

ARMR: Adaptively Responsive Network for Medication Recommendation

Feiyue Wu¹, Tianxing Wu^{1,2*}, Shenqi Jing³

¹School of Computer Science and Engineering, Southeast University, Nanjing, China

²Key Laboratory of New Generation Artificial Intelligence Technology and Its Interdisciplinary Applications (Southeast University), Ministry of Education, China

³The First Affiliated Hospital with Nanjing Medical University
(JiangSu Province Hospital), Nanjing, China
{wufeiye, tianxingwu}@seu.edu.cn, jingshenqi@jsph.org.cn

Abstract

Medication recommendation is a crucial task in healthcare, especially for patients with complex medical conditions. However, existing methods often struggle to effectively balance the reuse of historical medications with the introduction of new drugs in response to the changing patient conditions. In order to address this challenge, we propose an Adaptively Responsive network for Medication Recommendation (ARMR), a new method which incorporates 1) a piecewise temporal learning component that distinguishes between recent and distant patient history, enabling more nuanced temporal understanding, and 2) an adaptively responsive mechanism that dynamically adjusts attention to new and existing drugs based on the patient’s current health state and medication history. Experiments on the MIMIC-III and MIMIC-IV datasets indicate that ARMR has better performance compared with the state-of-the-art baselines in different evaluation metrics, which contributes to more personalized and accurate medication recommendations. The source code is publicly available at: <https://github.com/seucoin/armr2>.

1 Introduction

In the era of big data, abundant health information enables researchers and clinicians to develop sophisticated predictive methods for clinical decision support [Choi *et al.*, 2016a; Xiao *et al.*, 2018; Ali *et al.*, 2023]. Medication recommendation aims to provide personalized and effective medication plans based on a patient’s medical history, current diagnoses, and other relevant information. However, this task is particularly challenging for patients with complex conditions [Wu *et al.*, 2022; Shang *et al.*, 2019a], as their medical histories are often intricate and their responses to medications can be highly variable.

Deep learning methods have emerged as a powerful tool for modeling complex patterns in healthcare data [Zhang *et al.*, 2017; Le *et al.*, 2018; He *et al.*, 2020]. Their ability to capture intricate patterns and dependencies [Shang *et al.*, 2019b;




Figure 1: The proportion of newly introduced drugs in all drug prescriptions in the datasets. The x-axis represents the proportion range of new drugs, and the y-axis represents the proportion of prescriptions with the corresponding ratio in the dataset.

Bhoi *et al.*, 2021; Bajor and Lasko, 2022] makes them well-suited for medication recommendation tasks. Unlike traditional rule-based or instance-based methods, deep learning methods can learn from large datasets and adapt to the unique characteristics of individual patients. This adaptability is crucial for developing personalized medication plans that can evolve with the patient’s changing health status. Despite the advantages of deep learning methods, existing methods in medication recommendation face several limitations:

Neglecting the Importance of New Drug Patterns. Figure 1 shows the distribution of new and old drugs in prescribed medications. New drugs represent the drugs that has never been used before, while old drugs are the opposite. Each time the prescribed drug set contains about 30% new drugs on average. Many methods focus on modeling how to fully reuse old drugs [Wu *et al.*, 2022; Bhoi *et al.*, 2021; Yang *et al.*, 2021a; Choi *et al.*, 2016b] or recommending full medication sets [Zhang *et al.*, 2017; Shang *et al.*, 2019b], but neglect the importance of prescribing new drugs. For example, if a patient with hypertension is prescribed a new drug hydrochlorothiazide at some point, amlodipine is likely

*Corresponding author.




Figure 2: The similarity between two prescriptions for patients shows a decreasing trend with the increase of the interval between two admissions. The x-axis represents the interval between admission times for prescriptions, and the y-axis denotes the average similarity between all prescription pairs at the corresponding interval.

to be prescribed later. Learning the pattern of introduced newly prescribed drugs may help the method better understand the evolving health conditions as physicians adapt treatment plans to the patient’s changing requirements.

Insufficient Exploitation of History Records. Figure 2 shows a statistical analysis of the datasets. Generally, the larger the time lag between admission records, the greater the difference between medical code sets. Many existing methods focus primarily on recent patient data, neglecting the potential value of long-term historical information. However, distant records can still be crucial. For example, a patient’s severe health event from several years ago may be crucial for understanding their current condition, yet traditional methods may overlook this due to their focus on recent records.

