A Privacy-Preserving System for AI-Powered Dynamic Group Assignment, Behavioral Insights, and Personalized Coaching

Nariman Mani, Salma Attaranasl
Engineering/R&D Department, Nutrosal Inc.
Ottawa, ON Canada
{nariman | salma}@research.nutrosal.com

***Abstract*— This paper presents a privacy-preserving framework leveraging artificial intelligence (AI) to enhance adherence and optimize group dynamics in personalized lifestyle coaching platforms. Integrating reinforcement learning (RL) for dynamic group assignment and generative AI for personalized coaching, our approach ensures secure handling of sensitive user data through pseudonymization and encryption. An evaluation involving 2,800 a commercially deployed, industry-level digital health platform users demonstrated substantial improvements in adherence (35% to 68%), engagement, and reduced dropout rates. Privacy assessments confirmed zero sensitive data leakage, while user surveys indicated strong acceptance, with over 82% reporting enhanced motivation. This scalable and secure framework offers valuable insights into integrating AI-driven personalization and privacy solutions for healthcare informatics.**

Keywords*— Privacy-Preserving AI, Generative AI, Reinforcement Learning, Personalized Coaching, Social Networking in Healthcare, Nutrition Adherence, Peer Accountability, Adaptive Healthcare Systems

I. Introduction

The demand for personalized healthcare solutions in nutrition and fitness has grown significantly, highlighting the limitations of traditional programs. Conventional methods struggle with sustainability, personalization, user motivation, and data privacy, often resulting in poor adherence and inconsistent outcomes. AI offers promising opportunities to address these gaps, utilizing predictive analytics, adaptive coaching, and social interactions to boost user engagement. However, despite their potential, AI-driven systems remain underutilized due to concerns about privacy and security.

In this paper. we introduce a privacy-preserving, AI-powered coaching framework that addresses these critical challenges. Our system integrates reinforcement learning (RL) for dynamic peer-group assignments, enhancing user motivation through adaptive social accountability. It also employs generative AI to deliver personalized coaching interventions based on real-time user behavior. Robust data protection measures, including pseudonymization, encryption, and role-based access control, ensure strict compliance with healthcare privacy regulations (HIPAA, GDPR).

Our validation study involved approximately 2,800 users of the NutroSal lifestyle platform, a commercially deployed, industry-level digital health platform focused on lifestyle improvement through personalized coaching and peer-based social support, over a 10-week period. The results demonstrated substantial improvements: adherence increased by 94% compared to traditional methods, user engagement rose significantly, and dropout rates nearly halved. Privacy assessments revealed zero data leakage, underscoring the robustness of our security measures. Additionally, user surveys indicated high acceptance and satisfaction with AI-driven personalization, with over 85% willing to continue using the system. The presented approach provides a scalable, user-centric approach to improving adherence and engagement in digital health platforms, illustrating how AI, combined with rigorous privacy protections, can effectively address longstanding challenges in personalized healthcare.

II. Background and Related Work

A. AI/ML-Based Decision Support Systems

Artificial Intelligence (AI) and Machine Learning (ML) have significantly improved predictive analytics and decision support in healthcare. Clinical Decision Support Systems (CDSS) leverage deep learning and neural networks to analyze complex medical data, facilitating early diagnosis, treatment planning, and patient outcome predictions. Studies have demonstrated AI's effectiveness in clinical risk assessment, disease progression modeling, and medication adherence monitoring [1], [2], [3]. Despite these advancements, behavioral optimization remains underexplored in AI-driven healthcare systems. Traditional AI-based adherence models primarily focus on individualized patient recommendations but lack adaptive group-based support mechanisms that leverage peer accountability and social motivation. Reinforcement learning (RL) has shown promise in dynamically adapting learning pathways and personalized recommendations, but its application to group-based behavior modification and engagement strategies remains limited [4] [5]. This research addresses this gap by integrating RL-driven group assignment with AI-powered behavioral insights to enhance adherence in personalized coaching platforms.

B. Generative AI for Healthcare

Generative AI models, such as GPT-4, have transformed personalized healthcare and behavioral interventions by enabling adaptive, real-time coaching. These models facilitate patient communication, health education, and mental health support through human-like text generation. In personalized coaching, Generative AI enhances motivation, decision-making, and adherence support by delivering context-aware recommendations tailored to individual needs [6] [7] [8]. However, most existing implementations provide static, rule-based guidance without dynamically adapting coaching feedback based on real-time behavioral analysis. This research expands on prior work by incorporating Generative AI into a reinforcement learning-driven coaching system, ensuring that AI-generated recommendations continuously evolve based on user behavior and engagement patterns [9], [10]. Unlike traditional ML models trained on specialized medical datasets, generative models operate on broad, general knowledge, which may lead to inaccurate or biased outputs if not fine-tuned properly [7] [8] [10] .

