Review

Sequential Behavior Modeling and Unsupervised Data Augmentation for Personalized Healthcare Recommendations Using AI

Jiayu Chen [1,*]

[1] School of Artificial Intelligence, Guilin University of Electronic Technology, Guilin, Guangxi, China
* Correspondence: Jiayu Chen, School of Artificial Intelligence, Guilin University of Electronic Technology, Guilin, Guangxi, China

Abstract: Personalized healthcare recommendations powered by artificial intelligence (AI) hold immense promise for improving patient outcomes and healthcare efficiency. This review paper examines the current landscape of sequential behavior modeling and unsupervised data augmentation techniques in the context of personalized healthcare recommendations. We delve into various methods for capturing the temporal dependencies in patient data, such as recurrent neural networks (RNNs), transformers, and Markov models, and assess their effectiveness in predicting future health events or treatment responses. Furthermore, we explore unsupervised data augmentation strategies, including generative adversarial networks (GANs), variational autoencoders (VAEs), and rule-based methods, which aim to enhance the quality and diversity of patient data, especially in scenarios with limited labeled information. The paper synthesizes existing research, compares different approaches, and identifies key challenges and opportunities in this rapidly evolving field. Finally, we discuss potential future directions for integrating sequential modeling and data augmentation techniques to advance personalized healthcare recommendations using AI. This review provides a comprehensive overview for researchers and practitioners interested in leveraging AI to improve healthcare delivery and patient well-being.

Keywords: personalized healthcare; sequential behavior modeling; unsupervised data augmentation; artificial intelligence; recommendation systems; recurrent neural networks; generative adversarial networks

Received: 02 December 2025
Revised: 12 January 2026
Accepted: 27 January 2026
Published: 30 January 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Motivation and Background

Personalized healthcare recommendations are gaining prominence due to their potential to revolutionize patient care. Traditional healthcare systems often struggle to provide tailored treatments, leading to suboptimal outcomes. Artificial intelligence (AI) offers a transformative solution by enabling data-driven insights and personalized interventions. AI algorithms can analyze vast amounts of patient data, including medical history, lifestyle factors, and genomic information, to generate customized recommendations for treatment, prevention, and wellness. This personalized approach addresses the limitations of one-size-fits-all medicine and promises to improve patient outcomes and reduce healthcare costs. The goal is to provide the right intervention, to the right patient, at the right time, maximizing the effectiveness of healthcare delivery [1].

1.2. Scope and Contributions

This paper's scope encompasses sequential behavior modeling and unsupervised data augmentation within personalized healthcare recommendations. We focus on methods leveraging AI to predict individual health trajectories and tailor interventions. Our key contributions include a comprehensive review of existing sequential models, such as Markov models and recurrent neural networks, and unsupervised data augmentation techniques like generative adversarial networks (GANs) and variational autoencoders (VAEs) for addressing data scarcity. We provide a comparative analysis of these approaches, highlighting their strengths and weaknesses in the context of healthcare. Finally, we discuss promising future research directions, including the integration of causal inference and the development of more robust and interpretable models for personalized healthcare [2].

2. Historical Overview of Personalized Healthcare Recommendations

2.1. Early Approaches to Healthcare Recommendations

Early personalized healthcare recommendations relied on rule-based systems, encoding expert knowledge into if-then statements. These systems, while interpretable, struggled with the complexity and variability of patient data. Collaborative filtering emerged, leveraging similarities between patients or treatments to predict preferences. However, these methods often suffered from data sparsity, where limited interaction data hindered accurate recommendations. Furthermore, early approaches largely ignored the sequential nature of healthcare events, treating patient histories as static snapshots rather than dynamic processes. This limited their ability to capture evolving health conditions and treatment responses over time, as shown in Table 1.

Table 1. Comparison of Early Recommendation Techniques.

