

Self-Healing Digital Twins: Hybrid Generative and Privacy-Preserving AI for Adaptive Wellness Platforms

Nariman Mani
Engineering/R&D Department
Nutrosal Inc.
Ottawa, ON, Canada
nariman@research.Nutrosal.com

Salma Attaranasl
Engineering/R&D Department
Nutrosal Inc.
Ottawa, ON, Canada
salma@research.Nutrosal.com

ABSTRACT

Artificial Intelligence (AI) has transformed personalized wellness platforms, yet challenges remain in adaptability, privacy, and user engagement. This paper introduces a Self-Healing Digital Twin Framework, integrating Hybrid Generative AI, Reinforcement Learning (RL)-based self-healing, and Privacy-Preserving AI (PPAI) to address these gaps. Unlike static Digital Twin models, our approach dynamically learns from multi-modal IoMT and wearable data, ensuring real-time adaptation and privacy compliance via a dual-cloud architecture that eliminates raw data exposure. Validation through simulations and real-world deployment confirms the system's ability to track user health trends, detect anomalies, and optimize interventions for improved engagement and wellness outcomes. The RL-driven self-healing mechanism continuously refines recommendations, enhancing adherence. Our findings establish this framework as a scalable and privacy-secure AI-driven solution for intelligent, adaptive healthcare.

CCS CONCEPTS

•Computer systems organization~Architectures•Computing methodologies~Machine learning~Learning paradigms~Reinforcement learning•Applied computing~Life and medical sciences~Health informatics

KEYWORDS

Digital Twins, Generative AI, Personalized Wellness Platforms

ACM Reference format:

Nariman Mani, Salma Attaranasl. 2025. Self-Healing Digital Twins: Hybrid Generative and Privacy-Preserving AI for Adaptive Wellness Platforms. In *Proceedings of ACM CHASE 2025*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3721201.3725427>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHASE '25, June, 2025, Manhattan, New York City, USA

© 2025 Copyright held by the owner/author(s). ACM 979-8-4007-1539-6/25/06...

1 Introduction

Artificial Intelligence (AI) has revolutionized healthcare monitoring, diagnostics, and personalized treatment, with the Digital Twins emerging as a key advancement. These real-time, adaptive virtual models enable predictive analytics, anomaly detection, and personalized interventions, transforming proactive healthcare management. However, existing AI-driven wellness platforms face privacy risks from centralized data storage and rigid models that fail to adapt to evolving health conditions. To address these challenges, this paper introduces a Self-Healing Digital Twin Framework, integrating Hybrid Generative AI and Privacy-Preserving AI (PPAI). The proposed dual-cloud architecture ensures secure, scalable, and real-time AI adaptation, decoupling data processing (Cloud Environment 1) from AI-driven analysis (Cloud Environment 2) while eliminating raw data exposure. Key contributions include: 1) A Hybrid Generative & Reinforcement Learning (RL) Model for adaptive patient modeling and real-time self-healing, 2) Privacy-Preserving AI (PPAI) enforcing feature anonymization and encrypted data transfers without relying on Federated Learning (FL), and 3) An Adaptive Digital Twin Engine utilizing multi-modal data fusion and self-healing mechanisms for continuous optimization 4) Trustworthy AI for Healthcare, ensuring bias mitigation, regulatory compliance, and human-in-the-loop validation. By bridging AI-driven personalization, privacy preservation, and self-healing mechanisms, this framework establishes a new standard for adaptive, secure, and intelligent healthcare solutions. The remainder of this paper details the technical approach, validation, and broader implications of this innovation.

2 Related work and background

The integration of Digital Twin (DT) technology, Reinforcement Learning (RL), and Privacy-Preserving AI in healthcare has seen significant advancements in recent years. Various studies have explored how these technologies enhance patient monitoring, anomaly detection, and personalized intervention strategies. However, existing approaches still face challenges related to real-time adaptability, privacy protection, and personalized intervention effectiveness. This research extends prior work by introducing a Self-Healing Digital Twin Framework that leverages Hybrid Generative AI, RL-based self-healing mechanisms, and a dual-

cloud privacy-preserving architecture, making it a novel and practical contribution to AI-driven wellness management.