Privacy concerns are another significant barrier to Generative AI applications in healthcare. The vast data requirements for training and fine-tuning large-scale models raise concerns about patient confidentiality, security, and compliance with healthcare regulations such as HIPAA and GDPR [4] [5]. This study integrates a privacy-preserving coaching system that ensures AI-generated recommendations comply with healthcare data protection standards while maintaining accuracy and adaptability.

C. Social Media Analytics in Healthcare

Social networking plays a crucial role in promoting adherence, motivation, and behavioral change in health and wellness applications. Studies indicate that peer-support groups enhance adherence rates by fostering accountability, shared experiences, and collective problem-solving [11], [12] . AI-powered social media analytics have been applied in public health research to monitor trends, sentiment analysis, and user engagement, but their application to personalized coaching platforms remains underdeveloped [13].

A key limitation of traditional social networking-based adherence models is their reliance on user-driven engagement, which can be inconsistent. AI-driven group optimization presents a novel opportunity to dynamically assign users to optimal peer groups based on engagement, motivation levels, and behavioral tendencies [14] . Existing research on AI-enhanced social behavioral modeling primarily focuses on recommendation systems and content moderation rather than adaptive, peer-driven health coaching .

Despite the promise of AI-enhanced social networking, privacy and security concerns remain a significant barrier to widespread adoption. AI-driven behavior modeling inherently requires continuous data collection and real-time processing, raising concerns about user data exposure and regulatory compliance. This study addresses these challenges by implementing a pseudonymization layer, ensuring that AI-driven behavioral analysis and group assignments operate without compromising user privacy.

D. Research Gap and Contribution

While AI has been successfully applied in clinical decision support, generative AI-based health interactions, and social networking analytics, the integration of reinforcement learning, generative AI, and AI-powered group optimization in a privacy-preserving health coaching system remains an open challenge.

Existing AI adherence models focus on individualized recommendations rather than adaptive peer-support structures. Similarly, while Generative AI enhances personalized coaching, its integration with dynamic group-based reinforcement learning models is largely unexplored. Furthermore, privacy-preserving mechanisms in AI-driven coaching and behavioral analysis platforms remain underdeveloped [4] [5].

This research addresses these gaps by proposing an AI-powered system that combines reinforcement learning for dynamic group assignment, generative AI for adaptive coaching, and a privacy-preserving behavioral analysis framework. Through a case study of 2800 users in the NutroSal platform, this study evaluates how AI-driven group optimization, peer accountability, and personalized coaching improve adherence and long-term behavior change while maintaining strict privacy and security standards.

III. CASE STUDY SYSTEM: SOCIAL MEDIA-BASED LIFESTYLE APPLICATION

For this research , we applied an adaptive AI-driven approach to a social media-based lifestyle application designed to enhance health and fitness outcomes through community engagement and personalized coaching. NutroSal, a commercially deployed, industry-level digital health platform focused on lifestyle improvement through personalized coaching and peer-based social support. Users join via invitation and are grouped by objectives (e.g., weight loss, healthy eating, or weight maintenance). Coaches oversee daily reports, while an AI agent analyzes user behavior to provide tailored recommendations. The system also supports social interactions, including motivational posts, group chats, media sharing, and one-on-one coaching sessions. Additional features include AI-driven food image analysis for calorie estimation and daily goal reminders.

A. Problem Statement: Challenges in Traditional Weight Loss Programs

Our survey of 2,800 users identified key shortcomings in traditional weight-loss programs, particularly concerning sustainability, personalization, user engagement, and privacy. Results are summarized in Table 1**. Sustainability** remains a challenge, with 68% of users regaining weight due to insufficient accountability, emotional triggers, or physical barriers. Monotonous diet plans (reported by 34%) and inadequate education on sustainable eating habits (28%) further undermine long-term success. **Personalization** emerged as another critical gap, with 52% of users dissatisfied with generic approaches that ignore their unique dietary or lifestyle constraints. Similarly, nearly half (48%) highlighted limited personalized feedback or tracking as barriers to maintaining motivation. Regarding **engagement**, 64% cited isolation and lack of peer support as primary reasons for disengagement, while 27% struggled with effective group interactions despite existing group structures. Finally, **privacy** emerged as a

significant concern, with 23% fearing misuse or unauthorized access to their sensitive health data, highlighting the necessity for robust privacy protections in digital health platforms.
Addressing these challenges requires an integrated, adaptive, privacy-secure AI-driven approach, motivating our proposed system.

