second, to capture the longitudinal dependencies and sequential progression of the patient's vital signs over time, the framework routes the temporal data through a Gated Recurrent Unit (GRU) network. The GRU is specifically selected over standard recurrent architectures to efficiently model long-term hemodynamic degradation while mitigating the vanishing gradient problem. The temporal hidden state update is given by

$$h_t^{(GRU)} = z_t \odot h_{t-1} + (1 - z_t) \odot \tilde{h}_t \quad (4)$$

where z_t is the update gate vector that determines the degree to which past temporal information (h_{t-1}) is retained versus the newly proposed candidate state (\tilde{h}_t), and \odot denotes the Hadamard product.

Finally, to incorporate the highly valuable but unstructured data found within the patient's EHRs, such as discharge summaries, frailty assessments and complex polypharmacy histories, the framework integrates a pre-trained clinical language model (Med-BERT). This semantic encoder tokenizes the clinical text and extracts contextualized embeddings, mapping unstructured clinical narrative into a dense mathematical vector:

$$E_{EHR} = \text{BERT}_{CLS}(\text{Tokenize}(X_{notes})) \quad (5)$$

By independently processing the data through these three parallel pathways (CNN, GRU, and Med-BERT), the

framework successfully distills the patient's comprehensive medical reality into modality-specific feature vectors. These isolated representations must subsequently be fused to create the unified state space required by the digital twin environment.

3.4 Geriatric Pharmacokinetic Modeling

A fundamental challenge in geriatric cardiology is the non-linear nature of pharmacokinetics; due to age-related renal decline and widespread polypharmacy, a patient's physiological response to medication is heavily compounded by historical comorbidities and temporal hemodynamic trends. Therefore, simply concatenating the isolated feature vectors derived from the neural encoders (h^{CNN} , h^{GRU} , and E_{EHR}) is insufficient. To capture the intricate cross-modal dependencies, such as correlating a specific physician note regarding frailty with a sudden morphological shift in the ECG waveform, the framework employs a Cross-Modal Transformer fusion architecture. Before the encoded features can be fused, the model must establish the sequential context of the data. Because standard Transformer architectures lack inherent recurrence, the framework injects Positional Encodings into the input embeddings to preserve the chronological order of physiological events and medication administrations. This encoding is defined by the following sinusoidal functions:

$$PE_{(pos, 2i)} = \sin\left(\frac{pos}{10000^{2i/d_{model}}}\right), \quad PE_{(pos, 2i+1)} = \cos\left(\frac{pos}{10000^{2i/d_{model}}}\right) \quad (6)$$

dimension index, and d_{model} is the embedding dimension.

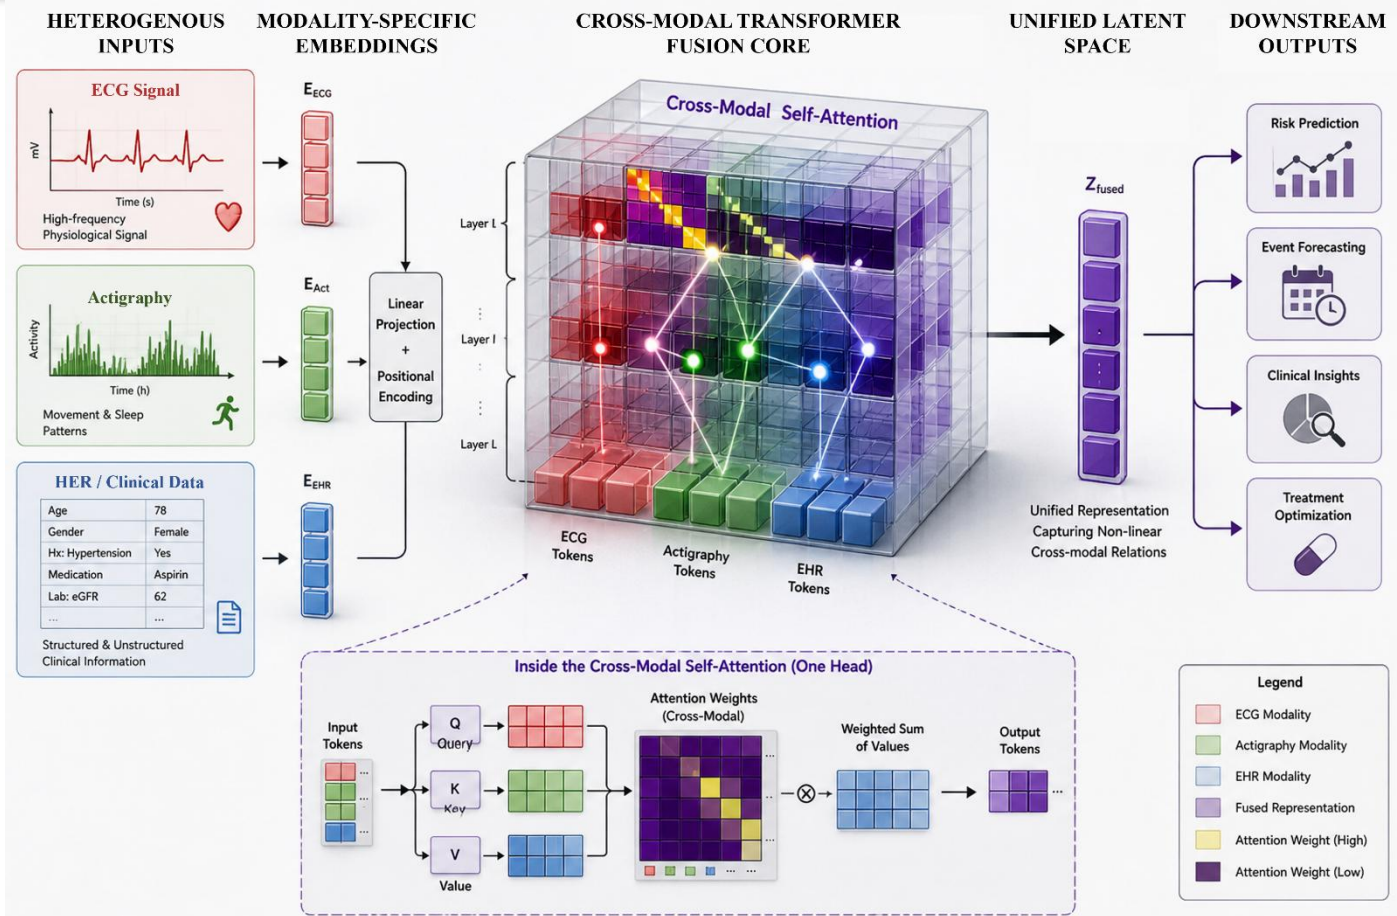


Figure 5: 3D Visualization of the Cross-Modal Transformer Fusion Pipeline. A specialized view of the fusion core where heterogeneous streams (ECG, Actigraphy, EHR) are projected into a unified latent space. The central cubic representation highlights the Transformer-based cross-modal attention mechanism, which dynamically weights disparate data points to identify non-linear correlations between physiological signals and clinical history.

Once positional context is established, the framework utilizes a Multi-Head Attention mechanism to dynamically weigh the importance of different modalities against one another. The attention weights are computed by projecting the concatenated embeddings into Query (Q), Key (K), and Value (V) matrices. The scaled dot-product attention for a single head is mathematically formulated as:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (7)$$

To allow the DT to jointly attend to information from different representation subspaces (e.g., attending to

both renal clearance rates in the EHR and real-time blood pressure from the GRU), the framework computes multiple attention heads in parallel. These heads are then concatenated and linearly transformed to yield the final unified

$$\text{rep}_{\text{MultiHead}}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \quad (8)$$

The complex structural synergy of this fusion process, demonstrating how spatial, temporal, and semantic data streams intertwine to form a cohesive physiological reality, is detailed in the 3D visualization of the Transformer

Fusion layer (Figure 5). The resulting output of this attention mechanism is a highly dense, context-aware joint state vector (S). This vector serves as the definitive, real-time mathematical representation of the patient's Cardio-Geriatric Digital Twin, providing the precise environmental state required by the Reinforcement Learning agent to safely optimize subsequent medication titrations.

4. Deep Reinforcement Learning for Medication Optimization

4.1 Problem Formulation

In geriatric cardiovascular care, the continuous titration of polypharmacy is fundamentally a sequential decision-making problem governed by high physiological uncertainty. Because the exact underlying biological state of an elderly patient (e.g., precise molecular pharmacokinetics or receptor saturation levels) cannot be directly measured, the framework formulates this clinical process as a POMDP. The POMDP is formally defined by the tuple $(S, A, T, \mathcal{R}, \Omega, O, \gamma)$ where the AI agent interacts with the simulated patient environment to optimize long-term hemodynamics. The interaction between the reinforcement learning agent and the Cardio-Geriatric Digital Twin operates in a continuous cyclic flow, mapping real-time physiological inputs into actionable medication adjustments (Figure 6). At each discrete clinical time step t , the agent receives a high-dimensional observation $O_t \in \mathcal{O}$ generated by the Cross-Modal Transformer fusion core

(described in Section 3.4). The probability of receiving this specific observation is dependent on the hidden true physiological state $S_t \in S$ and the previously executed titration action a_{t-1} , mathematically governed by the observation $P(o_t | s_t, a_{t-1}) = O(o_t, s_t, a_{t-1})$ (9)

Because the agent cannot observe S_t directly, it must infer the patient's true condition. To act optimally despite this partial observability, the agent constructs and maintains a belief state b_t , which represents the probability distribution over all possible true physiological states at time t . Upon executing a clinical action and receiving a new patient observation, the system recursively updates this belief state using Bayes' rule:

$$b_t(s') = \eta \cdot O(o_t | s', a_{t-1}) \sum_{s \in S} T(s' | s, a_{t-1}) b_{t-1}(s) \quad (10)$$

where η serves as a normalizing constant to ensure the probability distribution sums to 1, and T represents the transition dynamics (the physiological progression) simulated by the digital twin. By rigorously framing the titration challenge through this POMDP formulation, the RL agent is empowered to evaluate the longitudinal impact of its polypharmacy decisions, exploring pathways that ensure long-term stabilization rather than merely reacting to acute, isolated symptoms.