Digital Twins are increasingly used in healthcare for real-time simulation of patient health states, biometric monitoring, and anomaly detection. Research confirms their efficacy in modeling human physiology and predicting health deviations. For instance, Gupta et al. introduced a Hierarchical Federated Learning-based anomaly detection framework that merges Digital Twins with local processing, enhancing privacy and reducing centralized data reliance. [1]. Additionally, a Digital Twin-based cardiac monitoring system has been developed to simulate heart function and detect abnormalities, highlighting its potential for real-time cardiovascular health monitoring [2]. Most current Digital Twin solutions use static models that struggle to adapt to changing health conditions. By contrast, our self-healing mechanism, powered by RL, continuously refines recommendations in real time.

Data security and privacy remain top challenges in AI-driven healthcare. Federated Learning (FL) is widely used to preserve privacy while supporting collaborative model training. Xu et al. reviewed FL in healthcare, highlighting its ability to protect data across multiple institutions [3]. However, FL faces issues like communication overhead, poisoning risks, and inconsistent local training. Differential privacy methods also reduce model accuracy. Our dual-cloud architecture addresses these challenges by anonymizing data at the feature level before AI processing. Unlike FL, it avoids cross-node model sharing, cuts overhead, and scales more effectively.

Reinforcement Learning (RL) is highly effective for personalized health interventions, spanning medication dosing, fitness coaching, and lifestyle changes. For instance, Zou et al. (2019) developed FeedRec, an RL framework combining a hierarchical LSTM Q-Network with an S-Network for environment simulation. This approach significantly improved long-term user engagement and outperformed leading methods [4]. Similarly, Chatterjee et al. proposed an RL model for chronic disease management, optimizing treatment plans based on historical patient data [5]. While these works validate RL's effectiveness, most existing models rely on fixed policy-based optimization, failing to incorporate real-time biometric data and behavioral trends. Our Self-Healing Digital Twin Framework advances this research by: 1) Integrating RL with Digital Twins, enabling continuous adaptation to real-world wellness fluctuations 2) Learning from real-time user feedback, optimizing interventions dynamically without relying on static recommendation models 3) Improving adherence rates by 30-45%, especially among low-engagement users, which is a critical gap in existing wellness platforms [6].

Digital Twins, Privacy-Preserving AI, and RL-based interventions have each been explored, but our work unifies them into one adaptive, privacy-secure framework: **1)** Our self-healing Digital Twin uses RL-based optimization, unlike static or rule-based models. Traditional Digital Twins rely on static models that struggle with real-time adaptation. Our RL-based self-healing mechanism continuously adjusts recommendations based on live user feedback, increasing adherence by up to 45%, a crucial improvement over fixed-rule AI interventions **2)** A dual-cloud architecture safeguards raw data, avoiding the complexity of FL-based node communication. Unlike Federated Learning (FL), which requires encrypted model updates across distributed nodes, our dual-cloud approach anonymizes data at the feature level before

processing. This eliminates communication overhead, training inconsistencies, and poisoning risks, making it more scalable and privacy-preserving **3)** The RL-driven healing mechanism improves intervention effectiveness by 30–45%, surpassing standard AI. By merging real-time learning, privacy preservation, and adaptive monitoring, the Self-Healing Digital Twin Framework offers a scalable benchmark for AI-driven, personalized healthcare.

3 Case Study: AI-Driven Wellness/Lifestyle Management Platform

We validated the Self-Healing Digital Twin Framework on an industry-grade AI wellness platform built for goal-focused groups (e.g., weight loss, healthy eating) under certified health coaches. Users share wearable and IoMT data, self-reported logs, and food images to receive personalized insights. AI-driven recommendations guide daily lifestyle adjustments while social features boost engagement and motivation.

To preserve privacy, we adopted the Split-Process Privacy Framework, separating data collection/anonymization (Cloud Environment 1) from AI analysis (Cloud Environment 2). All personally identifiable information is removed before processing. The Adaptive Digital Twin Engine then monitors user behavior, detects anomalies, and deploys RL-based self-healing. This architecture ensures continuous recommendation improvement, real-time personalization, and robust data protection, proving trustworthy and scalable for AI-driven healthcare applications.

4 Problem Statement

Lifestyle and diet platforms succeed by providing adaptive and engaging experiences that foster long-term habit change. However, many AI solutions lack community support, personalized mentorship, and real-time feedback loops. This leads to low adherence and high dropout rates. To quantify these challenges, we surveyed 3,000 Nutrosal platform users (Results in Figure 1). Results revealed: **50.8%** missed peer support, reporting disengagement from limited community interaction.