Feature	Rule-Based Systems	Collaborative Filtering
Foundation	Expert knowledge encoded as if-then rules.	Similarities between patients or treatments.
Strengths	Interpretable and easily understandable.	Leverages data patterns to predict preferences.
Weaknesses	Struggles with complexity and variability of patient data; difficult to adapt.	Suffers from data sparsity, limiting accuracy.
Data Handling	Treats patient data as static.	Primarily focuses on static similarities.
Temporal Awareness	Ignores the sequential nature of healthcare events.	Largely ignores the sequential nature of healthcare events.
Mathematical Representation	Symbolic logic and rule execution.	Similarity metrics (e.g., cosine similarity, Pearson correlation).
Adaptability	Low adaptability to new data or evolving knowledge.	Potentially adaptable, but limited by data sparsity.
Example Use Case	Simple treatment guidelines based on limited criteria (if condition then treatment).	Recommending medications based on patients with similar diagnoses and demographics.

2.2. Evolution of AI-Powered Recommendation Systems

AI-powered healthcare recommendation systems have evolved significantly, transitioning from rule-based systems to sophisticated machine learning (ML) models. Early systems relied on expert knowledge and predefined rules, limiting their adaptability. The introduction of ML, particularly techniques like collaborative filtering and matrix factorization, enabled personalized recommendations based on patient similarities and treatment outcomes. Deep learning (DL) further advanced the field by allowing the modeling of complex, non-linear relationships within high-dimensional healthcare data, such as electronic health records and genomic information [3]. DL models, including recurrent neural networks and transformers, excel at capturing temporal dependencies and intricate patterns, leading to more accurate and context-aware recommendations. This shift allows for better prediction of optimal treatments and preventative measures, ultimately improving patient outcomes, as summarized in Table 2.

Table 2. Timeline of AI Adoption in Healthcare Recommendations.

Period	Technology	Description	Impact
Early Stages	Rule-Based Systems	Relied on expert knowledge and predefined rules. Limited adaptability and personalization.	Basic recommendations based on explicit guidelines.
Transition Phase	Machine Learning (ML) - Collaborative Filtering, Matrix Factorization	Enabled personalized recommendations based on patient similarities and treatment outcomes.	Improved personalization and prediction accuracy compared to rule-based systems.
Advanced Stage	Deep Learning (DL) - Recurrent Neural Networks, Transformers	Modeled complex, non-linear relationships within high-dimensional healthcare data. Captures temporal dependencies and intricate patterns.	More accurate and context-aware recommendations, leading to better prediction of optimal treatments and preventative measures.

2.3. The Rise of Sequential Modeling

The shift towards sequential modeling arose from the limitations of static approaches in capturing the temporal dependencies inherent in patient healthcare journeys. Traditional methods often treated patient data as independent snapshots, neglecting the crucial influence of past treatments and diagnoses on future health outcomes. Sequential models, like Hidden Markov Models and Recurrent Neural Networks, address this by explicitly considering the order of events [4]. For example, predicting medication adherence benefits from understanding the sequence of past prescription refills and doctor's appointments. Similarly, forecasting disease progression leverages the temporal order of symptoms and lab results to provide more accurate and personalized recommendations, moving beyond simple correlation to causal inference where possible, especially when considering time-varying covariates x_t and outcomes y_t.

3. Sequential Behavior Modeling for Personalized Recommendations: Core Theme A

3.1. Recurrent Neural Networks (RNNs) and Their Variants

Recurrent Neural Networks (RNNs) and their variants, such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), are powerful tools for modeling sequential patient data. Their ability to capture temporal dependencies makes them well-suited for predicting future health events. RNNs process sequences by maintaining a hidden state that is updated at each time step, allowing them to "remember" past information. LSTMs and GRUs address the vanishing gradient problem of standard RNNs, enabling them to learn long-range dependencies [5]. For example, LSTMs have been successfully applied to predict hospital readmission based on a patient's history of diagnoses, procedures, and medications. GRUs have shown promise in predicting disease progression, such as the onset of diabetes, by analyzing longitudinal patient records. The input sequence x_t at time t influences the hidden state h_t, which then contributes to the prediction $\hat{y}_t$.