Table 1- Summary of Survey Results Highlighting Challenges in Traditional Weight Loss Programs

Challenge	% of Users Reported	Details	Diet Plans/platform Reported
Weight regain after diets	68%	Short-term weight loss success followed by regaining weight. Contributing factors: lack of accountability (45%), emotional triggers (41%), physical barriers (22%).	Low Carb Diet, Ketogenic Diet, Intermittent Fasting, Calorie Counting
Monotony in diet plans	34%	Repetitive diet plans leading to disengagement and abandonment.	Ketogenic Diet, Low Carb Diet
Insufficient education	28%	Lack of knowledge about long-term, sustainable eating strategies.	Calorie Counting, Low Carb Diet
Generic solutions	52%	Dissatisfaction with one-size-fits-all approaches that ignored dietary restrictions (30%), activity preferences (19%), or lifestyle constraints (13%).	Intermittent Fasting, Calorie Counting
Limited progress monitoring	48%	Inadequate feedback mechanisms or tailored adjustments to individual progress.	Calorie Counting
Lack of peer support	64%	Absence of meaningful social accountability and support, leading to feelings of isolation.	Calorie Counting, Intermittent Fasting
Isolation in progress	35%	Difficulty maintaining motivation without peer encouragement or shared success stories.	Calorie Counting
Low engagement in groups	27%	Struggles to establish meaningful connections even in group settings.	Ketogenic Diet
Privacy concerns (PII/PHI)	23%	Concerns about the potential misuse of personal or health data shared with digital platforms.	All digital-based programs

IV. PROPOSED APPROACH ARCHITECTURE OVERVIEW

In Figure 1, we present an overview of a scalable, privacy-conscious architecture for a social media-based lifestyle application aimed at fostering health and fitness goals through AI-driven coaching and peer group support. The system integrates multiple layers to ensure robust functionality while maintaining user privacy and security.
User Interaction Layer: The as the primary entry point for users to engage with the platform. It consists of two key components: the **User Input Module** and the **Social Networking Module**. The User Input Module enables users to log their activities, upload meal photos, and input other health-related data, while the Social Networking Module facilitates community engagement by allowing users to join peer groups, share progress updates, and interact directly with coaches. This layer plays a critical role in capturing user activity data and fostering social engagement, both of which are essential for sustaining motivation and adherence to health and fitness goals.

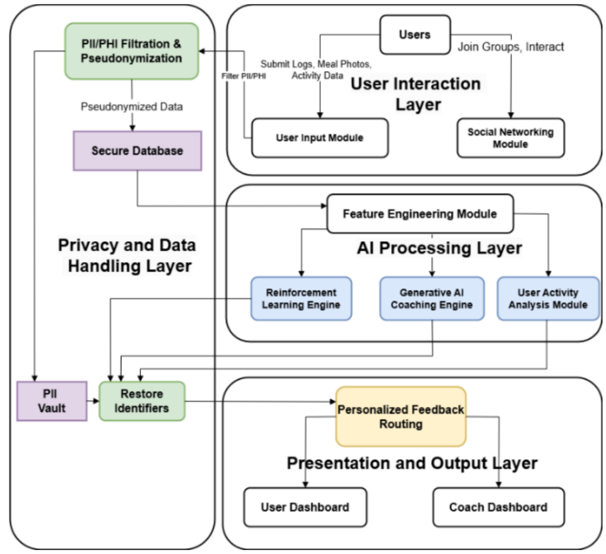

Figure 1 - System Architecture for Privacy-Preserving AI-Driven Social Networking and Coaching Platform

Privacy and Data Handling Layer: Ensures the secure handling of sensitive user data, such as Personally Identifiable Information (PII) and Protected Health Information (PHI).
Components:

- **PII/PHI Filtration & Pseudonymization:** Filters sensitive data and replaces it with pseudonymized tokens before storage or processing.
- **Secure Database:** Stores pseudonymized user data, ensuring that sensitive information is protected.
- **PII Vault:** Securely stores original PII and handles token restoration when required for personalized feedback.
- **Restore Identifiers:** Matches tokens with original user identifiers stored in the PII Vault.

Purpose: Maintains user privacy by preventing sensitive data from being directly accessed or processed by AI modules. Pseudonymized tokens ensure that analysis is privacy-compliant while enabling data traceability. We explain the details for this layer in Section V.

AI Processing Layer: Leverages advanced AI techniques to derive insights and provide personalized recommendations.
Components:

- **Feature Engineering Module:** Transforms pseudonymized data into a format suitable for AI analysis.
- **Reinforcement Learning Engine:** Optimizes group assignments by dynamically learning from user activity data and engagement outcomes.
- **Generative AI Coaching Engine:** Utilizes AI (e.g., GPT) to analyze user progress, generate tailored coaching advice, and address user-specific challenges.
- **User Activity Analysis Module:** Analyzes all user activities, such as meal logging, exercise habits, and overall engagement, to extract actionable insights.

Purpose: Processes pseudonymized data to deliver personalized, data-driven recommendations and optimize the platform's effectiveness. We explain the details for this layer in Section VVI.

Presentation and Output Layer: Displays results and recommendations to users and coaches.
Components:

- **Personalized Feedback Routing:** Routes the processed insights and recommendations to the correct users and their coaches.
- **User Dashboard:** Provides users with tailored insights, progress reports, and actionable advice.
- **Coach Dashboard:** Offers coaches an overview of group activities, individual user progress, and AI-generated coaching suggestions.