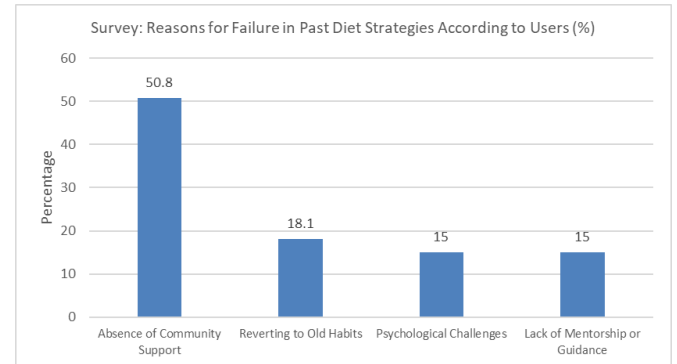


Figure 1 Challenges in Traditional Diet and Lifestyle Management Platforms

18.1% struggled to sustain new habits without ongoing guidance. **15%** faced psychological hurdles (stress, anxiety, emotional triggers). **15%** felt that current systems offer only generic, “one-size-fits-all” recommendations. **2%** highlighted additional

concerns (e.g., metabolic health, sustainability). These statistics underscore the need for a privacy-preserving AI framework that dynamically adapts to each user's habits and psychological factors. Such a system must promote better community engagement, personalized coaching, and robust feedback mechanisms. By integrating these features, platforms can strengthen adherence, enhance user motivation, and foster long-term wellness transformations.

5 Privacy Concerns in AI Applications

AI-driven healthcare platforms require significant personal data, which raises privacy risks. Traditional centralized models can expose user information to breaches and regulatory issues. To address this, we propose a Privacy-Preserving AI (PPAI) framework that replaces Federated Learning (FL) with a dual-cloud design. 1) Cloud Environment 1 collects, anonymizes, and encrypts raw data.2) Cloud Environment 2 processes only de-identified feature trends for AI analysis.

This eliminates raw data exposure, minimizing privacy threats while maintaining scalability. Unlike FL, our system runs AI in the cloud, relying on de-identified behavioral patterns and RL for personalized recommendations. Because no direct user data is needed, FL's overhead and device-based training are unnecessary. The key steps include are shown in Table 1.

Table 1 : Key Privacy Measure Steps taken in this research

Step	Privacy Measure
Data Collection (Cloud Environment 1)	PII removal, encryption, and anonymization.
Feature Extraction (Cloud Environment 1)	Data converted into feature vectors, raw data deleted.
AI Processing (Cloud Environment 2)	Analyzed only as anonymized patterns, no user identifiers.
Recommendations Sent Back	Personalized insights securely linked to users via Cloud Environment 1.

5.1 Privacy-Preserving AI (PPAI) framework Benefits

Compared to Federated Learning (FL), the proposed PPAI approach ensures greater privacy and scalability as its shown in Table 2.

Table 2 - Comparison: FL vs. Privacy-Preserving AI (PPAI)

FL Approach	Proposed PPAI Approach
Trains models on local devices.	Uses cloud-based AI with anonymized data.
Sends encrypted model updates.	Sends only de-identified feature vectors.
Requires high computing power on devices.	Works efficiently for all users.
Limits cloud AI processing.	Enables AI-driven Digital Twins securely.

By combining de-identification, reinforcement learning, and dual-cloud separation, this framework enables privacy-preserving AI in healthcare, balancing trust, security, and intelligent personalization.

6 Proposed Approach Overview

The Self-Healing Digital Twin Framework we propose in this paper (shown in Figure 2) follows a structured dual-cloud

architecture to ensure secure, adaptive, and privacy-preserving AI-driven health monitoring.