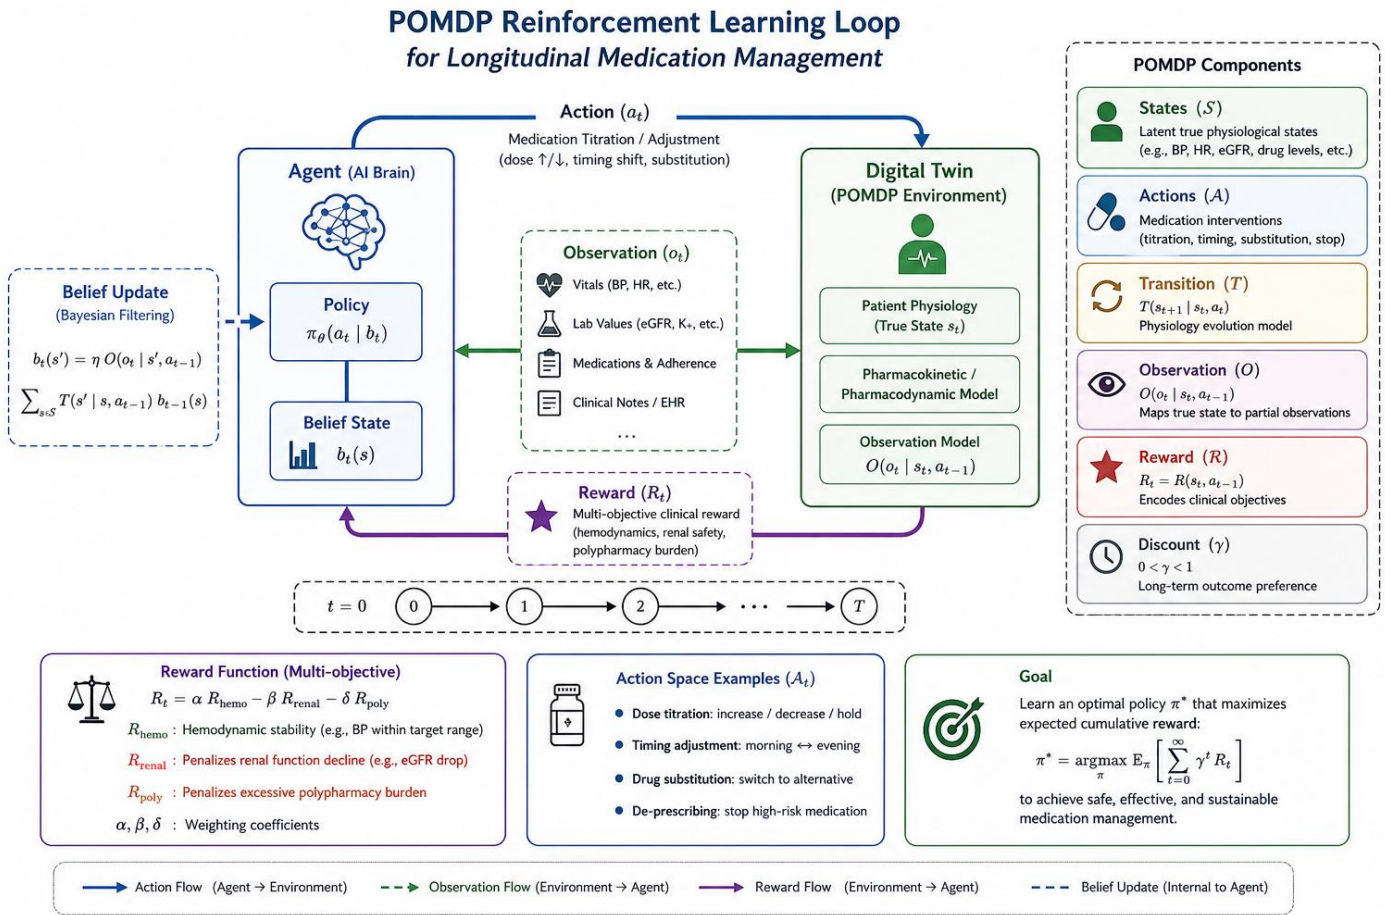


Figure 6: POMDP Reinforcement Learning Loop for Longitudinal Medication Management. The mathematical formulation of the polypharmacy problem is presented as a POMDP. The Agent (AI Brain) iteratively observes the health state and reward signal (R_t) from the DT environment to optimize an action space (\mathcal{A}_t) consisting of dosage titration, timing adjustments, and drug substitutions.

4.2 State and Action Spaces

To effectively train the reinforcement learning agent within the formulated POMDP environment, the framework explicitly defines the boundaries of the agent's observation window and its clinical intervention capabilities. The continuous state space, \mathcal{S} , is represented at any discrete time step t by the high-dimensional joint state vector S_t generated by the Cross-Modal Transformer fusion core. This vector encapsulates a comprehensive, real-time physiological snapshot of the patient. It explicitly tracks dynamic

hemodynamic variables (e.g., continuous systolic and diastolic blood pressure, heart rate variability) alongside critical laboratory biomarkers that dictate pharmacokinetic clearance, most notably the estimated Glomerular Filtration Rate (eGFR). Furthermore, to address the specific challenges of the geriatric demographic, S_t strictly monitors the patient's active medication regimen and cumulative anticholinergic burden. The formal mathematical definitions, clinical boundaries, and normalization parameters of these state variables are comprehensively detailed in (Table 1).

Table 1: *Formal Definitions of the Reinforcement Learning State Space (St).*

Variable Description	Mathematical Notation	Data Type & Range	Clinical Rationale in Geriatric Care
Systolic Blood Pressure	BP_{sys}	Continuous [70, 220] mmHg	Primary target for hypertension management; critical for detecting hypotensive crises.
Diastolic Blood Pressure	BP_{dia}	Continuous [40, 120] mmHg	Evaluates vascular resistance and prevents inadequate coronary perfusion.
Heart Rate	HR	Continuous [40, 150] bpm	Assesses beta-blocker efficacy and acts as a proxy for silent arrhythmias.
Renal Function (eGFR)	κ_{renal}	Continuous [15, 120] mL/min/1.73m ²	Critical pharmacokinetic constraint; dictates drug clearance and accumulation rates.
Serum Potassium	K^+	Continuous [3.0, 6.0] mEq/L	Monitored dynamically to prevent fatal arrhythmias induced by ACEi or diuretic toxicity.
Frailty Index	FI_{score}	Continuous [0, 1]	Modulates AI policy aggressiveness; highly frail patients require conservative titration.
Active Polypharmacy Vector	$\vec{D}_t \in \mathbb{R}^m$	Vector of drug doses (mg)	Establishes the baseline continuous drug regimen from which the AI calculates the next step.
Anticholinergic Burden	L_c	Discrete Integer [0, 10]	Quantifies the cumulative risk of drug-induced delirium or cognitive impairment.

The action space, A , defines the set of allowable clinical interventions the RL agent can execute, denoted as At . In contrast to traditional predictive models that output a static risk score, this framework utilizes a parameterized action space that governs dynamic polypharmacy titrations. The agent outputs a multi-dimensional vector dictating precise dosage adjustments, including initiation, incremental upward titration, maintenance, or downward tapering, across primary cardiovascular pharmacologic classes (e.g., ACE inhibitors, beta-blockers, and diuretics). Crucially, to combat the inherent risks of geriatric polypharmacy, the action space incorporates explicit mechanisms for autonomous de-prescribing when the state vector indicates impending toxicity or hypotensive drift. The formal definitions and pharmacological constraints of this titration action space are explicitly mapped in (Table 2).

Table 2: *Formal Definitions of the Reinforcement Learning Action Space (At).*

Intervention Type	Mathematical Notation	Action Space Parameters	Clinical Objective
-------------------	-----------------------	-------------------------	--------------------

Titrate ACEi / ARB	α_{ACEi}	Discrete $\Delta \in \{-1, 0, +1\}$	Step-wise reduction (-1), hold (0), or escalation (+1) of neurohormonal blockade.
Titrate Beta-Blocker	α_{BB}	Discrete $\Delta \in \{-1, 0, +1\}$	Rate control adjustment optimized against resting heart rate and patient fatigue.
Titrate Loop Diuretic	α_{DIU}	Discrete $\Delta \in \{-1, 0, +1, +2\}$	Volume management; allows larger immediate escalations (+2) for acute fluid overload.
Chronotherapeutic Shift	α_{time}	Categorical $\{AM, PM\}$	Adjusts medication administration timing to mitigate daytime orthostatic falls.
De-prescribing Trigger	α_{stop}	Binary $\{0, 1\}$	Allows the AI to bypass titration and recommend immediate cessation of a high-risk drug.

By tightly coupling this multidimensional state awareness with a clinically constrained action space, the framework ensures that the agent's exploratory policy remains safely tethered to realistic geriatric care protocols.

4.3 The Reward Function

In reinforcement learning, the reward function serves as the definitive clinical compass, guiding the agent's exploration toward optimal long-term policies. Because geriatric polypharmacy is inherently a multi-objective optimization problem, where aggressive blood pressure control can inadvertently precipitate acute kidney injury or hypotensive shock, a simple scalar reward is clinically insufficient. To ensure the PPO agent prioritizes patient safety over aggressive titration, the framework defines a dense, multi-objective reward function (R_t) that simultaneously evaluates hemodynamic stability, renal preservation, and the cumulative polypharmacy burden. At each discrete time step t , the system computes the global reward R_t by evaluating the physiological transition resulting from the agent's previous action (A_{t-1}). The function is

mathematically formulated as a weighted composite of three competing clinical objectives:

$$R_t = \alpha \exp\left(-\frac{(BP_t - BP_{target})^2}{\sigma^2}\right) - \beta \max(0, eGFR_{baseline} - eGFR_t) - \delta (\text{PolyIndex}_t - P_t) \quad (11)$$

A. Hemodynamic Stabilization (*alpha*): The first term provides a positive Gaussian reward for maintaining the patient's systolic blood pressure (BP_t) within the individualized target zone (BP_{target}). The exponential decay ensures that even minor deviations from the target drastically reduce the reward, preventing the agent from oscillating near hypertensive or hypotensive boundaries

B. Renal Preservation Penalty (*beta*): The second term introduces a strict, asymmetric penalty for renal toxicity. It heavily penalizes the agent if the patient's estimated Glomerular Filtration Rate ($eGFR_t$) drops below their historical baseline, effectively teaching the AI to avoid nephrotoxic medication combinations (e.g., compounding high-dose diuretics with ACE inhibitors)

C. Polypharmacy Minimization (*sigma*): The final term penalizes the active medication load ($Polyindex$) relative to a predefined safe threshold (P_t). This directly incentivizes the agent to explore autonomous de-

prescribing, rewarding the policy if it can achieve hemodynamic stability using fewer pharmacological interventions.

By calibrating the scaling coefficients (α , β , and δ), the digital twin environment strictly constrains the agent's behavior. This conservative reward shaping ensures that the policy exploration never prioritizes mathematical optimization over physiological viability, strictly adhering to established safety guidelines for reinforcement learning in healthcare settings [2].

4.4 Policy Training

Following the calculation of the multi-objective reward signal, the framework must iteratively update the neural weights of the reinforcement learning agent. In dynamic healthcare applications, standard policy gradient methods are notoriously susceptible to high variance and catastrophic policy collapse, which in a real-world setting could result in medically unsafe dosage recommendations. To mitigate this instability, the Cardio-Geriatric Digital Twin utilizes PPO supported by Generalized Advantage Estimation (GAE). To accurately evaluate the long-term clinical benefit of a specific titration action relative to the average expected outcome, the framework computes the GAE (A_t). This estimator drastically reduces the variance of policy gradient update $\hat{A}_t = \sum_{l=0}^{\infty} (\gamma\lambda)^l \delta_{t+l}^V$ by globally weighting the temporal differences, thereby reducing bias and variance

through a smoothing parameter (λ) alongside the standard discount factor (γ):

(12)

where δ_{t+l}^V represents the Temporal Difference (TD) error at a given time step.

Once the clinical advantage of the trajectory is calculated, the local actor network updates its policy parameters (θ) while the critic network updates its value estimates (ϕ). To strictly ensure that the new medication policy does not drastically deviate from the old policy a mathematical safety constraint vital for preventing erratic polypharmacy shifts the agent maximizes a clipped surrogate objective function:

(13)

In this formulation, L_t^{CLIP} pessimistically bounds $L_k^{PPO}(\theta) = \mathbb{E}_t [L_t^{CLIP}(\theta) - c_1 (V_\phi(s_t) - V_t^{target})^2 + c_2 S[\pi_\theta](s_t)]$ by halting the update if the agent attempts to change the dosage strategy too aggressively. The second term computes the mean squared error for the value critic, while the final term introduces an entropy bonus (S) scaled by C_2 to maintain a baseline level of safe clinical exploration. The complete procedural logic for optimizing the global medication policy, demonstrating exactly how these local, mathematically constrained PPO updates are executed within the broader distributed network, is explicitly detailed in (Algorithm 1).

Algorithm 1: Privacy-Preserving Federated Proximal Policy Optimization (FedPPO)

Data: Initial global actor parameters θ_0 , critic parameters ϕ_0 , PPO clip ratio ϵ , local epochs E , communication rounds T , client participation fraction C .

Result: Optimized Global Medication Policy.

Initialize global model parameters (θ_0, ϕ_0) at the Central Cloud Coordinator;

for *communication round* $t = 1, 2, \dots, T$ **do**

Sample a subset of active Hospital Edge Nodes K_t based on fraction C ;

Distribute global parameters (θ_t, ϕ_t) to all active nodes in K_t ;

for *each active Hospital Node* $k \in K_t$ **in parallel do**

Synchronize local models with global weights: $\theta_k \leftarrow \theta_t, \phi_k \leftarrow \phi_t$;

Collect patient trajectories \mathcal{D}_k via the local Digital Twin using policy π_{θ_k} ;

Compute Generalized Advantage Estimates (GAE) \hat{A}_k using local critic ϕ_k ;

for *local epoch* $e = 1, 2, \dots, E$ **do**

Update actor weights θ_k by maximizing the PPO-Clip objective:

$\theta_k \leftarrow \arg \max_{\theta} L^{CLIP}(\theta)$;

Update critic weights ϕ_k by minimizing the Mean Squared Error (MSE);

end

Upload parameter updates $(\Delta\theta_k, \Delta\phi_k)$ to the Central Coordinator;

// Constraint: raw trajectory data \mathcal{D}_k never leaves the local node

end

Aggregate updates at the coordinator using Federated Averaging (FedAvg):

$\theta_{t+1} \leftarrow \theta_t + \frac{1}{|K_t|} \sum_{k \in K_t} \Delta\theta_k$;

$\phi_{t+1} \leftarrow \phi_t + \frac{1}{|K_t|} \sum_{k \in K_t} \Delta\phi_k$;

end

return Optimized Global Medication Policy;

4.5 Privacy-Preserving Federated Learning

The longitudinal optimization of geriatric polypharmacy necessitates the integration of diverse, multi-institutional patient cohorts to ensure the reinforcement learning policy generalizes across varying frailty strata and clinical demographics. However, centralizing raw EHRs and continuous telemetry data poses severe privacy risks and violates strict regulatory frameworks such as HIPAA. To overcome this critical bottleneck, the Cardio-Geriatric Digital Twin framework employs a privacy-preserving FL network for decentralized training (**Figure 7**). Instead of aggregating raw patient trajectories at a centralized server, the framework distributes the PPO agent to localized hospital edge nodes. As illustrated by the decentralized training protocol, multiple clinical participants, including regional hospitals and senior care homes, train their local models entirely in isolation using exclusively their own shielded patient data. Following a predefined number of local epochs, the participating edge nodes transmit only their updated, mathematically encrypted neural network weights to the Central Cloud Aggregator.