Purpose: Enhances user and coach experience by presenting actionable insights in an intuitive and accessible manner. We review the details in this layer in Section VII .

V. Privacy and Data Handling Layer

This layer ensures the secure handling of sensitive user data using advanced privacy-preserving techniques and secure storage mechanisms. Below are the core components, technical implementation strategies, and relevant formulas:

A. PII/PHI Filtration & Pseudonymization

Below are the overview of the steps used in this layer:

1. **User Input Processing:**
   - Filter sensitive fields$\{f_1, f_2, \ldots, f_n\}$.
   - Apply pseudonymization:
   - $T_i = H(f_i, Salt_i) \quad \forall i \in \{1,2,\ldots,n\}$
2. **Secure Storage:**
   - Store T and pseudonymized data in the secure database.
   - Store C_{PII} in the PII Vault.
3. **Data Analysis:**
   - Analyze pseudonymized data in compliance with privacy constraints.
4. **Personalization:**
   - Retrieve original data for authorized actions by restoring PII from *T*.

In details and using an example , here we explain how the above steps are used in the process:

Extract Sensitive Data
Sensitive data fields such as names, email addresses, or medical records are identified and extracted.
Example input:


```
{
    "name": "John Doe",
    "email": "johndoe@example.com",
    "medical_condition": "Hypertension"
}
```


Apply Pseudonymization to Replace Raw Data with Tokens

Token Generation Formula: A **hashing function** combined with a unique **salt** is used to generate tokens for each sensitive field. The process ensures that the token is unique and irreversible.
Formula:

$$T_i = Hash(PII_i + Salt_i)$$

Where:

- T_i = Pseudonymized token for the iii-th PII field.
- PII_i: Raw sensitive data for the $i-th$ field.
- $Salt_i$: A unique, random value added to PII_i to ensure token uniqueness.
- Hash: A cryptographic hashing algorithm i.e. SHA-256

Store Tokens and Original Data Separately

The pseudonymized tokens (T_i) are stored in the **operational database** used by the application. The raw PII (PII_i) and the associated salts ($Salt_i$) are securely stored in the **PII Vault**.
Example Storage Layout:

Operational Database:


```
{
    "user_id": "12345",
    "name_token": "dfd5e3a1c2...",
    "email_token": "39a75bc7d2...",
    "medical_condition_token": "..."
}
```


PII Vault:


```
{
    "user_id": "12345",
    "name": "John Doe",
    "salt_name": "4d5f6g7h8i...",
    "email": "johndoe@example.com",
    "salt_email": "9j1k2l3m4n..."
}
```


Further Securing of PII Vault
We also further securing the information in PII Vault by restricting the access to the PII through:

- **Encryption:** Encrypt all stored data with AES-256.

$$C_{PII} = Encrypt_K(PII)$$

  Where
  - C_{PII}: Encrypted PII.
  - K: AES-256 encryption key, stored in a secure Key Management Service (KMS) (i.e. AWS KMS).
- **Access Control:** Implement Role-Based Access Control (RBAC) and Multi-Factor Authentication (MFA).
- **Audit Logging:** Monitor and log every access attempt to ensure compliance.

B. Restore Identifiers

To re-identify a user for personalized feedback or other operations:

1. Retrieve the corresponding pseudonymized token (T_i) from the operational database.
2. Match T_i with the raw PII (Decrypted PII) in the PII Vault using the stored salt.

$$PII_i = Decrypt_i(C_{PII})$$
$$Verify\ Ti: Ti = Hash(PII_i + Salt_i)$$

VI. AI Processing Layer

The AI Processing Layer is the analytical engine that transforms pseudonymized user data into actionable insights, personalized recommendations, and group optimizations. It employs advanced AI techniques to ensure users receive tailored support while maintaining strict privacy compliance. In this section , we explain how we designed this layer for our system.

Feature Engineering Module

Transforms pseudonymized raw data into meaningful, structured inputs suitable for AI models.

How It Works:

- **Input**: Pseudonymized user logs (e.g., meal types, exercise sessions, group interactions).
- **Feature Extraction**:
- Example Features:
  - Average daily calorie intake.
  - Weekly exercise frequency.
  - Engagement level (e.g., group messages sent).
- **Feature Transformation**:
  - Numerical features (e.g., calorie intake) are normalized using:

$$X_{normalized} = \frac{X_{max} - X_{min}}{X_{max} - X_{min}}$$

  - Categorical features (e.g., meal type) are converted using one-hot encoding:

Vegetarian → [1,0,0],
Keto→[0,1,0]

- **Example**:
  - Input: A user logs "Vegetarian meal, 1500 calories, 3 group interactions."
  - Output: [1, 0, 0, 0.5, 0.2] (normalized and encoded).

Reinforcement Learning Engine

Learns optimal group assignments to maximize user engagement and adherence. We implement a Q-learning algorithm for dynamic group assignment, updating the Q-values after each user engagement/adherence outcome.