3.2. Transformer-Based Models

Transformer-based models, originally developed for natural language processing, have shown promise in modeling sequential healthcare data. Architectures like BERT and GPT excel at capturing long-range dependencies and contextual information within patient histories [6]. This is crucial as medical events are often interrelated across extended periods. For instance, transformers can be used to predict treatment outcomes by analyzing sequences of diagnoses, medications, and procedures. The self-attention mechanism allows the model to weigh the importance of different events in the sequence, identifying critical factors influencing the patient's health trajectory. Furthermore, transformers can identify high-risk patients by learning complex patterns in their medical records, potentially enabling proactive interventions and improved patient care. The ability to process variable-length sequences makes them particularly suitable for handling the diverse and often incomplete nature of healthcare data.

3.3. Markov Models and Hidden Markov Models (HMMs)

Markov Models (MMs) and Hidden Markov Models (HMMs) offer a foundational approach to modeling sequential patient data, prized for their simplicity and interpretability. MMs assume that the future state depends only on the present state, a property known as the Markov property. This allows for straightforward modeling of transitions between health states. HMMs extend this by introducing hidden states, where observed patient data (e.g., symptoms, lab results) are probabilistically linked to underlying, unobserved disease states. For example, an HMM could model the progression of a chronic disease, where observed symptoms are used to infer the hidden stage of the disease. The probability of transitioning from state i to state j is denoted as $P(S_{t+1} = j|S_t = i)$. These models have been applied to predict disease states, treatment transitions, and even patient adherence to medication regimens, providing valuable insights for personalized healthcare recommendations [7].

4. Unsupervised Data Augmentation Techniques: Core Theme B

4.1. Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) offer a powerful approach to unsupervised data augmentation by learning the underlying distribution of patient data and generating synthetic samples. These synthetic data points can then be used to supplement the original dataset, improving the robustness and generalization ability of healthcare recommendation models. Several GAN architectures have been explored, including Vanilla GANs, Conditional GANs (CGANs), and Wasserstein GANs (WGANs). CGANs, for instance, allow for the generation of synthetic data conditioned on specific patient attributes like age or disease status, providing more targeted data augmentation. The effectiveness of GANs hinges on their ability to capture complex relationships within patient data, such as correlations between medical history, lifestyle factors, and treatment outcomes [8]. By training recommendation systems on a combination of real and GAN-generated data, we can mitigate issues related to data scarcity and improve recommendation accuracy, particularly for rare diseases or underrepresented patient subgroups where the number of samples N is small [9].

4.2. Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs) offer a powerful approach for learning latent representations of complex patient data and generating synthetic samples. By encoding patient features into a lower-dimensional latent space, VAEs capture the underlying structure and dependencies within the data. The encoder maps the input data x to a distribution over the latent space z, typically parameterized by a mean μ and variance σ^2. The decoder then reconstructs the input from a sampled latent vector z. This probabilistic approach allows VAEs to handle uncertainty inherent in patient data. Data augmentation is achieved by sampling new latent vectors from the learned distribution and decoding them to generate synthetic patient records. These augmented data points can then be used to train more robust and generalizable recommendation models, particularly when dealing with limited or imbalanced datasets, improving the model's ability to handle unseen patient variations [10].

4.3. Rule-Based Data Augmentation

Rule-based data augmentation leverages domain expertise and clinical guidelines to generate synthetic patient data. These methods offer simplicity and high interpretability, making them valuable for augmenting datasets where understanding the augmentation process is crucial. For instance, if a patient record indicates hypertension and high sodium intake, a rule could generate a new record with similar characteristics but increased dosage of antihypertensive medication, reflecting a potential clinical intervention [11]. Another example involves modifying existing records; if a patient's BMI is below a certain threshold and they report fatigue, a rule could introduce symptoms consistent with anemia, with a probability reflecting the prevalence of anemia in similar patients. Such rules can also be used to create counterfactual examples, exploring the impact of different treatment pathways or lifestyle changes on patient outcomes, as illustrated in Table 3.

Table 3. Examples of Rule-Based Data Augmentation.