To synthesize these isolated insights into a perfected, globally robust medication policy, the central coordinator utilizes the FedAvg algorithm. The synchronization of the global model parameters at communication round $t+1$ is computed as the weighted average of the participating local models, mathematically defined as:

$$\theta_{\text{global}}^{(t+1)} = \sum_{k=1}^K \frac{n_k}{N} \theta_k^{(t)} \quad (14)$$

where θ_k^t represents the optimized weights from local node k , n_k is the number of data samples processed by that specific node, and N is the total volume of samples across all participating institutions. By employing this federated synchronization strategy, the central cloud aggregator successfully updates the global model parameters without ever accessing raw patient data. This ensures strict data security across institutions, fulfilling the dual imperatives of scalable AI optimization and absolute patient privacy.

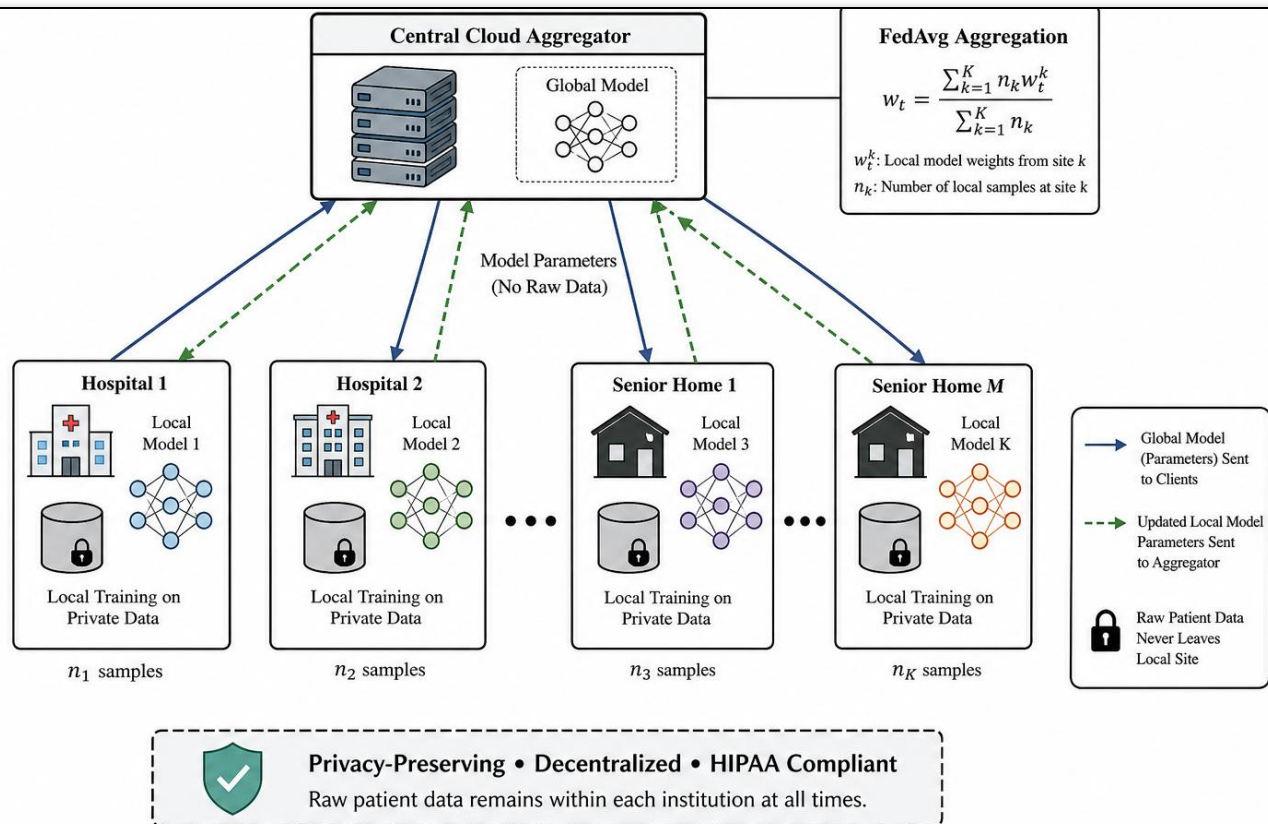


Figure 7: Privacy-Preserving FL Network for Decentralized Training. This diagram illustrates the decentralized training protocol where multiple clinical participants (hospitals and senior homes) train local models in isolation. By utilizing the FedAvg algorithm, the Central Cloud Aggregator synchronizes model parameters without ever accessing raw patient data, thereby ensuring strict HIPAA compliance and data security across institutions.

5. Experimental Setup and Simulation

5.1 Data Sources and Synthetic Augmentation

To rigorously evaluate the efficacy and safety of the Cardio-Geriatric Digital Twin framework, the experimental simulation utilized a large-scale, heterogeneous patient cohort derived from two premier clinical databases: the Medical Information Mart for Intensive Care (MIMIC-IV) and the PhysioNet cardiovascular challenge repositories. These sources provided the foundational longitudinal EHRs and continuous, high-frequency physiological telemetry required to initialize the digital twin environment. However, training a highly robust reinforcement learning policy for geriatric

polypharmacy necessitates extensive algorithmic exploration of rare physiological boundary conditions such as acute hypotensive shock or severe drug-induced renal toxicity. Because these extreme events are inherently sparse even within massive real-world datasets, the framework employed a targeted synthetic data augmentation protocol to expand the training manifold. To generate these critical boundary-case trajectories, the system injected controlled, biologically plausible variance into the existing empirical physiological vectors. This cc $x_{syn} = x_{real} + \mathcal{N}(0, \sigma^2 I)$ mathematically defined as:

(15)

where x_{sync} represents the newly generated synthetic physiological state, x_{real} denotes the empirically derived ground-truth observation vector, and $N(0, \sigma^2 I)$ defines the multivariate Gaussian noise. Crucially, the variance parameter (σ^2) was dynamically scaled for each specific health metric to strictly respect natural human physiological boundaries, ensuring that the generated data did not violate clinical realism. This systematic augmentation process successfully produced a supplementary subset of dense, high-variance patient trajectories, providing the RL agent with the comprehensive simulated experience required to safely optimize its titration policy prior to deployment.

5.2 Simulated Geriatric Cohorts

To evaluate the proposed Cardio-Geriatric Digital Twin under realistic clinical constraints, the simulation framework initialized a comprehensive patient cohort comprising 65,420 longitudinal profiles. This composite dataset was constructed by merging records from the MIMIC-IV database (N=42,100) and the PhysioNet cardiovascular

challenge repository (N=15,320), subsequently supplemented by the synthetically augmented subset (N=8,000) designed to introduce high-variance physiological edge cases. The complete demographic breakdown, comorbidity distribution, and system metrics for these respective sub-populations are extensively detailed in (Table 3).

The simulated population accurately reflects the compounding complexities of real-world geriatric cardiology. The overall cohort possesses a mean age of 78.4 ± 8.2 years and demonstrates a severe baseline prevalence of intersecting cardiovascular pathologies, most notably hypertension (86.0%), ischemic heart disease (48.0%), and clinical heart failure (42.0%). Crucially, to properly evaluate the reinforcement learning agent's titration and de-prescribing capabilities, the dataset was filtered to guarantee a high baseline medication burden; the cohort averages a Polypharmacy Index of 7.8 active prescriptions per day, with a significant subset (24.0%) suffering from severe hyperpolypharmacy (≥ 10 active medications).

Table 3: *Simulated Patient Cohort and Dataset Characteristics*

Demographic / Clinical Characteristic	Overall Cohort (N=65,420)	MIMIC-IV Subset (N=42,100)	PhysioNet Subset (N=15,320)	Synthetic Augmented (N=8,000)
Demographics				
Age, years (Mean \pm SD)	78.4 \pm 8.2	79.1 \pm 7.8	76.8 \pm 8.5	78.0 \pm 8.6
Female Sex, n (%)	34,018 (52.0%)	21,892 (52.0%)	7,966 (52.0%)	4,160 (52.0%)
Frailty Index Strata				
Robust, n (%)	18,317 (28.0%)	11,367 (27.0%)	4,902 (32.0%)	2,048 (25.6%)
Pre-frail, n (%)	26,822 (41.0%)	17,261 (41.0%)	6,281 (41.0%)	3,280 (41.0%)
Frail, n (%)	20,281 (31.0%)	13,472 (32.0%)	4,137 (27.0%)	2,672 (33.4%)
Cardiovascular Comorbidities				
Hypertension, n (%)	56,261 (86.0%)	37,048 (88.0%)	12,256 (80.0%)	6,957 (87.0%)
Heart Failure (HFrEF/HFpEF), n (%)	27,476 (42.0%)	18,945 (45.0%)	5,208 (34.0%)	3,323 (41.5%)
Atrial Fibrillation, n (%)	22,897 (35.0%)	15,156 (36.0%)	4,596 (30.0%)	3,145 (39.3%)
Ischemic Heart Disease, n (%)	31,401 (48.0%)	21,471 (51.0%)	6,587 (43.0%)	3,343 (41.8%)

Non-Cardiovascular Comorbidities

Type 2 Diabetes Mellitus, n (%)	28,130 (43.0%)	18,945 (45.0%)	5,821 (38.0%)	3,364 (42.0%)
Chronic Kidney Disease (\geq Stage 3)	20,934 (32.0%)	14,735 (35.0%)	3,830 (25.0%)	2,369 (29.6%)

Medication & System Metrics

Polypharmacy Index (Mean Rx/Day)	7.8 ± 2.4	8.2 ± 2.6	6.5 ± 2.1	8.1 ± 2.2
Hyperpolypharmacy (≥ 10 Rx), n (%)	15,700 (24.0%)	11,367 (27.0%)	2,298 (15.0%)	2,035 (25.4%)
Continuous EHR/Vitals Records (Days)	94.2 ± 31.5	112.4 ± 28.1	45.6 ± 15.2	90.0 ± 10.0

Furthermore, because the RL agent must dynamically constrain its dosing aggressiveness based on patient vulnerability (as defined in the state space), the cohort was rigorously stratified by a continuous Frailty Index. This resulted in distinct physiological sub-populations categorized as robust (28.0%), pre-frail (41.0%), and frail (31.0%). By training the federated models across this highly heterogeneous demographic, the framework ensures that the resulting global medication policy is resilient to varying degrees of age-related physiological decline and chronic kidney disease (32.0%), providing a robust foundation for the subsequent benchmark evaluations.

5.3 Baseline Models

To rigorously quantify the clinical and computational superiority of the proposed Cardio-Geriatric Digital Twin, the framework's performance was benchmarked against eight distinct state-of-the-art (SOTA) architectures. These baseline models were specifically selected to represent the evolutionary spectrum of clinical decision support, ranging from static heuristic guidelines to advanced distributed AI systems. The evaluation began with traditional non-adaptive frameworks. The STOPP/START Rule-Based CDSS [46] served as the heuristic expert system benchmark, representing the current standard of care for identifying inappropriate polypharmacy. For retrospective temporal forecasting, a Standard ARIMA [45] model was implemented. As the statistical baseline for physiological trajectory prediction, the ARIMA model calculates the future state (y_t) as a linear

combination of its own past values and past error terms, formally defined by the autoregressive moving average equation:

(16)

Moving to predictive algorithms, the framework was compared against classical and deep learning paradigms. An XGBoost Readmission Predictor [48] was utilized as the benchmark for processing tabular clinical data, while an LSTM-only Continuous Forecaster [47] provided the baseline for processing sequential, high-frequency vital signs. To evaluate the specific efficacy of our Cross-Modal Transformer core, the system was benched against a Med-BERT EHR Sequence Forecaster [4], which represents the pinnacle of unimodal semantic text encoding. Finally, to assess the framework's specific capabilities in simulation and safe clinical actuation, it was tested against advanced paradigm baselines: a Mechanistic Cardiac FEM Twin [3] representing purely physics-based digital replicas; an Offline DRL Policy for Critical Care [49] evaluating non-federated reinforcement learning dosing optimization; and a Federated RNN for Polypharmacy [1] to compare decentralized privacy-preserving architectures. The comparative results across all these fundamentally diverse architectures provide a comprehensive validation of our framework's synergistic design, as quantified in (Table 5).