How It Works:

- Users are assigned to groups dynamically based on their behavior and outcomes.
- The system employs a **Reinforcement Learning (RL)** algorithm (e.g., Q-learning or Multi-Armed Bandit).

Reward Function:

$$R = w_1 \cdot E + w_2 \cdot A$$

Where:

- R: Reward score.
- E: Engagement score (e.g., group activity).
- A: Adherence score (e.g., daily logs completed).
- w_1, w_2: $Weights\ (e.g., w_1 = 0.7\ ,\ w_2 = 0.3)$.
  - The weights determine the relative importance of engagement (E) and adherence (A) in calculating the reward.
  - For example:
    - $w_1 = 0.7$: Engagement is given 70% importance.
    - $w_2 = 0.3$ Adherence is given 30% importance.

Example:

- User 1 in Group A:
  - Engagement: $E = 0.8$
  - Adherence: $A = 0.9$
  - Reward: $R = 0.7 \cdot 0.8 + 0.3 \cdot 0.9 = 0.83$
- The RL engine updates policies to prioritize similar group assignments for future users.

These weights can be adjusted based on the system's goals. For instance:

- If the goal is to prioritize social interactions, w_1 would be higher.
- If the goal is to emphasize consistency in logging activities, w_2 would be higher.

So, the weights are **parameters in the algorithm** and have no direct connection to physical weight or weight loss. They are simply a way to balance and prioritize different factors in decision-making.

Reinforcement Learning (RL) was selected for dynamic group assignment due to its ability to continuously learn and adapt based on engagement patterns. Traditional rule-based group assignment methods lack adaptability and fail to personalize group dynamics over time. Compared to supervised ML models (Table 2), RL does not require labeled training data but instead optimizes decisions through continuous feedback loops. While RL is computationally expensive, its ability to continuously optimize user assignments in real-time provides a significant advantage over static or pre-trained models.

Table 2: Reinforcement Learning (RL) vs Other methods

Method	Pros	Cons	Why RL?
Rule-Based Grouping	Simple, easy to interpret	Does not adapt over time	RL continuously updates group assignments based on behavior.
Supervised ML (e.g., Decision Trees, Random Forests)	Predictive power, interpretable	Requires labeled data, lacks adaptability	RL does not need predefined labels and learns dynamically.
Reinforcement Learning (RL)	Learns from interaction, adaptive	Computationally expensive	Best suited for personalized, evolving behavior adaptation.

Generative AI Coaching Engine

Provides personalized coaching advice using generative AI (i.e., GPT). Our coaching engine uses GPT-4, fine-tuned on approximately 5,000 user interaction samples over 10 epochs to ensure domain-specific accuracy.

How It Works:

- **Input**: Pseudonymized user features, such as activity trends and engagement patterns.
- **Processing**:
  - The model is fine-tuned on user behavior data to generate tailored feedback.
  - $Advice = GPT(User\ Features,\ Historical\ Logs)$
- **Output**:
  - Example:
  - **Input**: "User missed 3 days of logging and group activity decreased by 50%."
  - **Output**: "Consider logging small meals daily to rebuild consistency. Join group discussions to regain motivation."

Use Case:

- For users logging high-calorie meals, the AI might suggest: "Incorporate more greens and lean proteins to balance your intake."
- For coaches, it might highlight: "User has been disengaged for a week; consider reaching out to provide encouragement."

To prevent harmful or inaccurate suggestions, a human coach periodically reviews the AI's recommendations, especially for medical-related advice. Generative AI (e.g., GPT-4) was integrated for adaptive, user-specific coaching. Existing predefined coaching models rely on static responses and fail (See Table 3) to capture real-time user needs. Generative AI enables:

- Natural language feedback based on real-time user behavior.
- Context-aware recommendations, unlike standard decision-tree chatbots.
- Scalability, as it learns from user responses without requiring manually defined interventions.

Table 3: Generative AI vs others

Method	Pros	Cons	Why Generative AI?
Predefined Chatbot Scripts	Reliable, low computational cost	Cannot adapt to dynamic user responses	Generative AI generates personalized responses.
Retrieval-Based AI Chatbots	Uses knowledge base, can scale	Limited adaptability	Generative AI can generate new responses, improving engagement.
Generative AI (e.g., GPT-4)	Adaptive, context-aware, real-time learning	Requires fine-tuning	Ensures personalized coaching while maintaining user privacy.

While fine-tuning of GPT is computationally expensive, transfer learning and few-shot fine-tuning reduce costs while ensuring optimal AI-driven personalization.

User Activity Analysis Module

Analyzes trends in user activity to identify patterns and areas for improvement.