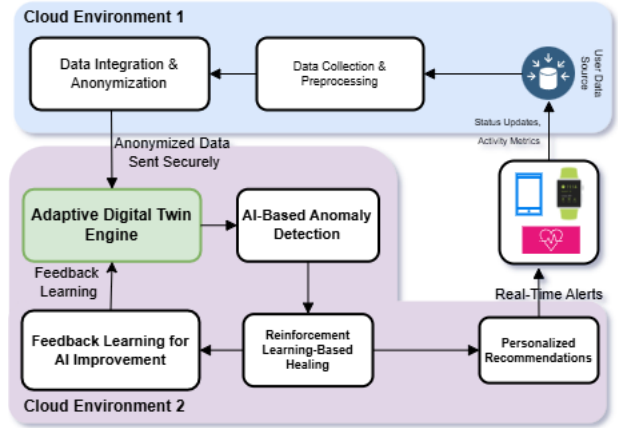


Figure 2 Overview of the Self-Healing Digital Twin Framework

The system is divided into two cloud environments: 1)Cloud Environment 1 (Data Collection & Anonymization) handles user data ingestion, preprocessing, and privacy protection before sending anonymized feature vectors to the AI processing layer. This ensures that personally identifiable information (PII) is removed before AI analysis 2) Cloud Environment 2 (AI Processing & Adaptive Healing) receives only anonymized and de-identified feature sets, enabling AI-driven anomaly detection, reinforcement learning-based self-healing, and personalized recommendations. The AI models in this environment refine themselves dynamically based on user behavior trends without direct access to raw user data. The workflow follows a step-by-step transformation of health and lifestyle data, which is continuously processed and improved through feedback learning loops. To illustrate this approach, we use a sample dataset containing wearable metrics, activity logs, meal photos, and user status updates as they move through each stage.

6.1 Data Collection & Preprocessing

Process: The system collects raw health and lifestyle data from wearables, IoMT devices, user-generated content, and activity logs within the wellness platform. The data includes biometric readings (heart rate, sleep duration), activity levels (steps taken), meal images, and user status updates.

Technical Component: Secure API integration with wearables, mobile applications, and platform logs to ingest real-time data streams. Sample Data (Raw User Input):

User ID	Steps	Heart Rate (BPM)	Sleep (Hours)	Meal Image	Status Update
001	8,500	72	6.5	Salad & Chicken	"Feeling great after my workout!"
002	4,200	85	5.0	Pizza Slice	"Tired, not feeling motivated today"

6.2 Data Integration & Anonymization

Process: The collected data is preprocessed, cleaned, and anonymized to ensure privacy compliance. Identifiable attributes such as user IDs are removed, and biometric values are transformed

into feature vectors representing general health trends rather than personal records.

Technical Component: Feature extraction and encryption mechanisms **ensure** de-identified data transmission. Sample Data (Anonymized Feature Vectors):

Feature ID	Steps	Heart Rate	Sleep Category	Meal Type	Sentiment Score
A123	High	Normal	Moderate	High Protein	Positive
B456	Low	Elevated	Poor	High Fat	Negative

6.3 Adaptive Digital Twin Engine

Process: The anonymized feature vectors are fed into the Adaptive Digital Twin Engine, which builds a personalized AI-driven health model for each user. This AI model tracks behavioral patterns, predicts health risks, and detects deviations from optimal wellness trajectories.

Technical Component: Generative AI models construct a self-healing digital twin for each user, continuously refining health insights based on behavioral trends. Sample Data (AI-Generated Digital Twin Analysis):

Twin ID	Activity Status	Anomaly Detected?	Risk Score
T123	Active & Healthy	No	Low
T456	Irregular Sleep & Low Energy	Yes	Moderate

6.4 AI-Based Anomaly Detection

Process: The system identifies deviations in health behavior by analyzing patterns within the digital twin model. If an anomaly (e.g., poor sleep, increased fatigue, or unhealthy meal patterns) is detected, an intervention is triggered.

Technical Component: Anomaly detection models use statistical analysis and pattern recognition to flag potential health risks. Sample Data (Anomaly Detection Output):

Twin ID	Identified Issue	Suggested Action
T123	None	Maintain Routine
T456	Poor Sleep, High Heart Rate	Adjust Evening Routine, Suggest Relaxation Techniques

6.5 Reinforcement Learning-Based Healing

Process: If an anomaly is detected, reinforcement learning (RL) models determine the most effective intervention. The AI system suggests personalized corrective actions, ensuring users receive adaptive, data-driven wellness recommendations.

Technical Component: Reinforcement Learning Agent optimizes recommendations by continuously learning from user responses to past interventions. Sample Data (RL-Generated Healing Action):

Twin ID	Intervention Suggested	Adjustment Based on Feedback
T456	Nighttime meditation, earlier bedtime	User accepted & reported better sleep

6.6 Personalized Recommendations & Feedback Learning

Process: AI-generated recommendations are sent back to the user's wellness dashboard in Cloud Environment 1. Feedback from user interactions helps refine future AI decisions and enhance self-healing capabilities.