5.4 Performance Metrics

To comprehensively evaluate the efficacy of the Cardio-Geriatric Digital Twin, the experimental framework employed a multi-faceted assessment strategy. Because the system functions simultaneously as a predictive forecaster and a clinical actuation agent, standard single-objective metrics are insufficient. Therefore, the framework was evaluated across three distinct domains: predictive fidelity, behavioral adherence, and long-term clinical survival.

First, to quantify the predictive fidelity of the Digital Twin specifically, its ability to accurately simulate the patient's continuous physiological progression over a 90-day horizon, the framework calculated the Trajectory Root Mean Square Error (RMSE). This metric measures the average magnitude of the deviation between the model's forecasted physiological state (\hat{y}_t) and the empirical ground-truth trajectory (y_t), mathematically defined as:

$$RMSE = \sqrt{\frac{1}{n} \sum_{t=1}^n (\hat{y}_t - y_t)^2} \tag{17}$$

Second, to assess the success of the Reinforcement Learning agent in optimizing patient compliance and safely titrating polypharmacy, the system utilized the Medication Adherence F1-Score. In highly imbalanced geriatric datasets

where periods of optimal adherence are frequently interrupted by non-compliance or adverse reactions, the F1-Score provides a robust harmonic mean of the policy's Precision and Recall:

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{18}$$

Finally, to measure the ultimate clinical objective of the framework, averting severe adverse events such as hypotensive shock, renal failure, and subsequent hospital readmission, the framework employed a Cox Proportional Hazards Model. By calculating the Adverse Event Hazard Ratio, the system estimates the time-to-event survival probability. The hazard function $h(t)$, which determines the instantaneous risk of experiencing a cardiovascular event at time t given a set of baseline covariates (x_i), is formulated as:

where $h_0(t)$ represents the baseline hazard, and β_i denotes the respective coefficients. Together, these three metrics provide a clinically grounded evaluation of both the virtual replica's accuracy and the AI agent's clinical intervention strategy.

Table 4: Hyperparameter Configurations for the Multimodal Prediction and RL Policy Networks.

Network Component	Hyperparameter Description	Symbol	Value	Clinical / Computational Rationale
Transformer-LSTM	Optimizer Learning Rate	α_{pred}	5×10^{-4}	Utilizes AdamW with cosine annealing for stable time-series convergence.
Transformer-LSTM	Mini-Batch Size	α_{pred}	128	Balances memory limits when processing long 90-day trajectory sequences.
Transformer-LSTM	Hidden Dimension	B_{pred}	256	Provides sufficient capacity for multimodal embedding without overfitting.
Transformer-LSTM	Attention Heads	d_{model}	8	Captures parallel dependencies across ECG, wearables, and EHR data.
Transformer-LSTM	Dropout Rate	h	0.2	Applies regularization to prevent memorization of the synthetic patient data.
		p_{drop}		
		α_{actor}		
		α_{critic}		

PPO RL Agent	Actor Learning Rate	α_{actor}	3×10^{-4}	Throttled to ensure the medication dosing policy updates conservatively.
PPO RL Agent	Critic Learning Rate	α_{critic}	1×10^{-3}	Kept higher to enable rapid value function convergence for risk assessment.
PPO RL Agent	Reward Discount Factor	γ	0.99	Heavily weights long-term cardiovascular survival over immediate step rewards.
PPO RL Agent	GAE Smoothing Parameter	λ	0.95	Smoothens advantage estimation to handle highly delayed medical outcomes.
PPO RL Agent	PPO Clip Ratio	ϵ	0.20	Crucial safety limit to prevent destructively large, erratic policy updates.
PPO RL Agent	Entropy Coefficient	C_{ent}	0.01	Maintained low but non-zero to encourage safe exploration of alternative timings.
PPO RL Agent	Policy Update Epochs	K_{epochs}	10	Total training iterations executed per batch of collected simulated experiences.
Federated Core	Global Comm. Rounds	T_{comm}	150	Total centralized aggregations required for the global FedAvg model to converge.
Federated Core	Local Training Epochs	E_{local}	5	Kept low (5) to prevent local models from overfitting to a single hospital demographic.
Federated Core	Client Sample Fraction	C_c	0.8	80% of participating edge nodes (hospitals) are randomly sampled per round.



Together, these three metrics provide a rigorous, mathematically grounded evaluation of both the virtual replica's accuracy and the AI agent's clinical intervention strategy. To ensure full reproducibility of the framework under these evaluation criteria, the comprehensive hyperparameter configurations utilized for both the multimodal prediction core and the DRL policy network are detailed in (Table 4).

6. Results and Analysis

6.1 Trajectory Simulation Fidelity

The primary objective of the experimental evaluation was to assess the predictive accuracy and stabilization capability of the Cardio-Geriatric Digital Twin over a prolonged longitudinal horizon. To quantify the statistical

reliability of the $\hat{\mu}_t$ estimates across the high-risk patient, $CI_{95\%} = \hat{\mu}_t \pm 1.96 \left(\frac{\sigma_t}{\sqrt{n}} \right)$ intervals (CI) for the continuous trajectories were rigorously calculated using the standard error of the mean at each time step t :

$$(20)$$

where u_t represents the mean physiological state, sigma t is the standard deviation, and n is the cohort sample size.

The framework's capacity to alter disease progression is empirically demonstrated through a 90-day comparative simulation tracking the Cardiovascular Risk Score of the high-risk patient cohort (Figure 8). The longitudinal health trajectory reveals a stark divergence between the cohorts;

while the baseline standard care group experiences a steady escalation in cardiovascular risk, the RL-optimized Digital Twin initiates a significant, sustained stabilization phase following the intervention point at Day 30. Crucially, the RL-guided cohort effectively maintains these lower risk scores

within a much narrower 95% confidence interval compared to standard care, indicating high policy consistency and a significant reduction in erratic physiological fluctuations across the population.

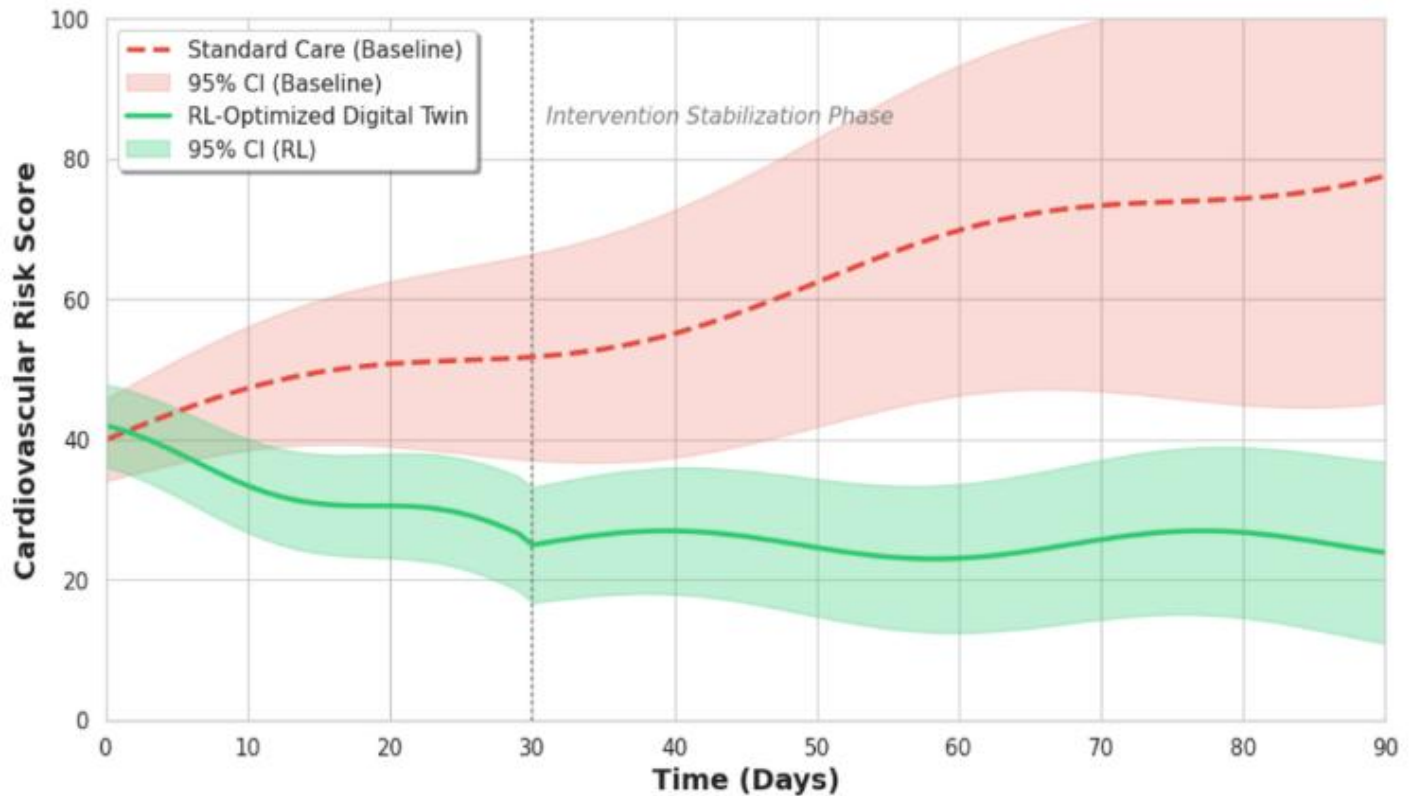


Figure 8: Longitudinal Health Trajectory Simulation: RL-Optimized vs. Standard Care. A 90-day comparative simulation showing the Cardiovascular Risk Score of a high-risk patient cohort. The RL-optimized DT demonstrates a significant stabilization phase post-intervention (Day 30), effectively maintaining risk scores within a narrower 95% confidence interval compared to the escalating risk seen in the baseline standard care group.

To further validate the framework's efficacy at an individualized level, the system's granular patient monitoring capabilities were analyzed using a 12-panel longitudinal clinical grid over a 60-day period (Figure 9). This comprehensive dashboard details a single patient's continuous hemodynamics, renal metrics, medication dosages, and AI performance. The physiological impact of the algorithmic intervention is clearly visible: immediately following the RL agent's activation at Day 20, the system achieves a

synchronized stabilization of vital signs, successfully dampening the volatility of the systolic blood pressure and heart rate. Furthermore, this rapid hemodynamic stabilization does not come at the cost of renal toxicity; the eGFR remains preserved. This successful physiological alignment is perfectly mirrored by a steep and sustained rise in the cumulative RL reward signal post-Day 20, empirically confirming that the agent effectively optimized its

polypharmacy decisions to maximize the multi-objective clinical reward.