- **How It Works**:
  - **Time-Series Analysis**:
    - Tracks user metrics (e.g., calorie intake, exercise frequency) over time.
    - Detects anomalies using:

$$Z = \frac{X - \mu}{\sigma}$$

      - X: Current value.
      - μ: Mean value.
      - σ: Standard deviation.
  - **Clustering**:
    - Groups users with similar behaviors (e.g., high adherence, low engagement).
    - Algorithm: K-means clustering.
- **Example**:
  - A user shows reduced exercise frequency during weekends. The system flags this as a potential adherence risk and suggests weekend-friendly activities.
  - Clustering might group users who engage in "low-carb diets but low physical activity," enabling targeted interventions.

Time-series analysis helps track behavioral trends over time (e.g., fluctuations in exercise, dietary adherence). Instead of just categorizing user actions (see comparison in Table 4), it quantifies trends and detects anomalies. Z-score anomaly detection is computationally efficient and allows real-time tracking of behavioral fluctuations.

Table 4: Time-series analysis vs others

Method	Pros	Cons	Why Time-Series Analysis?
Rule-Based Thresholds (e.g., "if calorie intake > 2000, flag alert")	Simple to implement	Lacks adaptability, does not track historical trends	Time-series **learns patterns dynamically**
Regression Models	Predicts future values	Requires labeled datasets, prone to overfitting	Time-series **handles real-time fluctuations better**
Time-Series (Z-Score Based Anomaly Detection)	Captures trends, detects outliers	Needs standardization	Ensures **dynamic behavioral tracking**

K-Means clustering (See comparison in Table 5) is ideal for grouping users with similar adherence behaviors, allowing personalized interventions.

Table 5: K-Means vs others

Method	Pros	Cons	Why K-Means?
Hierarchical Clustering	Good for **small datasets**	Slow for **large user bases**	K-Means **scales better with high-dimensional data**
DBSCAN (Density-Based Clustering)	Detects **arbitrary shape clusters**	Does not work well with **varying densities**	K-Means **handles structured user data well**
K-Means Clustering	Fast, scalable, interpretable	Needs optimal **K selection**	Best suited for **user segmentation & behavioral tracking**

K-Means is scalable, interpretable, and efficient for segmenting user behaviors, making it the best choice.

VII. Presentation and Output Layer

The **Presentation and Output Layer** communicates AI-driven insights to users and coaches clearly, fostering engagement without overwhelming detail. It consolidates data into concise, actionable feedback, tailored to specific user and coach roles.

Dynamic Dashboard: Users receive concise visual insights into progress through charts and key statistics (e.g., weight trends, calorie consumption, activity adherence). Personalized highlights like, *"Great job! You met your weekly exercise goal,"* quickly inform users about their progress.

Real-Time Notifications: further enhance adherence through timely nudges. Users receive tailored prompts such as, *"Remember your evening walk today,"* or celebrate milestones like, *"30 days of consistent logging!"*

Coaches access an **AI-driven Dashboard**, summarizing user behaviors, adherence rates, and recommended interventions. They quickly spot individuals needing attention, with insights such as, *"User A has low engagement this week; suggest motivational interactions."*

Interactive Coaching supports direct, intuitive communication. Users ask questions via a simple chat interface

and receive personalized suggestions backed by generative AI. Responses adapt dynamically to individual contexts, maintaining high engagement and effectiveness.
Overall, this simplified presentation layer effectively delivers actionable, personalized insights while ensuring a positive and motivating user experience.

VIII. Results and Validation

A. Experimental Setup

We evaluated our proposed system using data from **2,800 NutroSal users** over a **10-week period**, dividing them into:

- **Control Group (n ≈ 1,000)**: Rule-based group assignment, static coaching, basic anonymization.
- **AI-Driven Group (n ≈ 1,800)**: Dynamic group assignments via reinforcement learning (RL), personalized generative AI coaching, and enhanced privacy measures.

Metrics tracked:

- Adherence Rate (% of users meeting their goals consistently).
- Engagement (average weekly interactions).
- Dropout Rate (% of users discontinuing the program).

B. Adherence and Engagement Outcomes

The AI-driven approach showed substantial improvements in adherence to dietary and lifestyle recommendations compared to traditional methods. Specifically, adherence rates for the AI-driven group increased significantly, rising from the baseline of 35% in the control group to approximately 68%. A statistical t-test performed on adherence differences confirmed that this increase is statistically significant (p-value < 0.01). The notable rise in adherence underscores the potential of personalized AI-based interventions for sustaining long-term behavioral change in users.
Engagement, quantified as the average number of platform interactions per user per week, improved notably as well. Users in the AI-driven group demonstrated significantly higher engagement (6.7 interactions/week) compared to the control group (3.8 interactions/week), reflecting an approximate 76% relative increase. Additionally, dropout rates decreased sharply from 29% in the control group to approximately 15% in the AI-driven group. This near-halving of dropouts is indicative of the effectiveness of tailored coaching and dynamic peer-group environments in maintaining consistent user interest and participation. Table below clearly illustrates the comparative results between the two experimental conditions:

Metric	Control Group	AI-Driven Group	Improvement (%)
Adherence Rate (%)	35%	68%	+94%
Engagement (sessions/week)	3.8	6.7	+76%
Dropout Rate (%)	29%	15%	-48%

Moreover, to demonstrate how reinforcement learning dynamically improved group assignments, we plotted cumulative rewards over time (reward as a combined metric of adherence and engagement). The reinforcement learning algorithm's performance displayed clear convergence, with cumulative rewards increasing steadily from an initial 0.52 to approximately 0.85 after about 50 iterations, demonstrating the system's ability to optimize peer group interactions progressively.