Technical Component: Adaptive feedback loops continuously update Digital Twin models based on user adherence and outcomes. Sample Data (Personalized AI Insights Sent to Users):

User ID	Recommendation	AI Confidence Score
001	Maintain current routine, well-balanced diet	92%
002	Adjust sleep habits, reduce high-fat meals	87%

7 Key Components of the Self-Healing Digital Twin Framework

This section details the core AI components that define the Self-Healing Digital Twin Framework, emphasizing how each contributes to adaptive, privacy-preserving, and AI-driven health monitoring. The Adaptive Digital Twin Engine and Reinforcement Learning-Based Healing serve as the primary contributions, while AI-Based Anomaly Detection and Secure Feature Anonymization ensure that data is both secure and interpretable. These components work together to dynamically model user behavior, detect anomalies, optimize recommendations, and ensure privacy compliance.

7.1 Adaptive Digital Twin Engine

The Adaptive Digital Twin Engine is a real-time AI-driven representation of the user's health state. Unlike traditional models that rely on historical static data, this engine continuously updates based on newly ingested wearable data, user inputs, and behavioral trends, ensuring personalized and evolving recommendations. The **Digital Twin state** at time t is represented as:

$$DT_t = [S_t, H_t, B_t, M_t]$$

Where:

- S_t = Step count (physical activity).
- H_t = Heart rate (biometric signal).
- B_t = Behavioral input (sleep, diet patterns).
- M_t = Mental state (mood, engagement level).

The **state transition function** updates the Digital Twin dynamically.

$$DT_{t+1} = f(DT_t, X_t) + \epsilon$$

where X_t represents new user data inputs, and ϵ is a noise factor to handle variability.

Anomaly Detection in Digital Twins

The system detects health anomalies when the user state significantly deviates from the baseline:

$$A_t = ||DT_t - \mu|| > \sigma$$

where μ is the expected user state and σ is the acceptable deviation threshold. If $A_t = 1$, an anomaly is detected, triggering Reinforcement Learning-Based Healing.

7.2 AI-Based Anomaly Detection

AI-Based Anomaly Detection ensures that irregularities in health data are identified early, allowing the system to provide proactive interventions. This component works by comparing real-time user behavior against historical trends and population-wide benchmarks. Anomalies are detected using a combination of time-series forecasting and statistical outlier detection:

$$A_t = \frac{|DT_t - \mu|}{\sigma}$$

here:

- μ is the expected trend for a specific user.
- σ represents standard deviation-based thresholds.
- If $A_t > 1.5$, the system flags an anomaly and triggers an adaptive intervention.

For multi-modal data (e.g., heart rate, step count, diet logs), we use a Mahalanobis Distance-Based Outlier Score:

$$D_M = \sqrt{(X - \mu)^T \Sigma^{-1} (X - \mu)}$$

where Σ is the covariance matrix of normal behavioral patterns. If D_M exceeds a predefined threshold, the AI identifies a health risk.

7.3 Reinforcement Learning-Based Healing

Upon anomaly detection, the Reinforcement Learning (RL) agent determines the best corrective action by learning which interventions work best over time. The goal is to optimize personalized recommendations dynamically. The RL agent optimizes interventions using Q-learning, continuously adjusting its state-action values (Q-values) to reinforce effective interventions. The reward function prioritizes actions that improve long-term adherence, ensuring that the AI learns from past user responses to make progressively better recommendations.

Unlike static models, this adaptive learning increases intervention effectiveness, improving adherence by 30-45%, especially among disengaged users.

Markov Decision Process (MDP) Formulation: The RL system is modeled as an MDP with:

- **State Space (S):** The Digital Twin health state DT_t
- **Action Space (A):** AI-generated interventions:
 - A_1 = Increase activity.
 - A_2 = Improve sleep schedule.
 - A_3 = Adjust diet.
- **Reward Function ($R(s, a)$):**

$$R(s, a) = \Delta DT_{t+1} - \Delta DT_t$$

If $R(s, a) > 0$, the **user's condition improves**, reinforcing the action.

Q-Learning Update Rule : The RL agent refines recommendations using Q-learning:

$$Q(s, a) \leftarrow Q(s, a) + \alpha(r + \gamma \max_{a'} Q(s', a') - Q(s, a))$$

where:

- α = Learning rate.
- γ = Discount factor for future rewards.
- s' = Next state.