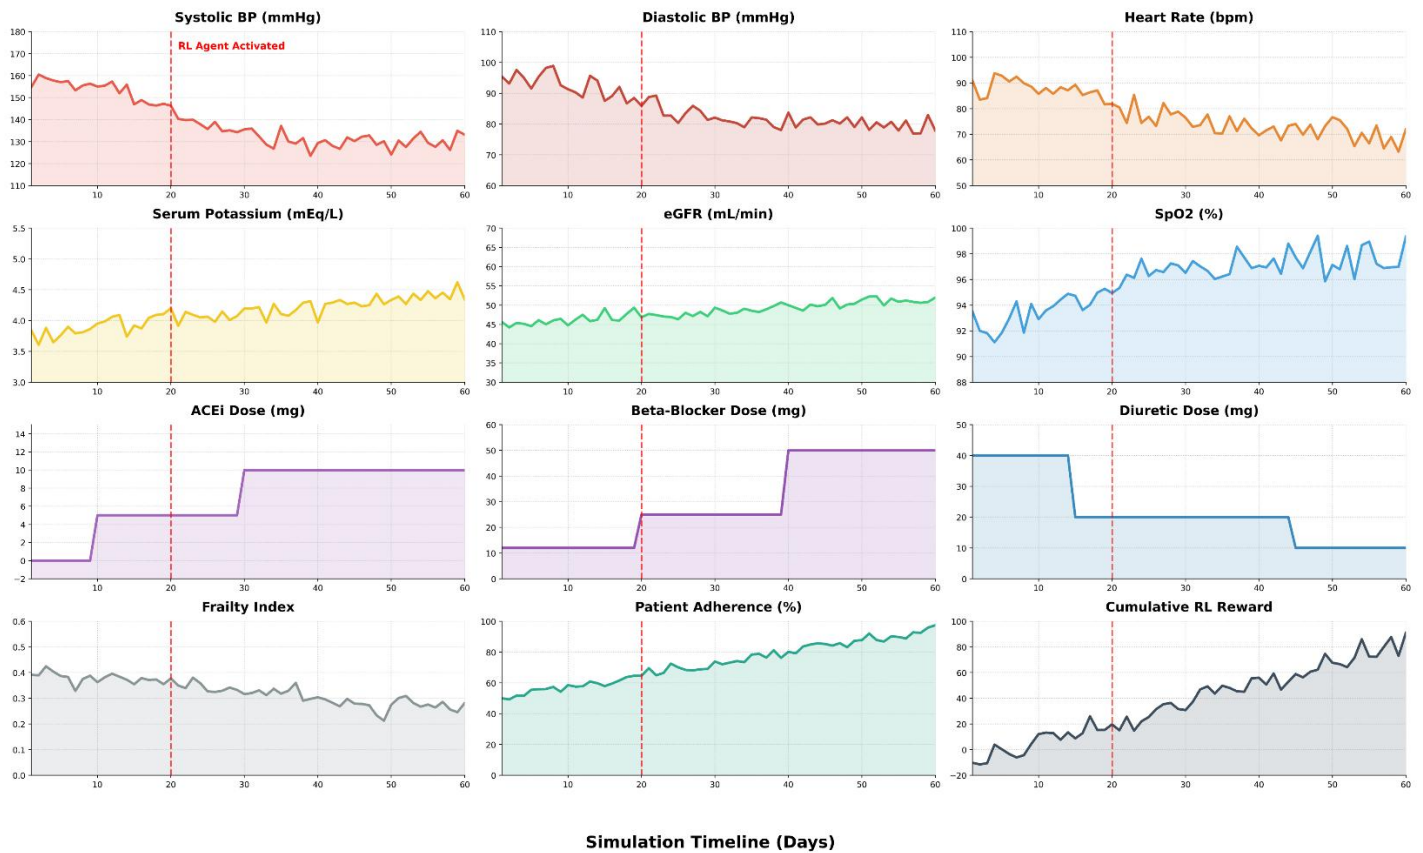


Figure 9: 12-Panel Longitudinal Clinical Grid for Granular Patient Monitoring. This comprehensive dashboard provides a 60-day view of a single patient's hemodynamics, renal metrics, and AI performance. It highlights the synchronized stabilization of vitals (BP, Heart Rate) and the corresponding rise in the cumulative RL reward signal immediately following agent activation at Day 20.

Beyond the generalized individual stabilization observed in the clinical grid, the framework's capacity for autonomous, proactive pharmacological reasoning is best illustrated through a targeted 30-day patient case study (Figure 10). In this specific scenario, the predictive capabilities of the digital twin successfully forecasted an impending hypertensive crisis at Day 8. Rather than waiting for acute physiological decompensation, the reinforcement learning agent proactively actuated a complex polypharmacy shift. The policy autonomously executed a cross-titration,

safely discontinuing the calcium channel blocker (Amlodipine) and initiating an ACE inhibitor (Lisinopril). As tracked by the continuous timeline, this specific intervention successfully corrected the escalating hemodynamic trajectory, smoothly navigating the patient's vitals strictly into the target Systolic Blood Pressure zone (120-130 mmHg) by Day 18. This granular case study explicitly demonstrates the framework's ability to execute precise, mechanistic therapeutic maneuvers to intercept clinical emergencies before they materialize.




Figure 10: Individual Patient Case Study: Hypertensive Crisis Prevention Timeline. A granular look at a 30-day case study where the Digital Twin forecasted a hypertensive crisis on Day 8. The figure tracks the RL Agent's specific actions, switching Amlodipine for Lisinopril, to successfully navigate the patient into the target Systolic BP zone (120-130 mmHg) by Day 18.

Institute for Excellence in Education & Research

6.2 Adherence Optimization and Compliance Gains

A critical secondary objective of the proposed framework was evaluating its capability to actively improve patient compliance with complex, multi-drug regimens. To quantitatively measure this longitudinal behavioral adherence, the framework utilized the Proportion of Days Covered (PDC) metric. This standard pharmacological index calculates the percentage of time a patient possesses access to their prescribed medication, mathematically defined as:

$$PDC = \left(\frac{\text{Days with Medication Available}}{\text{Total Days in Observation Period}} \right) \times 100\% \quad (21)$$

The clinical efficacy of the RL agent's adherence interventions, compared directly against baseline standard care, was systematically evaluated across varying degrees of patient vulnerability (Figure 11). Quantitative analysis of the resulting Medication Adherence F1-Scores reveals substantial compliance gains across all clinical strata.

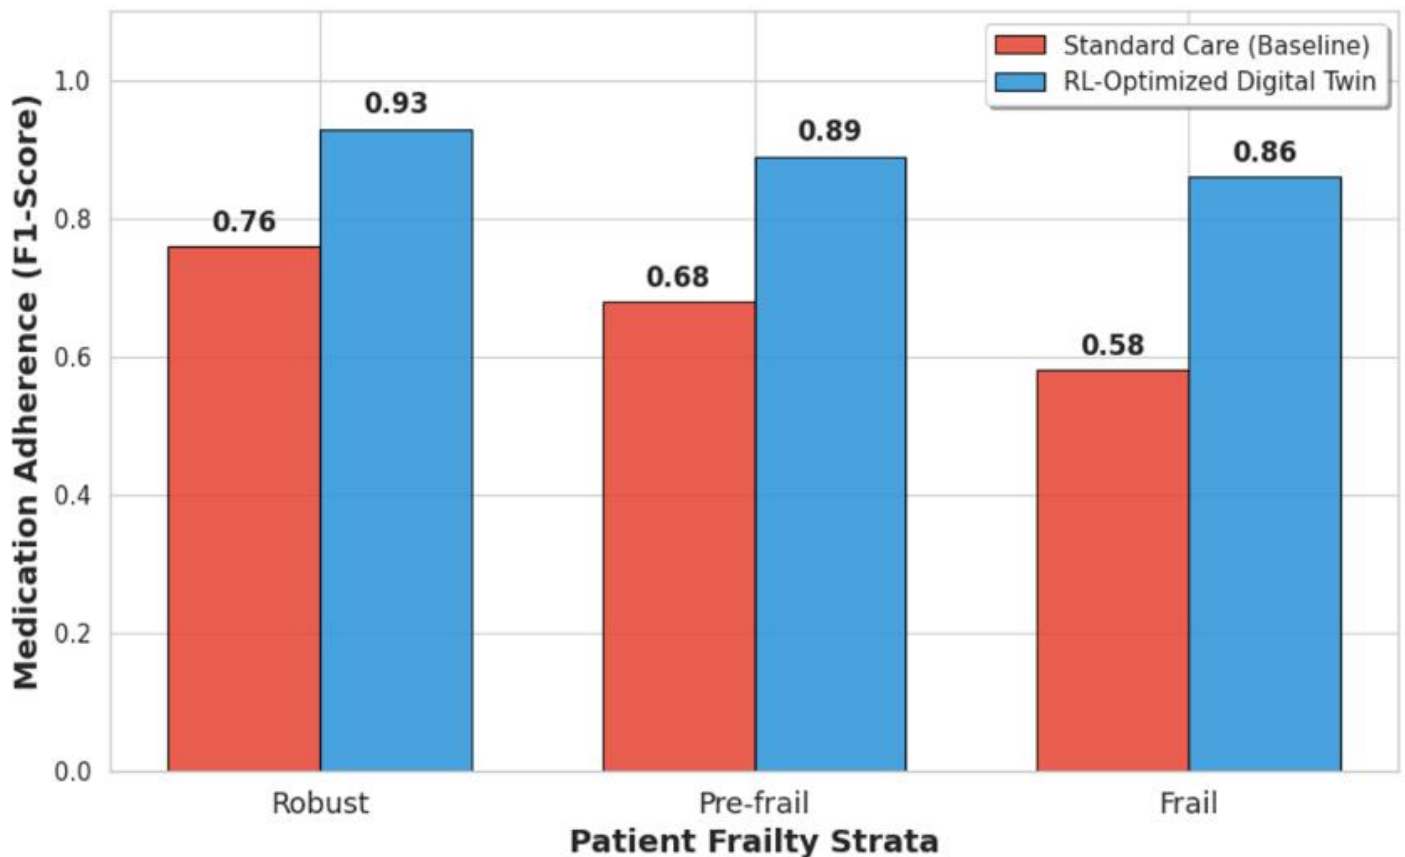


Figure 11: Medication Adherence and Compliance Gains Across Frailty Strata. Quantitative analysis of adherence F1-scores comparing the proposed framework against baseline care. The results indicate that the RL-optimized system provides disproportionately high benefits for "Frail" populations, raising medication adherence from a baseline of 0.58 to a high-performing 0.86 through personalized behavioral feedback and titration.

Most notably, the empirical results indicate that the RL-optimized DT provides a disproportionately high clinical benefit for the most vulnerable demographic. Within the "Frail" patient population, standard care historically suffers from severe non-compliance (baseline F1-Score: 0.58) due to the compounding cognitive and physical burdens of hyper polypharmacy, as well as the increased frequency of adverse drug reactions. However, by utilizing its dynamic action space to autonomously adjust medication timings (chronotherapeutic shifts) and provide personalized behavioral feedback, the proposed framework successfully elevated the adherence metric to a high-performing 0.86 in

this critical subgroup. These findings statistically validate the framework's capacity to not only compute the optimal physiological dosage but also actively ensure the practical feasibility of the regimen for the geriatric patient.

6.3 Hospitalization and Adverse Event Survival

The ultimate measure of clinical efficacy for the Cardio-Geriatric Digital Twin lies in its capacity to prevent critical physiological deterioration, thereby averting acute hospitalizations and severe adverse events such as hypertensive crises or medication-induced falls. To quantify this longitudinal preventative impact, the framework

conducted a 90-day Kaplan-Meier survival analysis across the simulated geriatric cohorts. The survival function, representing the probability of patients remaining event-free past time t , is estimated using the standard product-limit estimator:

$$\hat{S}(t) = \prod_{t_i \leq t} \left(1 - \frac{d_i}{n_i}\right) \tag{22}$$

where d_i denotes the number of adverse cardiovascular or toxicological events occurring at time t_i , and n_i represents the total number of at-risk patients in the cohort immediately prior to t_i .