To address these limitations, we propose the Adaptively Responsive Network for Medication Recommendation (ARMR). This method adjusts attention to new and old drugs through an adaptively responsive mechanism, according to the patient’s current health condition and drug usage history. Furthermore, we introduce a PTL component to effectively distinguish and process a patient’s long-term and short-term health records, thereby improving the accuracy and personalization of medication recommendations. Our major contributions are:

- We propose a medication recommendation method that combines a dynamic adaptive mechanism and piecewise temporal learning. Compared to existing methods, our method is the first to explicitly distinguish the recommendation needs for new and history drugs, and dynamically adjust the recommendation strategy based on changes in the patient’s current health state.
- We develop a unique piecewise temporal learning component which can more accurately process and analyze information in long-term health records, providing more personalized and accurate medication recommendations.

- We conduct comprehensive experiments on two public medical datasets. ARMR outperforms the state-of-the-art baselines by 2.16% improvement in Jaccard similarity and 2.55% in PRAUC, respectively.

2 Related Work

In recent years, medication recommendation has garnered significant attention due to its potential to improve patient care. Existing methods can be broadly classified into three categories based on the information they utilize, which are rule-based methods, instance-based methods, and longitudinal methods.

Rule-based methods. These methods typically rely on pre-defined clinical guidelines or attempt to derive rules from electrical health records [Almirall *et al.*, 2012; Chen *et al.*, 2016]. For instance, [Lakkaraju and Rudin, 2017] employs Markov Decision Processes to learn mappings between patient characteristics and treatments, optimizing a series of if-then-else rules. Another study [Solt and Tikk, 2009] focuses on extracting rules from discharge summaries. However, these methods often require substantial input from healthcare professionals and may lack generalizability.

Instance-based methods. These methods focus solely on current visit data to extract patient features. [Zhang *et al.*, 2017] conceptualizes medication recommendations as a multi-instance multi-label task, proposing a sequence-to-sequence method with a content-attention mechanism. [Gong *et al.*, 2021] approaches the problem as a link prediction task, jointly embedding diseases, medicines, patients, and their relationships into a shared space using knowledge graphs. However, these methods often neglect valuable historical patient data.

Longitudinal methods. This category represents popular methods that leverages sequential dependencies in patient treatment histories [Choi *et al.*, 2016b; Wu *et al.*, 2022; Yang *et al.*, 2021a; Chen *et al.*, 2023; Yang *et al.*, 2023; Liu *et al.*, 2024; Shang *et al.*, 2019b; Tan *et al.*, 2022; Zhang *et al.*, 2023; Shang *et al.*, 2019a; Bhoi *et al.*, 2023]. Many of these methods employ RNN-based architectures to capture longitudinal patient information. For example, [Choi *et al.*, 2016b] introduces a two-level neural attention network to effectively identify relevant information from patient visits. [Shang *et al.*, 2019b] incorporates a memory bank to integrate global medication interaction knowledge. MICRON [Yang *et al.*, 2021a] utilizes residual networks to predict health condition changes by analyzing differences between consecutive visits. Some studies [Chen *et al.*, 2023; Yang *et al.*, 2023] combine RNNs with graph neural networks to enhance medication representation. COGNet [Wu *et al.*, 2022] employs Transformer architecture to encode medical code sets and introduces a copy module for explicit reuse of historical medications. More recently, [Liu *et al.*, 2024] explores the potential of large language models for medication recommendation.

Our proposed method distinguishes itself from existing work by explicitly differentiating between historical and new medications and by incorporating fine-grained temporal

learning to capture evolving patient conditions. This strategy aims to provide more precise and personalized medication recommendations.

3 Preliminaries

3.1 Problem Formulation

Definition 1 (Electrical Health Records) *Electrical Health Records (EHRs) collect the medical information of patients over time. An EHR is composed of a series of medical visits denoted as $\mathcal{V}^{(i)} = \{\mathcal{V}_1^{(i)}, \mathcal{V}_2^{(i)}, \dots, \mathcal{V}_{T^{(i)}}^{(i)}\}$, where $T^{(i)}$ represents the total number of visits for the i -th patient. For notational simplicity, we omit the patient index i when the context is clear. Each visit contains three essential sets: the diagnosis set $\mathcal{D} = \{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_{|\mathcal{D}|}\}$, the procedure set $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_{|\mathcal{P}|}\}$, and the medication set $\mathcal{M} = \{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_{|\mathcal{M}|}\}$. Here, $|\cdot|$ denotes the cardinality of each set. Each visit $\mathcal{V}_t$ is then characterized by a triplet $\{\mathcal{D}_t, \mathcal{P}_t, \mathcal{M}_t\}$, where $\mathcal{D}_t \subseteq \mathcal{D}$, $\mathcal{P}_t \subseteq \mathcal{P}$, and $\mathcal{M}_t \subseteq \mathcal{M}$ represent the diagnoses, procedures, and medications associated with that visit, respectively.*