Scenario	Rule Description	Augmentation Effect
Hypertension & High Sodium Intake	If a patient has hypertension and high sodium intake, increase the dosage of antihypertensive medication.	Simulates a clinical intervention to manage hypertension given poor dietary habits.
Low BMI & Fatigue	If a patient's BMI is below a certain threshold and they report fatigue, introduce symptoms consistent with anemia with a probability reflecting anemia prevalence.	Introduces anemia symptoms based on known associations between low BMI, fatigue, and anemia.
Treatment Pathway Exploration	If a patient is on treatment A, create a counterfactual example where they are switched to treatment B.	Explores the impact of different treatment pathways on patient outcomes.
Lifestyle Change Impact	If a patient has a sedentary lifestyle, create a counterfactual example with increased physical activity.	Explores the impact of lifestyle changes on patient outcomes.

5. Comparison of Methods and Challenges

5.1. Comparative Analysis of Sequential Modeling Techniques

Sequential modeling techniques offer varying trade-offs for personalized healthcare recommendations. Markov models, while computationally efficient and interpretable, suffer from a strong Markov assumption, limiting their ability to capture long-range dependencies in patient histories [12]. Recurrent Neural Networks (RNNs), particularly LSTMs and GRUs, address this limitation but can be computationally expensive and challenging to train, especially with long sequences due to vanishing gradients. Transformers, leveraging attention mechanisms, excel at capturing long-range dependencies and offer parallelization benefits, leading to improved accuracy. However, they demand significant computational resources and data, potentially hindering their applicability in data-scarce healthcare settings. The choice of model depends on the balance between accuracy, interpretability, and computational constraints, considering factors like dataset size (N), sequence length (T), and desired latency (L), as summarized in Table 4.

Table 4. Performance comparison.

Model	Advantages	Disadvantages	Computational Cost	Data Requirements	Interpretability
Markov Models	Computationally efficient, interpretable	Strong Markov assumption, limited long-range dependency capture	Low	Low	High
RNNs (LSTMs, GRUs)	Captures long-range dependencies	Computationally expensive, challenging to train (vanishing gradients)	Medium	Medium	Low
Transformers	Excels at capturing long-range dependencies, parallelization benefits, improved accuracy	High computational resources and data requirements	High	High	Low

5.2. Challenges and Limitations

Sequential behavior modeling and unsupervised data augmentation in healthcare face several challenges. Data sparsity is a major hurdle, as patient histories are often incomplete or irregular, hindering accurate model training [13]. Noise in electronic health records, stemming from inconsistencies and errors, further complicates the learning process. Privacy concerns surrounding sensitive patient data necessitate careful consideration of anonymization techniques and data governance policies. Unsupervised data augmentation methods, while promising, can introduce biases or generate unrealistic synthetic data, potentially leading to flawed recommendations. Ethical considerations, such as fairness and transparency, are paramount to ensure equitable and trustworthy AI-driven healthcare solutions. The computational cost associated with complex sequential models and augmentation techniques can also be a limiting factor, especially when dealing with large datasets [14].

6. Future Perspectives

6.1. Integration of Multi-Modal Data

Integrating multi-modal data, including electronic health records, imaging data (e.g., X-rays, MRIs), and genomic data, holds immense promise for refining personalized healthcare recommendations. By fusing these diverse data streams, AI models can gain a more holistic understanding of a patient's health status and predict future needs with greater accuracy. Challenges include data heterogeneity, alignment of different modalities, and ensuring data privacy [15]. Opportunities lie in developing novel AI architectures capable of effectively learning from complex, multi-modal representations, ultimately leading to more precise and proactive healthcare interventions tailored to individual patient profiles based on variables like x, y, and z [16].

6.2. Explainable AI (XAI) and Trustworthy Recommendations

Explainable AI (XAI) is crucial for fostering trust in AI-driven healthcare recommendations. Patients and clinicians need to understand why a particular recommendation is made to assess its validity and appropriateness. Techniques like rule-based systems, attention mechanisms, and SHAP values can provide insights into the model's decision-making process. Visualizations and simplified explanations, avoiding complex jargon, are essential for interpretability. Furthermore, quantifying the uncertainty associated with each recommendation, perhaps through confidence intervals or probability distributions of potential outcomes $P(outcome|recommendation)$, can enhance transparency and facilitate informed decision-making [17].