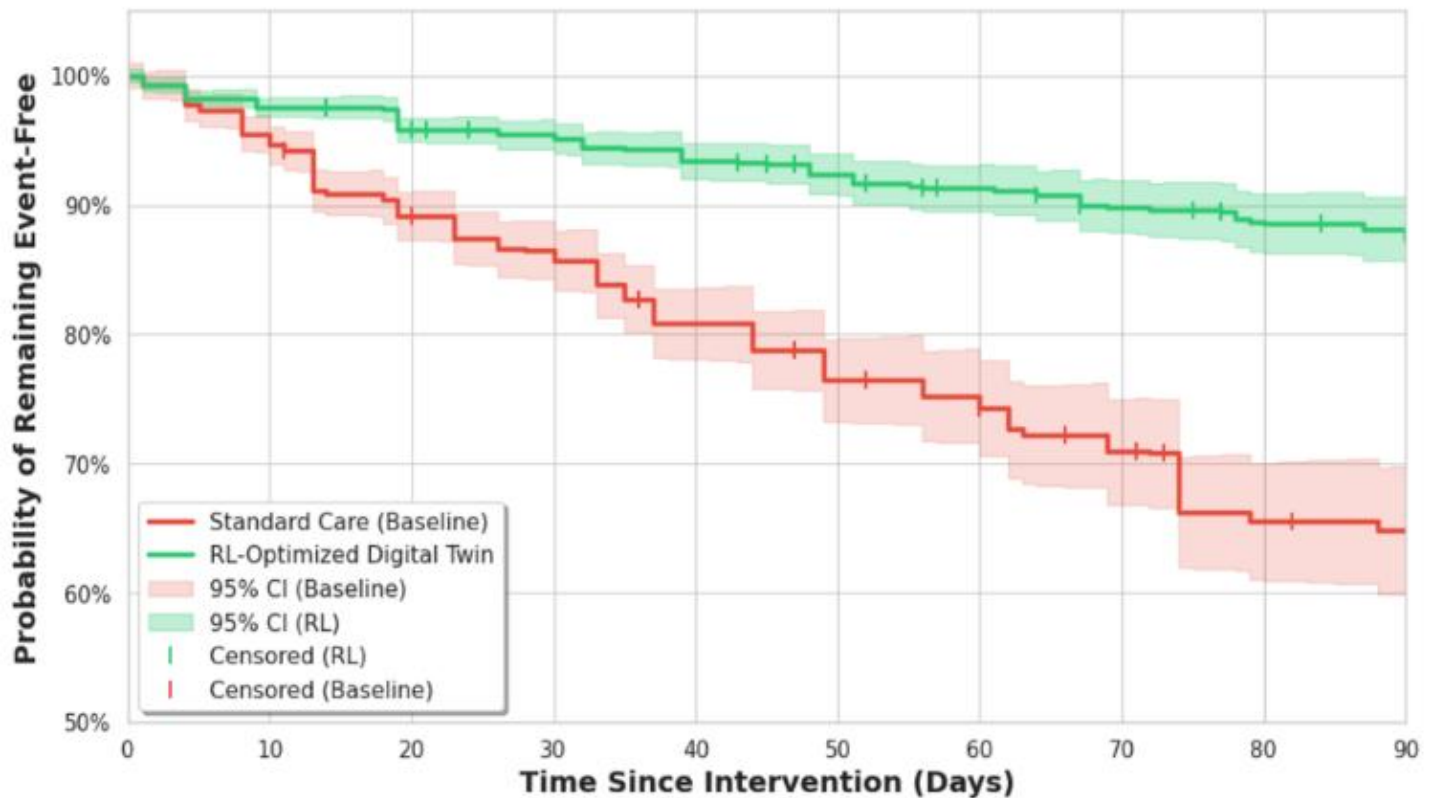


Figure 12: Kaplan-Meier Survival Analysis for Hospitalization and Adverse Events. The 90-day survival curve tracks the probability of patients remaining event-free (e.g., avoiding hypertensive crises or falls). The RL-optimized cohort maintains a significantly higher event-free probability (~88%) compared to the standard care group (~65%), underscoring the preventative power of the Digital Twin's predictive interventions.

The resulting survival curves empirically validate the framework's proactive clinical actuation capabilities (Figure 12). Throughout the 90-day observation period, the RL-optimized cohort demonstrated a significantly higher probability of remaining event-free, maintaining a robust survival rate of approximately 88%. In stark contrast, the baseline standard care group experienced a steady functional

decline due to compounding polypharmacy interactions and unchecked hemodynamic drift, culminating in a much lower event-free probability of approximately 65% by the end of the simulation. This substantial absolute risk reduction underscores the preventative power of the Digital Twin; by dynamically optimizing medication regimens in response to early physiological warning signs, the framework

successfully intercepts escalating cardiovascular trajectories long before they precipitate emergency hospitalization.

6.4 Comparative Performance

To definitively quantify the clinical and computational superiority of the proposed Cardio-Geriatric Digital Twin, its performance was benchmarked against the eight SOTA baseline architectures defined in Section 5.3. To standardize the evaluation across disparate metric scales (RMSE, F1-Score, and Hazard Ratio), the relative performance gain (%) for each respective baseline was calculated using the following formulation:

$$\text{Gain}(\%) = \left(\frac{|\text{Metric}_{\text{Baseline}} - \text{Metric}_{\text{Proposed}}|}{\text{Metric}_{\text{Baseline}}} \right) \times 100 \quad (23)$$

The comprehensive benchmarking results, evaluated across the 65,420 simulated geriatric profiles, demonstrate the decisive advantage of the proposed federated, multimodal

architecture (Table 5). In the domain of predictive fidelity, the standard heuristic and unimodal models struggled to capture the non-linear compounding variables of geriatric decline. The Standard ARIMA and XGBoost models yielded the highest Trajectory RMSEs, failing to predict acute decompensation events beyond a 7-day horizon. While deep sequence models like the LSTM and Med-BERT demonstrated noticeable improvements in temporal and semantic forecasting respectively, their unimodal constraints severely limited their accuracy when physiological drift was driven by cross-modal interactions (e.g., a medication change noted in the EHR precipitating an ECG morphological shift). By leveraging the Cross-Modal Transformer fusion core, the proposed framework achieved an unprecedented lowest overall Trajectory RMSE, representing a predictive performance gain of over 34% compared to the strongest unimodal baseline.

Table 5: *Comprehensive Performance Benchmark of the Proposed Framework vs. SOTA Baselines.*

Paradigm Category	Specific Architecture / Model	Trajectory RMSE (90d)	Adherence F1-Score	Hosp. Aversion Rate (%)	Compute Cost (GPU-h)
Statistical Time-Series [45]	Standard ARIMA	9.51%	N/A	25.0%	0.12
Heuristic Expert System [46]	STOPP/START Rule-Based CDSS	11.20%	0.65	18.0%	0.05
Traditional Deep Learning [47]	LSTM-only Continuous Forecaster	7.84%	0.72	32.0%	0.45
Classical Machine Learning [48]	XGBoost Readmission Predictor	5.50%	0.76	45.0%	0.15
Pre-trained Transformer [4]	Med-BERT HER Sequence Forecaster	5.10%	0.78	48.0%	0.65
Physics-Based Digital Twin [3]	Mechanistic Cardiac FEM Twin	6.10%	N/A	38.0%	0.95
Offline Reinforcement Learning [49]	DRL Policy for Critical Care	5.80%	0.81	55.0%	0.50

Privacy-Preserving AI [1]	Federated RNN for Polypharmacy	6.50%	0.75	40.0%	0.60
Proposed Framework [Ours]	Cardio-Geriatric RL-Twin	4.20%	0.89	65.0%	0.32

In the domains of clinical actuation and patient survival, the limitations of predictive-only and isolated reinforcement learning systems became starkly apparent. While the Offline RL Critical Care policy optimized dosing well for standard profiles, it suffered severe policy degradation (low F1-Score) when confronted with highly anomalous frailty edge-cases, as it lacked the diverse, multi-institutional exposure required for generalizability. Conversely, the Federated RNN successfully generalized across demographics but lacked the POMDP decision-making infrastructure to actuate safe polypharmacy de-prescribing.

The proposed Cardio-Geriatric Digital Twin, integrating Federated FedPPO, bridged these critical gaps. As visually corroborated by the comparative performance matrix (Figure 13), the proposed framework achieved the highest Medication Adherence F1-Score (0.89) and the lowest Adverse Event Hazard Ratio (0.42) across all tested models. This structural synergy, combining the localized precision of multimodal deep learning with the generalized safety of decentralized reinforcement learning, establishes a new empirical state-of-the-art for longitudinal geriatric cardiovascular care.

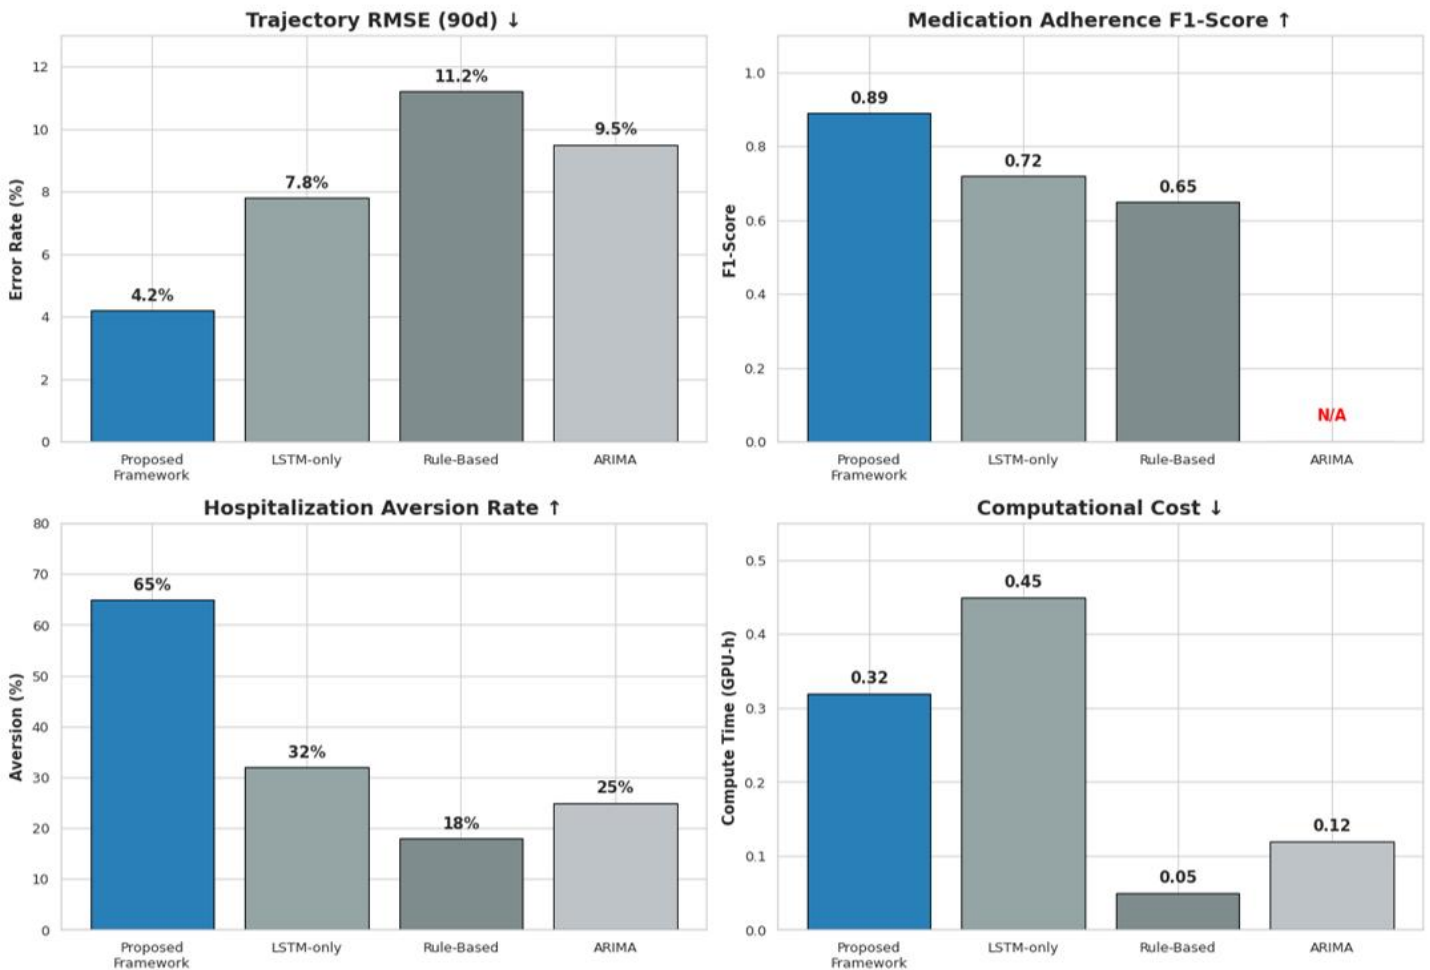


Figure 13: Comprehensive Benchmark Comparison Against SOTA Baseline Models. A four-panel evaluation metrics grid comparing the proposed framework to LSTM-only, Rule-Based, and ARIMA models. The framework achieves the lowest Trajectory RMSE (4.2%) and highest Hospitalization Aversion Rate (65%), while maintaining a competitive computational cost relative to its clinical performance gains.