Definition 2 (Medication Recommendation) *Given a patient’s current diagnoses $\mathcal{D}_t$, procedures $\mathcal{P}_t$, and historical visit information $\{\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_{t-1}\}$, medication recommendation aims to recommend an appropriate medication combination $\mathcal{M}_t$ for the patient’s current visit.*

Notation	Description
$\mathcal{V}$	The record for a single patient
$\mathcal{D}$	All diagnosis set
$\mathcal{P}$	All procedure set
$\mathcal{M}$	All medication set
T	The length of patient’s visit records
$\mathbf{d}$	The diagnosis set of a single visit
$\mathbf{p}$	The procedure set of a single visit
$\mathbf{m}$	The medication set of a single visit

Table 1: The notations used in ARMR.

4 The Proposed Method

4.1 Overall Architecture

As shown in Figure 3, our ARMR method comprises one plugable component and three modules: (1) a hybrid **piecewise temporal learning** (PTL) component that encodes temporal medical codes set sequences by considering two stages; (2) an **adaptively responsive module** (ARM) that learns dynamic medication representation from both historical medication usage and new medication in response to changing conditions; (3) a **patient representation learning module** that learns patients’ comprehensive representations from their EHR data; (4) the patient’s and medication’s representation vectors are combined in the **joint recommendation module**, where the final drug output is obtained from element-wise thresholding.

4.2 Piecewise Temporal Learning

The piecewise temporal learning component addresses the varying impact of long-term and short-term health records on drug recommendations. It divides health records into “recent” and “distant” segments. Given an input sequence $\mathbf{S} \in \mathbb{R}^{T \times \dim}$ of length T ($T \geq N$), we split it into a recent part: $\mathbf{S}_{near} = [\mathbf{S}^{(1)}, \mathbf{S}^{(2)}, \dots, \mathbf{S}^{(N)}] \in \mathbb{R}^{N \times \dim}$ and a distant part: $\mathbf{S}_{far} = [\mathbf{S}^{(N+1)}, \mathbf{S}^{(N+2)}, \dots, \mathbf{S}^{(T)}] \in \mathbb{R}^{(T-N) \times \dim}$, and N is a hyperparameter determining the split point. For $\mathbf{S}_{near}$, we apply a feedforward neural network to capture recent health changes $\mathbf{h}_{near}$ as follows:

$$\mathbf{h}_{near} = \text{LayerNorm}(\mathbf{S}_{near} + \text{FF}_{\text{PTL}}(\mathbf{S}_{near})) \quad (1)$$

where LayerNorm [Ba et al., 2016] is layer normalization function, and FF_{PTL} is a two-layer fully connected feedforward network. For distant health changes, we need an efficient sequence modeling architecture that can effectively capture long-range dependencies while maintaining computational efficiency. Mamba [Gu and Dao, 2023], a state-of-the-art selective state space model, has shown superior performance in modeling long sequences through its input-dependent selection mechanism. While transformer-based (with quadratic-complexity) architectures are possible alternatives, Mamba offers comparable sequence modeling capabilities with significantly better computational efficiency (near linear). This makes it particularly suitable for processing longitudinal patient records with long-term patterns. Thus, we utilize the Mamba block to obtain distant health changes representation $\mathbf{h}_{far}$ as follows:

$$\mathbf{h}_{far} = \text{Mamba}(\mathbf{S}_{far}^{(1)}, \dots, \mathbf{S}_{far}^{(T-N)}) \quad (2)$$