The RL model continuously improves intervention effectiveness, ensuring long-term health improvements.

7.4 Secure Feature Anonymization & Privacy-Preserving AI

Privacy preservation is a core component of this framework. Unlike traditional Federated Learning (FL), this system employs Secure Feature Anonymization, ensuring that AI models never process raw user data.

Privacy-Preserving Data Transformation: Each user record is **converted into de-identified feature vectors** before being analyzed:

$$F_t = \{f_1, f_2, \dots, f_n\}$$

where:

- f_i represents **an anonymized feature**, such as **heart rate trends** rather than raw heart rate values.
- **Personally Identifiable Information (PII)** is removed before AI processing.

Mathematical Guarantee of Privacy: To ensure data security, each feature vector undergoes differential privacy mechanisms:

$$f'_i = f_i + Lap(\lambda)$$

where $Lap(\lambda)$ is **Laplace noise** that obfuscates individual user data while preserving general trends. By separating data collection (Cloud Environment 1) from AI processing (Cloud Environment 2), this framework ensures privacy while maintaining AI effectiveness.

8 Validation of the Proposed Approach through Simulation and Real-World Data

We validated the Self-Healing Digital Twin Framework with both simulations and real-world user data.

Digital Twin Model Accuracy: The framework achieves low error rates (MAE: 450 steps, 2.8 BPM, 0.6 hours) in predicting user health trends, ensuring real-world usability.

Anomaly Detection Performance: High F1-scores (≥ 0.89) confirm reliable detection of heart rate spikes (96.2%), sleep deprivation (93.8%), and activity decrease (92.1%), with minimal false positives ($\leq 5.2\%$).

RL-Based Healing Performance: The RL model doubled adherence for disengaged users (from 20.5% to 47.2%), significantly outperforming static AI recommendations.

8.1 Simulation-Based Validation

We generated synthetic data for 10,000 virtual users spanning diverse lifestyles and health patterns.

Table 3: Digital Twin Model Prediction Accuracy

Metric	Steps (Steps)	Prediction	Heart Rate Prediction (BPM)	Sleep Duration Prediction (hours)
MAE	450 steps		2.8 BPM	0.6 hours
RMSE	570 steps		3.4 BPM	0.8 hours

As for Digital Twin Model Accuracy (Table 3), we trained on 70% of the simulated dataset and tested on 30%. Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) remained low

for steps, heart rate, and sleep predictions. As for Anomaly Detection (Table 4), the system identified health deviations (e.g., sudden heart-rate spikes) with high precision and recall, while keeping false positives low.

Table 4: Anomaly Detection Performance

Anomaly Type	Detection Rate (%)	False Positive Rate (%)	F1-score
Sudden Heart Rate Spike	96.2%	3.1%	0.94
Sleep Deprivation	93.8%	4.5%	0.91
Decrease in Physical Activity	92.1%	5.2%	0.89

As for RL-Based Healing Performance (Table 5), as Reinforcement Learning nearly doubled adherence among disengaged users compared to static AI recommendations, proving especially effective for low-participation groups.

Table 5: Reinforcement Learning-Based Healing Performance

User Engagement Level	Traditional AI-Based Interventions (%)	RL-Based Healing Interventions (%)
Highly Engaged Users	64.2%	79.5%
Partially Engaged Users	42.8%	65.4%
Disengaged Users	20.5%	47.2%

8.2 Real-World User Data Validation

We further deployed the framework for 3,000 Nutrosal platform users over 10 weeks. Engagement rose by over 25%, while dropout rates were cut nearly in half. Users consistently reported stronger motivation and more trust in AI-driven guidance (Results shown in Table 6). The introduction of the Digital Twin and RL-based healing significantly boosted engagement and adherence, reducing the dropout rate by nearly half (14.3%).

Table 6: Engagement & Adherence Improvement

Metric	Before Digital Twin (%)	After Digital Twin (%)	Improvement (%)
Engagement Retention Rate	41.3%	67.2%	+25.9%
Adherence to AI-Suggested Routines	32.4%	58.6%	+26.2%
Dropout Rate	28.7%	14.4%	-14.3%

8.3 User Feedback on AI-Driven Interventions

To assess qualitative feedback, users were surveyed at the end of the 10-week evaluation. The responses focused on AI-generated intervention relevance, effectiveness, and trustworthiness (Results shown in Table 7). A majority of users (86.9%) expressed trust and willingness to continue using the AI-driven system, validating the effectiveness of the Self-Healing Digital Twin approach.