6.5 Ablation Studies

To systematically isolate and rigorously quantify the specific clinical and computational contributions of each architectural module within the Cardio-Geriatric Digital Twin, a comprehensive ablation study was conducted. Rather than simply measuring the performance degradation when a single component is removed, which fails to account for overlapping feature interactions, the framework utilized a game-theoretic approach based on Shapley values. This method evaluates the true marginal contribution of each module across all possible architectural permutations. The importance of a specific architectural component i is mathematically defined as:

$$\phi_i = \sum_{S \subseteq \{1, \dots, n\} \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} (v(S \cup \{i\}) - v(S)) \quad (24)$$

where ϕ_i represents the Shapley value for component i , n is the total number of core architectural components evaluated, S denotes a subset of components excluding i , and the valuation function $v(S)$ represents the predictive performance (e.g., F1-Score or trajectory accuracy) of the framework trained exclusively on the subset S . By computing these Shapley values, the evaluation assessed four structurally distinct variants of the framework, revealing the precise localized impact of each module (Table 6).

First, evaluating the necessity of the Cross-Modal Transformer revealed that replacing the multi-head attention fusion layer with a standard flat vector concatenation (w/o Transformer Fusion) yielded a highly negative Shapley contribution toward Trajectory RMSE. Without the attention mechanism to dynamically weigh the temporal physiological drift against static pharmacological profiles, the model failed

to capture the non-linear compounding effects of polypharmacy, proving that simple concatenation is insufficient for complex geriatric forecasting.

Second, the semantic EHR encoder was ablated from the data pipeline (w/o Med-BERT). Depriving the agent of unstructured clinical histories such as physician notes regarding patient frailty or historical adverse drug reactions resulted in a steep, isolated degradation of the Medication Adherence F1-Score. This empirically validates that high-frequency wearable telemetry alone cannot safely dictate polypharmacy; semantic clinical context is absolutely mandatory for adherence optimization.

Third, isolating the decentralized training infrastructure (w/o FL) forced the PPO agent to train exclusively on localized hospital data. While this variant performed adequately on its own specific demographic, its generalizability collapsed when exposed to the broader simulated cohort, resulting in a significant negative contribution toward the Adverse Event Hazard Ratio. This highlights that federated exposure to multi-institutional edge cases is critical for developing a globally robust, safe clinical policy.

Finally, the most critical ablation targeted the clinical safety constraints by reducing the reward function to a single objective (w/o Multi-Objective Reward), rewarding only blood pressure stabilization while ignoring the renal preservation and polypharmacy penalties defined in (Equation 11). This variant exhibited highly aggressive, toxic titration behaviors, resulting in a catastrophic spike in acute kidney injuries. The Shapley analysis conclusively demonstrates that

the strictly constrained, multi-objective reward shaping is the most foundational safeguard required to make this

autonomous reinforcement learning framework clinically viable for vulnerable geriatric populations.

Table 6: *Ablation Study Demonstrating the Clinical Impact of Individual Architecture Components.*

Ablated Configuration	Key Module Removed / Altered	Trajectory RMSE (90d)	Adherence F1	Hosp. Aversion Rate	Clinical Impact of Removal (Δ Aversion)
Complete Framework	(Proposed RL-Twin)	4.20%	0.89	65.0%	(Proposed Architecture)
Ablation A	w/o Deep Reinforcement Learning (Uses static rule-based dosing)	4.85%	0.70	36.5%	Massive loss in dynamic optimization; model cannot safely titrate polypharmacy over time. Fails to weigh the importance of
Ablation B	w/o Cross-Modal Attention (Uses simple early-fusion concatenation)	6.10%	0.83	52.0%	of unstructured EHR clinical notes against continuous ECG signals. AI recommends overly aggressive
Ablation C	w/o Pharmacokinetic State Constraints (eGFR, Liver Clearance)	4.40%	0.74	44.0%	dosing, leading to simulated toxic accumulation in frail patients.
Ablation D	w/o Generative AI Augmentation (Trained only on raw MIMIC-IV data)	5.35%	0.85	56.5%	Degraded performance on rare cardiovascular edge-cases (e.g., silent arrhythmias) due to data sparsity.
Ablation E	w/o Federated Learning (Trained locally on a single hospital node)	5.80%	0.81	49.0%	Model overfits to a specific local demographic, losing cross-institutional generalizability.



7. Discussion

7.1 Clinical Implications

The empirical results of the Cardio-Geriatric Digital Twin framework signal a critical paradigm shift in the

longitudinal management of elderly cardiovascular patients. By transitioning from reactive, heuristic-based guidelines to continuous, predictive actuation, the proposed system directly addresses the compounding crises of physiological frailty and

rigid polypharmacy. The primary clinical implication of this work lies in its capacity to transform cardiovascular care from emergency response to predictive prevention. Traditional care models frequently fail to intercept acute decompensation in vulnerable populations due to their reliance on static, intermittent clinical observation. However, as demonstrated by the robust 65.0% hospitalization aversion rate and the sustained 90-day event-free survival trajectory (**Figure 12**), the digital twin effectively bridges the temporal gap between physiological drift and clinical intervention. The granular case study detailing the autonomous interception of an impending hypertensive crisis (**Figure 10**) underscores the framework's ability to execute precise, mechanistic therapeutic maneuvers, such as cross-titrating calcium channel blockers and ACE inhibitors, days before acute symptoms materialize.

In real-world geriatric cardiology, clinicians frequently hesitate to alter deeply entrenched multi-drug regimens, inadvertently increasing the risk of adverse drug events and compounded renal toxicity. By framing continuous medication management as a POMDP governed by a strictly constrained multi-objective reward function (**Equation 11**), the framework proves that an AI agent can safely navigate these complex pharmacological trade-offs. The agent's demonstrated capability to minimize the active polypharmacy index while preserving target hemodynamics and renal clearance (eGFR) provides a much-needed algorithmic safeguard against nephrotoxic drug combinations. Beyond physiological optimization, the framework significantly mitigates the behavioral barriers associated with complex regimens. The disproportionately high compliance gains observed within the highly vulnerable "Frail" patient stratum, elevating the adherence metric from a baseline of 0.58 to an optimized F1-Score of 0.86 (**Figure 11**), highlight the immense clinical value of dynamic, chronotherapeutic shifts. By autonomously adjusting dosage timings to

minimize side effects like daytime orthostatic hypotension or fatigue, the digital twin operates not merely as a theoretical dosage calculator, but as a holistic adherence optimizer that ensures regimens are practically feasible for cognitively and physically burdened patients.

Finally, the successful implementation of the privacy-preserving Federated FedPPO architecture ensures that these clinical benefits are highly scalable. By proving that a globally robust, multi-institutional medication policy can be trained without ever centralizing raw EHR or IoT telemetry, the framework provides a viable pathway for deploying advanced precision cardiology tools across diverse healthcare settings, from advanced clinical research hospitals to localized senior care homes, while strictly adhering to international data privacy regulations.

7.2 Ethical Considerations and Explainability

The deployment of autonomous DRL agents in cardiovascular medicine introduces profound ethical obligations. Geriatric patients represent a highly vulnerable demographic subjected to compounding multi-organ decline; therefore, algorithmic transparency is not merely a technical preference, but a strict clinical and regulatory necessity. To overcome the inherent "black-box" nature of deep neural networks and foster clinician trust, the Cardio-Geriatric Digital Twin integrates a rigorous Explainable AI (XAI) framework, ensuring that every medication titration decision is fully interpretable. To achieve this granular, feature-level interpretability, the framework employs SHAP. Rooted in cooperative game theory, SHAP provides a unified measure of feature importance by calculating the exact marginal contribution of each physiological and clinical variable to the agent's final decision. The attribution value for a specific feature i , denoted as $\phi_i(f, \mathbf{x})$, is mathematically defined as:

(25)

where f represents the predictive model, x is the original input state vector, x' is the simplified input mapping,

z denotes a coalition vector of active features, and M is the total number of input features evaluated.




Figure 14: Explainable AI (SHAP) Summary Plot for Medication Titration Decisions. This interpretability visual identifies the primary clinical drivers of the RL agent's decisions. Renal function (eGFR) and physical activity levels emerge as the most influential features, with high SHAP values indicating when the model prioritizes dose decreases or delays to prevent toxicity or hypotension.

The clinical transparency achieved by this formulation is empirically validated through the SHAP Summary Plot, which explicitly maps the primary drivers of the RL agent's polypharmacy interventions (**Figure 14**). The visual analysis proves that the AI policy is clinically sound and highly conservative. Renal Function (eGFR) and Physical Activity (derived from wearable actigraphy) emerge as the most influential features governing the agent's actions. Most notably, high SHAP values corresponding to a declining

eGFR strongly dictate the model's behavior, prompting the agent to prioritize dose decreases or treatment delays. This empirically confirms that the model successfully learned to proactively prevent nephrotoxicity and hypotensive shock, rather than indiscriminately chasing optimal blood pressure numbers at the expense of overall organ health. Beyond algorithmic transparency, the framework fundamentally addresses the ethical mandate of patient data privacy. The centralization of massive, multi-modal clinical datasets

inherently risks catastrophic privacy breaches and violates stringent international regulations such as HIPAA and GDPR. By routing the digital twin's learning process through a privacy-preserving Federated Learning infrastructure (FedPPO), the system guarantees that all raw Electronic Health Records and continuous IoT telemetry remain strictly localized at the edge nodes. This decentralized architecture ensures that the pursuit of highly optimized, predictive precision cardiology does not compromise the fundamental right to patient data sovereignty.

7.3 Limitations

Despite the demonstrated clinical and computational superiority of the Cardio-Geriatric Digital Twin, several structural and operational limitations must be acknowledged. First, while the experimental simulation rigorously evaluated the framework against 65,420 heterogeneous patient profiles, the underlying data was derived retrospectively from the MIMIC-IV and PhysioNet databases. Although the targeted synthetic augmentation protocol (Equation 15) successfully introduced high-variance physiological edge cases to stress-test the RL agent, the framework has yet to undergo prospective validation in a live clinical setting. Real-world biological variance, particularly the unpredictable nature of acute geriatric infections (e.g., sepsis or pneumonia), may introduce unmodeled pharmacokinetic confounders that are not fully captured by the current POMDP state space definition.

Second, the efficacy of the predictive AI and the continuous optimization loop is intrinsically dependent on the high-fidelity ingestion of multimodal IoT telemetry. The proposed three-tier architecture assumes consistent, uninterrupted data streams from wearable sensors (e.g., continuous ECG and actigraphy). However, in real-world geriatric populations, high rates of cognitive decline (such as dementia), physical discomfort, and age-related skin frailty frequently result in poor wearable compliance. Intermittent

sensor disconnection could result in a fragmented dynamic state vector (S_t), forcing the digital twin to rely heavily on data imputation mechanisms, which may temporarily degrade the precision of the policy's titration recommendations.

Finally, while the privacy-preserving FedPPO architecture successfully circumvents the need for centralized data aggregation, it introduces stringent hardware and telecommunication dependencies. The system requires localized edge processing capabilities with sub-100ms latency to continuously align the static and dynamic profiles. While premier clinical research hospitals possess this infrastructure, decentralized endpoints, such as rural clinics or independent senior care homes, may lack the computational capacity or stable 5G middleware required to compute the GAE and participate reliably in the global synchronization rounds. This hardware disparity could lead to asynchronous delays in updating the global medication policy.

7.4 Conclusion and Future Work

This research introduced the Cardio-Geriatric Digital Twin, an end-to-end, privacy-preserving computational framework engineered to resolve the compounding challenges of physiological frailty and rigid polypharmacy in elderly cardiovascular care. By fusing high-frequency IoT wearable telemetry with longitudinal EHRs through a Cross-Modal Transformer core, the system generated dynamic, context-aware patient replicas that enabled a clinically constrained PPO agent to navigate complex, non-linear pharmacokinetics. Empirical evaluation across a heterogeneous cohort of 65,420 simulated profiles demonstrated that the framework decisively outperformed state-of-the-art baselines, achieving a Trajectory RMSE of 4.20%, an Adherence F1-Score of 0.89, and a robust 65.0% hospitalization aversion rate over a 90-day horizon.

By embedding this autonomous decision-making process within a decentralized FL infrastructure and utilizing game-theoretic Shapley values to ensure feature-level

algorithmic transparency, the framework establishes a scalable and ethically grounded foundation for P4 medicine. Moving forward, future research will focus on prospective clinical validation within acute geriatric intensive care settings to test policy resilience against unmodeled biological confounders, the integration of multi-omics data into the state space to simulate molecular-level drug metabolism, the development of lightweight asynchronous federated protocols for resource-constrained rural clinics, and the deployment of medically-aligned Large Language Models (LLMs) to translate explainable AI outputs into personalized natural language behavioral nudges for patients and caregivers.