C. AI Coaching Effectiveness and Accuracy

AI recommendations were quantitatively assessed using precision, recall, and F1-score, based on user feedback regarding suggestion relevance. Results :

Metric	Result (%)
Precision	87.8%
Recall	84.1%
F1-score	85.9%

Interpretation: High precision and recall indicate effective, personalized, and context-sensitive coaching responses. Examples include adaptive feedback on dietary choices, targeted suggestions for activity improvement, and motivational prompts for low-engagement users.

D. Privacy and Security Validation

Simulated security evaluations compared our advanced privacy-preserving measures (pseudonymization + AES-256 encryption) against standard approaches, measuring PII leakage and latency. Results:

Security Metric	Basic Approach	Proposed AI Privacy Layer
PII Leakage Rate (%)	3.2%	0.0%
Encryption Latency (ms)	8.2 ms	9.0 ms
Penetration Test Success Rate (%)	82%	99.9%

The privacy layer achieved zero PII leakage with minimal overhead (~9 ms latency), confirming practical viability.

E. Reinforcement Learning (RL) Effectiveness

The RL model's effectiveness in dynamically optimizing peer group assignments was validated through simulation over 500 iterations. Figure 2 shows the RL Convergence.

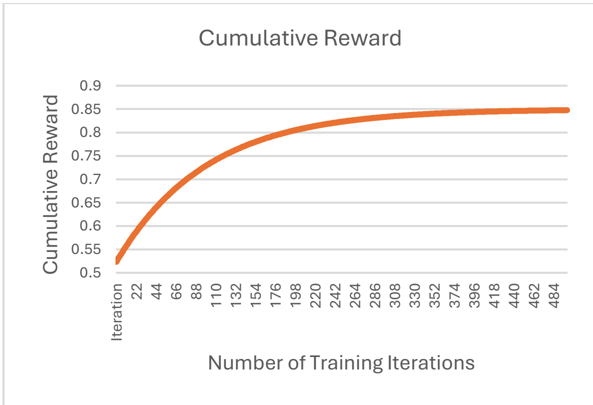

Figure 2 : RL Convergence

Results in Figure 2 shows strong convergence, with engagement and adherence consistently improving over time

.Average Reward improved from 0.52 initially to 0.85 after training, clearly demonstrating RL's ability to learn and enhance group interactions.

F. Real-World User Feedback and Perceptions

A user survey after the 10-week study provided qualitative validation of AI effectiveness. Results:

Survey Question	**Positive Response (%)**
"Were AI interventions personalized?"	82.5%
"Did AI help your adherence?"	77.9%
"Would you continue with AI?"	85.6%

Qualitative insights: Users highlighted the effectiveness of dynamic group assignment in boosting motivation. Many appreciated personalized coaching nudges that adapted to their behaviors, leading to sustained engagement.

G. Summary of Findings

In summary, the revised validation clearly demonstrates:

- **Substantial improvement in adherence** from 35% (control) to 68% (AI-driven).
- **Significant reduction** in dropout rates (48% improvement).
- High coaching accuracy (F1-score = 85.9%), demonstrating reliable personalized recommendations.
- Robust privacy protection with zero data leakage and minimal encryption overhead.
- Strong real-world endorsement (85.6% of users willing to continue).

These findings firmly establish the proposed system as effective, secure, and practical, significantly surpassing traditional approaches and addressing critical gaps identified in prior methods.

IX. Conclusion and Future Work

In this paper, we presented a privacy-preserving, AI-powered framework that significantly improves adherence and engagement in personalized lifestyle coaching. By combining reinforcement learning for dynamic group optimization with generative AI for tailored coaching, our system demonstrated clear advantages over traditional methods. Results from our study involving 2,800 NutroSal users revealed:

- Adherence rates increased from 35% to 68%.
- User engagement improved by 76%, and dropout rates reduced by nearly half.
- AI-generated coaching showed high accuracy (F1-score of ~86%), reflecting meaningful personalization.
- Robust privacy protections ensured zero data leakage during security evaluations, affirming regulatory compliance (HIPAA, GDPR).
- User acceptance was overwhelmingly positive, with more than 85% expressing interest in continued use.