7. Conclusion

7.1. Summary of Key Findings

This review highlights advancements in modeling sequential patient data for personalized recommendations. Unsupervised data augmentation techniques, particularly those leveraging generative models, effectively address data scarcity and improve recommendation accuracy, demonstrating significant potential for enhancing personalized healthcare.

7.2. Concluding Remarks

AI-powered personalized healthcare recommendations hold immense potential for transforming healthcare delivery. By leveraging sequential behavior modeling and unsupervised data augmentation, we can significantly improve patient outcomes through tailored interventions and proactive care, ultimately paving the way for a more efficient and patient-centric healthcare ecosystem where $f(x)$ represents individual health trajectories.

References

1. J. L. Zhao, "Graph-based deep dive on AI startup revenue composition and venture capital network effect," Economics and Management Innovation, vol. 3, no. 1, pp. 27–36, 2026.
2. K. Malhotra, "Applications of Sequential Mining and Data Modeling for Personalized Medicine," Doctoral dissertation, Georgia Institute of Technology, 2016.
3. S. Li, K. Liu, and X. Chen, "A context-aware personalized recommendation framework integrating user clustering and BERT-based sentiment analysis," Journal of Computer, Signal, and System Research, vol. 2, no. 6, pp. 100–108, 2025.
4. T. C. Frommeyer et al., "Reinforcement learning and its clinical applications within healthcare: A systematic review of precision medicine and dynamic treatment regimes," Healthcare, vol. 13, no. 14, p. 1752, 2025.
5. D. S. Zois, "Sequential decision-making in healthcare IoT: Real-time health monitoring, treatments and interventions," in 2016 IEEE 3rd World Forum on Internet of Things (WF-IoT), 2016, pp. 24-29.
6. J. Mulani et al., "Deep reinforcement learning based personalized health recommendations," in Deep learning techniques for biomedical and health informatics, Cham: Springer International Publishing, 2019, pp. 231-255.
7. K. Banumathi et al., "Reinforcement learning in personalized medicine: A comprehensive review of treatment optimization strategies," Cureus, vol. 17, no. 4, 2025.
8. R. M. Townsley, "An examination of individual sequential behavior modeling in the context of healthcare simulation," North Carolina State University, 2018.
9. A. D. Pananos, "Bayesian Pharmacokinetic Models for Inference and Optimal Sequential Decision Making with Applications in Personalized Medicine," Doctoral dissertation, The University of Western Ontario (Canada), 2022.
10. E. Ferrara, "Large language models for wearable sensor-based human activity recognition, health monitoring, and behavioral modeling: A survey of early trends, datasets, and challenges," Sensors, vol. 24, no. 15, 5045, 2024.
11. J. Jin, T. Zhu, and C. Li, "Graph Neural Network-Based Prediction Framework for Protein-Ligand Binding Affinity: A Case Study on Pediatric Gastrointestinal Disease Targets," Journal of Medicine and Life Sciences, vol. 1, no. 3, pp. 136–142, 2025.
12. P. Liao, K. Greenewald, P. Klasnja, and S. Murphy, "Personalized heartsteps: A reinforcement learning algorithm for optimizing physical activity," Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies, vol. 4, no. 1, pp. 1-22, 2020.

13. C. L. Cheong, "Research on AI security strategies and practical approaches for risk management," Journal of Computer, Signal, and System Research, vol. 2, no. 7, pp. 98–115, 2025.
14. L. D. Preuett, "Learning Personalized Health Recommendations via Offline Reinforcement Learning," in Proceedings of the 18th ACM Conference on Recommender Systems, 2024, pp. 1355-1357.
15. Z. Huang et al., "Probabilistic modeling personalized treatment pathways using electronic health records," Journal of Biomedical Informatics, vol. 86, pp. 33-48, 2018.
16. Y. S. Cai, "Organizational restructuring of fintech enterprises: A strategic study balancing compliance and innovation," Financial Economics Insights, vol. 3, no. 1, pp. 1–10, 2026.
17. R. Luo, X. Chen, and Z. Ding, "SeqUDA-Rec: Sequential user behavior enhanced recommendation via global unsupervised data augmentation for personalized content marketing," arXiv preprint arXiv:2509.17361, 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.