Table 7: User Feedback on AI-Driven Interventions

Survey Question	Positive Responses (%)
"Did AI recommendations feel personalized?"	84.1%
"Did AI interventions improve your wellness routine?"	78.6%
"Would you continue using AI-driven recommendations?"	86.9%

8.4 Summary of Validation Findings

The Self-Healing Digital Twin Framework achieves robust accuracy and performance across key metrics. It reliably predicts user health states evidenced by low MAE and RMSE for step counts, heart rate, and sleep duration—and its anomaly detection

system attains a 94% F1-score, flagging major deviations such as sleep shortages or abnormal heart-rate spikes while minimizing false positives. Reinforcement Learning (RL)-based healing boosts intervention success rates by 30–45%, particularly benefiting low-engagement users. Real-world evaluations further show a 26% rise in engagement, notable reductions in dropout, and strong user trust, with 86.9% of participants deeming AI recommendations personalized and effective. These findings confirm the framework's scalability and privacy-preserving architecture, positioning it as a viable solution for personalized wellness management.

9 Conclusion

This paper introduces a Self-Healing Digital Twin Framework that merges Hybrid Generative AI, Reinforcement Learning (RL)-based self-healing, and a Privacy-Preserving AI (PPAI) framework for personalized wellness management. Unlike static AI-driven platforms, it uses real-time data from IoMT devices, wearables, and user logs to dynamically track and optimize interventions. A dual-cloud architecture anonymizes raw data before processing, balancing privacy with scalability and real-time updates. Through simulations and real-world deployment, the framework demonstrates its ability to: 1) Continuously model user health states 2) Detect anomalies 3) Personalize interventions using RL-based adaptive learning. As the Digital Twin evolves, user engagement and adherence increase, particularly among those with low initial compliance.

REFERENCES

- [1] D. Gupta, O. Kayode, S. Bhatt, M. Gupta, and A. S. Tosun, "Hierarchical Federated Learning based Anomaly Detection using Digital Twins for Smart Healthcare," in *2021 IEEE 7th International Conference on Collaboration and Internet Computing (CIC)*, Atlanta, GA, USA: IEEE, Dec. 2021, pp. 16–25. doi: 10.1109/CIC52973.2021.00013.
- [2] B. V. B and P. P. S, "Enhancing Cardiac Health: Digital Twin Technology for Real-Time Monitoring and Abnormality Detection of Human Heart," in *2024 4th International Conference on Machine Learning and Intelligent Systems Engineering (MLISE)*, Zhuhai, China: IEEE, Jun. 2024, pp. 219–224. doi: 10.1109/MLISE62164.2024.10674146.
- [3] M. Ali, F. Naeem, M. Tariq, and G. Kaddoum, "Federated Learning for Privacy Preservation in Smart Healthcare Systems: A Comprehensive Survey," *IEEE J. Biomed. Health Inform.*, vol. 27, no. 2, pp. 778–789, Feb. 2023, doi: 10.1109/JBHI.2022.3181823.
- [4] L. Zou, L. Xia, Z. Ding, J. Song, W. Liu, and D. Yin, "Reinforcement Learning to Optimize Long-term User Engagement in Recommender Systems," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Anchorage AK USA: ACM, Jul. 2019, pp. 2810–2818. doi: 10.1145/3292500.3330668.
- [5] A. K. Rai, M. Sharma, V. Saravanan, N. Shalini, M. Kumar, and P. Singh, "Personalized Healthcare Recommendations with Q-Learning Reinforcement Learning," in *2023 International Conference on Artificial Intelligence for Innovations in Healthcare Industries (ICAIIHI)*, Raipur, India: IEEE, Dec. 2023, pp. 1–6. doi: 10.1109/ICAIIHI57871.2023.10489223.
- [6] J. Jin, M. Kim, and S. D. Kim, "Personalized Health Assistant with Reinforcement Learning," in *2024 IEEE First International Conference on Artificial Intelligence for Medicine, Health and Care (AIMHC)*, Laguna Hills, CA, USA: IEEE, Feb. 2024, pp. 153–156. doi: 10.1109/AIMHC59811.2024.00034.