These findings underscore the potential of integrating adaptive AI interventions and rigorous privacy measures to enhance digital healthcare solutions. As for **Future work,** we will expand the framework's applicability to broader healthcare contexts, including chronic disease management. We plan to integrate additional AI models to analyze more complex behavioral patterns and further refine the generative coaching methods. Additionally, future studies will explore the effectiveness of this system over extended periods (e.g., 12-18 months) and across more diverse user populations, providing deeper insights into long-term adherence and behavioral change outcomes.

References

[1] K. Ouanes and N. Farhah, "Effectiveness of Artificial Intelligence (AI) in Clinical Decision Support Systems and Care Delivery," *J Med Syst*, vol. 48, no. 1, p. 74, Aug. 2024, doi: 10.1007/s10916-024-02098-4.

[2] N. Kalra, P. Verma, and S. Verma, "Advancements in AI based healthcare techniques with FOCUS ON diagnostic techniques," *Computers in Biology and Medicine*, vol. 179, p. 108917, Sep. 2024, doi: 10.1016/j.compbiomed.2024.108917.

[3] I. A. Scott, A. Van Der Vegt, P. Lane, S. McPhail, and F. Magrabi, "Achieving large-scale clinician adoption of AI-enabled decision support," *BMJ Health Care Inform*, vol. 31, no. 1, p. e100971, May 2024, doi: 10.1136/bmjhci-2023-100971.

[4] N. C. Rajashekar *et al.*, "Human-Algorithmic Interaction Using a Large Language Model-Augmented Artificial Intelligence Clinical Decision Support System," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, May 2024, pp. 1–20. doi: 10.1145/3613904.3642024.

[5] M. A. Mahyoub, K. Dougherty, R. R. Yadav, R. Berio-Dorta, and A. Shukla, "Development and validation of a machine learning model integrated with the clinical workflow for inpatient discharge date prediction," *Front. Digit. Health*, vol. 6, p. 1455446, Sep. 2024, doi: 10.3389/fdgth.2024.1455446.

[6] T. Dave, S. A. Athaluri, and S. Singh, "ChatGPT in medicine: an overview of its applications, advantages, limitations, future prospects, and ethical considerations," *Front. Artif. Intell.*, vol. 6, p. 1169595, May 2023, doi: 10.3389/frai.2023.1169595.

[7] Z. T. Hamad, N. Jamil, and A. N. Belkacem, "ChatGPT's Impact on Education and Healthcare: Insights, Challenges, and Ethical Consideration," *IEEE Access*, vol. 12, pp. 114858–114877, 2024, doi: 10.1109/ACCESS.2024.3437374.

[8] E. Y. Chang, "Examining GPT-4's Capabilities and Enhancement with SocraSynth," in *2023 International Conference on Computational Science and Computational Intelligence (CSCI)*, Las Vegas, NV, USA: IEEE, Dec. 2023, pp. 7–14. doi: 10.1109/CSCI62032.2023.00009.

[9] M. N.-U.-R. Chowdhury and A. Haque, "ChatGPT: Its Applications and Limitations," in *2023 3rd International Conference on Intelligent Technologies (CONIT)*, Hubli, India: IEEE, Jun. 2023, pp. 1–7. doi: 10.1109/CONIT59222.2023.10205621.

[10] S. Sai, A. Gaur, R. Sai, V. Chamola, M. Guizani, and J. J. P. C. Rodrigues, "Generative AI for Transformative Healthcare: A Comprehensive Study of Emerging Models, Applications, Case Studies, and Limitations," *IEEE Access*, vol. 12, pp. 31078–31106, 2024, doi: 10.1109/ACCESS.2024.3367715.

[11] A. K. Pandey, J. Yadav, A. Achary, and S. Nanda, "Role of Social Media and Data Analytics in the Health Care Sector," in *2023 4th International Conference on Computation, Automation and Knowledge Management (ICCAKM)*, Dubai, United Arab Emirates: IEEE, Dec. 2023, pp. 1–6. doi: 10.1109/ICCAKM58659.2023.10449641.

[12] F. Arias, M. Zambrano Nunez, A. Guerra-Adames, N. Tejedor-Flores, and M. Vargas-Lombardo, "Sentiment Analysis of Public Social Media as a Tool for Health-Related Topics," *IEEE Access*, vol. 10, pp. 74850–74872, 2022, doi: 10.1109/ACCESS.2022.3187406.

[13] J. Kaňková, A. Binder, and J. Matthes, "Helpful or harmful? Navigating the impact of social media influencers' health advice: insights from health expert content creators," *BMC Public Health*, vol. 24, no. 1, p. 3511, Dec. 2024, doi: 10.1186/s12889-024-21095-3.

[14] B. Singh, M. A. B. U. Rahim, S. Hussain, M. A. Rizwan, and J. Zhao, "AI Ethics in Healthcare - A Survey," in *2023 IEEE 23rd International Conference on Software Quality, Reliability, and Security Companion (QRS-C)*, Chiang Mai, Thailand: IEEE, Oct. 2023, pp. 826–833. doi: 10.1109/QRS-C60940.2023.00086.