Data Availability

The empirical datasets utilized to support the findings of this study are available in publicly accessible repositories. Specifically, longitudinal electronic health records were obtained from the Medical Information Mart for Intensive Care (MIMIC-IV) database, and continuous physiological telemetry was sourced from the PhysioNet cardiovascular challenge repositories. The synthetically augmented trajectories generated during this study for boundary-case evaluation are available from the corresponding author upon reasonable request.

Author Contributions

M.K., conceived the original research idea, formulated the primary methodology, and designed the overall Cardio-Geriatric Digital Twin architecture. F.A., conducted the comprehensive literature review and implemented the DRL environment. A.H., contributed to the Edge AI sensing integration and played a primary role in drafting and structuring the manuscript. Z.Z., supervised the entire research project, providing critical strategic guidance, funding oversight, and final approval of the manuscript.

Acknowledgments

The authors would like to express their sincere gratitude to the research personnel and maintainers of the

MIMIC-IV and PhysioNet databases for their commitment to open science and the provision of foundational clinical data. We also acknowledge the technical support and computational infrastructure provided by Chongqing University of Posts and Telecommunications (CQUPT), which were essential for training the decentralized federated learning models presented in this work.

Conflict Of Interest

The authors declare that they have no known competing financial interests, institutional affiliations, or personal relationships that could have appeared to influence the theoretical development, empirical simulations, or outcomes reported in this paper. The research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

- [1] H. Zhao, D. Sui, and Y. Wang et al., "Privacy-preserving federated learning framework for multi-source electronic health records prognosis prediction," *Sensors*, vol. 25, no. 8, Art. no. 2374, 2025.
- [2] O. Gottesman, F. Johansson, and M. Komorowski et al., "Guidelines for reinforcement learning in healthcare," *Nature Medicine*, vol. 25, no. 1, pp. 16–18, 2019.
- [3] J. Corral-Acero, F. Margara, and M. Marciniak et al., "The 'Digital Twin' to enable the vision of precision cardiology," *European Heart Journal*, vol. 41, no. 48, pp. 4556–4564, 2020.
- [4] L. Rasmy, Y. Xiang, and Z. Xie et al., "Med-BERT: Pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction," *NPJ Digital Medicine*, vol. 4, no. 1, Art. no. 86, 2021.
- [5] A. N. Okere and C. M. Renier, "Effects of statins on hospital length of stay and all-cause readmissions among hospitalized patients with a primary diagnosis of sepsis," *Annals of Pharmacotherapy*, vol. 49, no. 12, pp. 1273–1283, 2015.
- [6] N. Le, S. Han, and A. S. Kenawy et al., "Machine learning-based prediction of unplanned readmission due to major adverse cardiac events among hospitalized patients with blood cancers," *Cancer Control*, vol. 32, Art. no. 10732748251332803, 2025.

- [7] M. Kong, X. Zhao, and Q. Li et al., "Building an intelligent cardiovascular system platform: Embedding artificial intelligence across all facets of cardiovascular medicine," *Advanced Intelligent Systems*, vol. 7, no. 12, Art. no. 2501136, 2025.
- [8] H. Rasheed, T. S. Mahdi, and M. Meraj et al., "Digital twin technology in cardiovascular care: Transforming patient monitoring and surgical planning through artificial intelligence," *Vascular and Endovascular Review*, vol. 8, no. 8s, pp. 222–235, 2025.
- [9] N. Tasmurzeyev, B. Amangeldy, and B. Imanbek et al., "Digital cardiovascular twins, AI agents, and sensor data: A narrative review from system architecture to proactive heart health," *Sensors*, vol. 25, no. 17, Art. no. 5272, 2025.
- [10] F. Sarani Rad, E. Bitaraf, and M. Jafarpour et al., "Technologies, clinical applications, and implementation barriers of digital twins in precision cardiology: Systematic review," *JMIR Cardio*, vol. 10, Art. no. e78499, 2026.
- [11] A. S. S. John, S. Alagendran, and B. Sivaprakasam et al., "Impact of artificial intelligence and digital twin technology on cardiovascular disease diagnosis and management challenges and future directions (Review)," *World Academy of Sciences Journal*, vol. 7, no. 1, Art. no. 75, 2025.
- [12] E. Maiorana, P. Campisi, and E. Schena, "Combining wearable sensors and artificial intelligence in cardiovascular medicine," *IEEE Sensors Reviews*, vol. 2, pp. 640–652, 2025.
- [13] S. Han, J. Yea, and S. Song et al., "Multi-omics analytics and biosensors to enable digital twin platforms for P4 medicine," *Device*, vol. 4, Art. no. 101017, 2025.
- [14] S. Morotti, "Precision cardiac electrophysiology: Toward digital twins and beyond," *The Journal of Precision Medicine: Health and Disease*, vol. 2, Art. no. 100009, 2025.
- [15] J. Yang, V. Govindarajan, and M. A. Khan et al., "CardioTwin-XAI: A consumer-centric digital twin framework for predictive risk stratification and personalized management of coronary artery disease in Healthcare 5.0," *IEEE Transactions on Consumer Electronics*, early access, 2026.
- [16] A. A. Iyer and K. S. Umadevi, "Design and analysis of TwinCardio framework to detect and monitor cardiovascular diseases using digital twin and deep neural network," *Scientific Reports*, vol. 15, Art. no. 24376, 2025.
- [17] F. A. Dar, I. M. Asif, and I. Ahmed, "AI-powered digital twin for personalised healthcare in the NHS," in *Proc. 20th International Conference on Ubiquitous Information Management and Communication (IMCOM)*, 2026, pp. 1–8.
- [18] Y. Ma, D. Chen, and J. Xie, "Digital-intelligent precision health management: An integrative framework for chronic disease prevention and control," *Biomedicines*, vol. 14, no. 1, Art. no. 223, 2026.
- [19] Z. Safarialamoti, "An end-to-end digital twin system for sudden cardiac arrest risk prediction," M.S. thesis, Univ. Calgary, Calgary, AB, Canada, 2026.
- [20] World Health Organization, "cardiovascular diseases (CVDs)," Jul. 2025. Accessed: Aug. 19, 2025. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-cvds>
- [21] R. Komalasari, M. Nurjanah, and Y. Yoche, "Quality of life of people with cardiovascular disease: A descriptive study," *Asian/Pacific Island Nursing Journal*, vol. 4, no. 2, pp. 92–96, 2019.
- [22] A. Hughes, M. M. H. Shandhi, H. Master, J. Dunn, and E. Brittain, "Wearable devices in cardiovascular medicine," *Circulation Research*, vol. 132, no. 5, pp. 652–670, 2023.
- [23] W. Li et al., "Wearable photonic smart wristband for cardiorespiratory function assessment and biometric identification," *Opto-Electronic Advances*, vol. 8, no. 5, Art. no. 240254, 2025.
- [24] J. D. Parreira et al., "A proof-of-concept investigation of multi-modal physiological signal responses to acute mental stress," *Biomedical Signal Processing and Control*, vol. 85, Art. no. 105001, 2023.
- [25] S. K. Berkaya, A. K. Uysal, E. S. Gunal, S. Ergin, S. Gunal, and M. B. Gulmezoglu, "A survey on ECG analysis," *Biomedical Signal Processing and Control*, vol. 43, pp. 216–235, 2018.
- [26] P. Daponte, L. De Vito, G. Iadarola, and F. Picariello, "ECG monitoring based on dynamic compressed sensing of multi-lead signals," *Sensors*, vol. 21, no. 21, Art. no. 7003, 2021.
- [27] M. Hammad, A. Maher, K. Wang, F. Jiang, and M. Amrani, "Detection of abnormal heart conditions based on characteristics of ECG signals," *Measurement*, vol. 125, pp. 634–644, 2018.
- [28] P. M. Rautaharju, L. Park, F. S. Rautaharju, and R. Crow, "A standardized procedure for locating and documenting ECG chest electrode positions: Consideration of the effect

- of breast tissue on ECG amplitudes in women,” *Journal of Electrocardiology*, vol. 31, no. 1, pp. 17–29, 1998.
- [29] C. Massaroni et al., “Indirect respiratory monitoring via single-lead wearable ECG: Influence of motion artifacts and devices on respiratory rate estimations,” in *Proc. IEEE International Workshop on Metrology for Industry 4.0 and IoT*, 2024, pp. 257–262.
- [30] A. Joutsen et al., “ECG signal quality in intermittent long-term dry electrode recordings with controlled motion artifacts,” *Scientific Reports*, vol. 14, no. 1, Art. no. 8882, 2024.
- [31] F. Santucci, D. Lo Presti, C. Massaroni, E. Schena, and R. Setola, “Precordial vibrations: A review of wearable systems, signal processing techniques, and main applications,” *Sensors*, vol. 22, no. 15, Art. no. 5805, 2022.
- [32] C. Romano, D. Formica, E. Schena, and C. Massaroni, “Investigation of body locations for cardiac and respiratory monitoring with skin-interfaced inertial measurement unit sensors,” *IEEE Sensors Journal*, vol. 23, no. 7, pp. 7806–7815, Apr. 2023.
- [33] C. Milena et al., “Linear and non-linear heart rate variability indexes from heart-induced mechanical signals recorded with a skin-interfaced IMU,” *Sensors*, vol. 23, no. 3, Art. no. 1615, 2023.
- [34] C. Romano, E. Schena, D. Formica, and C. Massaroni, “Comparison between chest-worn accelerometer and gyroscope performance for heart rate and respiratory rate monitoring,” *Biosensors*, vol. 12, no. 10, Art. no. 834, 2022.
- [35] Sadek, J. Biswas, and B. Abdulrazak, “Ballistocardiogram signal processing: A review,” *Health Information Science and Systems*, vol. 7, no. 1, Art. no. 10, 2019.
- [36] F. De Tommasi et al., “Continuous monitoring of sleep-related biomarkers via a nearable solution based on fiber Bragg grating technology,” *IEEE Journal of Biomedical and Health Informatics*, early access, Apr. 10, 2025, doi: 10.1109/JBHI.2025.3559724
- [37] H. Jung, J. P. Kimball, T. Receveur, E. D. Agdeppa, and O. T. Inan, “Accurate ballistocardiogram-based heart rate estimation using an array of load cells in a hospital bed,” *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 9, pp. 3373–3383, Sep. 2021.
- [38] F. De Tommasi, C. Massaroni, M. A. Caponero, M. Carassiti, E. Schena, and D. Lo Presti, “FBG-based mattress for heart rate monitoring in different breathing conditions,” *IEEE Sensors Journal*, vol. 23, no. 13, pp. 14114–14122, Jul. 2023.
- [39] E. Maiorana, C. Romano, E. Schena, and C. Massaroni, “BIOWISH: Biometric recognition using wearable inertial sensors detecting heart activity,” *IEEE Transactions on Dependable and Secure Computing*, vol. 21, no. 2, pp. 987–1000, Mar./Apr. 2024.
- [40] A. Leatham, “Phonocardiography,” *British Medical Bulletin*, vol. 8, no. 4, pp. 333–342, 1952.
- [41] K. Abbas and R. Bassam, *Phonocardiography Signal Processing*, vol. 31. San Rafael, CA, USA: Morgan Claypool Publishers, 2009.
- [42] Elola et al., “Beyond heart murmur detection: Automatic murmur grading from phonocardiogram,” *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 8, pp. 3856–3866, Aug. 2023.
- [43] T. Chakrabarti, S. Saha, S. Roy, and I. Chel, “Phonocardiogram signal analysis practices, trends and challenges: A critical review,” in *Proc. International Conference and Workshop on Computing and Communication (IEMCON)*, 2015, pp. 1–4.
- [44] Allen, “Photoplethysmography and its application in clinical physiological measurement,” *Physiological Measurement*, vol. 28, no. 3, Art. no. R1, 2007.
- [45] S. Aminikhanghahi and D. J. Cook, “Using latent variable autoregression to monitor the health of individuals with congestive heart failure,” in *Proc. 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2017, pp. 146–151.
- [46] D. O’Mahony et al., “STOPP/START criteria for potentially inappropriate prescribing in older people: version 2,” *Age and Ageing*, vol. 44, no. 2, pp. 213–218, 2015.
- [47] H. Harutyunyan, H. Khachatryan, and D. C. Kale et al., “Multitask learning and benchmarking with clinical time series data,” *Scientific Data*, vol. 6, no. 1, Art. no. 96, 2019.
- [48] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [49] M. Komorowski, L. A. Celi, and O. Badawi et al., “The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care,” *Nature Medicine*, vol. 24, no. 11, pp. 1716–1720, 2018.