We then combine the recent health change representation $\mathbf{h}_{near}$ and the distant health change representation $\mathbf{h}_{far}$ to get an enhanced distant health change representation $\mathbf{h}'_{far}$ using an attention mechanism as follows:

$$\mathbf{h}'_{far} = \text{LayerNorm}(\mathbf{h}_{near} + \text{Softmax}(\mathbf{h}_{near}(\mathbf{h}_{far})^T)\mathbf{h}_{far}) \quad (3)$$

where Softmax is a commonly used mathematical function to transform a vector of values into a probability distribution. This operation aims to make the component focus on relevant historical information based on recent context. Finally, we concatenate the processed representations to get an overall temporal representation $\mathbf{h}_S$ as follows:

$$\mathbf{h}_S = \text{CONCAT}(\mathbf{h}_{near}, \mathbf{h}'_{far}) \in \mathbb{R}^{2N \times \dim} \quad (4)$$

where CONCAT concatenation between vectors. We denote this entire process as $\text{PTL}(\mathbf{S}) = \mathbf{h}_S$. The PTL component integrates multiple techniques to analyze temporal data by distinguishing between recent and historical information. In this component, the feedforward network is utilized to capture immediate health status, while the Mamba block is employed to learn distant health status.

4.3 Adaptively Responsive Module

The adaptively responsive module generates dynamic medication embeddings based on the patterns of newly introduced drugs and historically reused drugs. By distinguish-




Figure 3: The overview of ARMAR's architecture. Given the input longitudinal patient EHR which comprises M , D and P , we predict the last medication set $\hat{o}^{(t)}$. **Step 1:** learning personalized medication representation E'_m , **Step 2:** learning the patient's comprehensive representation $h_{patient}$, **Step 3:** obtaining the final predicted medication set $\hat{o}^{(t)}$. The Mamba module is utilized in PTL, while the PTL component is used as a sequence temporal extractor.

ing both types of drugs, we aim to personalize recommendations to make the module adapt to both continuing treatments and newly introduced medications. Given a sequence of drug sets $\{m_1, m_2, \dots, m_{T-1}\}$ where $m_i \subseteq \mathcal{M}$, we define two set sequences, i.e., the existing drug set sequence: $\{m_1^o, m_2^o, \dots, m_{T-1}^o\}$ and the new drug set sequence: $\{m_1^n, m_2^n, \dots, m_{T-1}^n\}$ as follows:

$$m_i^o = \begin{cases} \emptyset & \text{if } i = 1 \\ m_i \cap \left(\bigcup_{k=1}^{i-1} m_k \right) & \text{if } i > 1 \end{cases} \quad (5)$$

$$m_i^n = \begin{cases} m_i & \text{if } i = 1 \\ m_i \setminus m_i^o & \text{if } i > 1 \end{cases} \quad (6)$$

where m_i^o represents the medications that were also prescribed in previous visits, while m_i^n represents newly introduced medications at each visit. We use $\cap$ to represent set intersection, and $\setminus$ to represent set difference. The separation of the existing drug set and the new drug set for each prescription aims to make the module differentiate between ongoing treatments and new interventions. We embed these sequences and apply two different PTL components PTL_m^o and PTL_m^n corresponding to m_i^o and m_i^n respectively, and then obtain time-sensitive representations h_m^o and h_m^n as follows:

$$h_m^o = PTL_m^o(\{m_i^o E_m | 1 \leq i < T\}) \quad (7)$$

$$h_m^n = PTL_m^n(\{m_i^n E_m | 1 \leq i < T\}) \quad (8)$$

where $E_m \in \mathbb{R}^{|\mathcal{D}| \times \dim}$ is the base medication embedding which is randomly initialized. The use of PTL aims to make the module capture temporal patterns in both existing and new medication sequences. We then define two medication set masks: used drugs mask $h^o \in \mathbb{R}^{|\mathcal{M}|}$ and unused drugs mask $h^N \in \mathbb{R}^{|\mathcal{M}|}$, to explicitly select the used drugs and unused drugs, respectively. These are derived from the historical medication sequence $m_h = \bigcup_{i=1}^{T-1} m_i$. h_j^o and h_j^N are defined as the j -th values of h^o and h^N respectively:

$$h_j^o = \sum_{i=1}^{T-1} (\alpha_i \cdot e^{-\beta_i \cdot (T-i)}) * \mathbf{1}\{\mathcal{M}_j \in m_h\}, 1 \leq j \leq |\mathcal{M}| \quad (9)$$

$$h_j^N = 1 - \mathbf{1}\{\mathcal{M}_j \in m_h\}, 1 \leq j \leq |\mathcal{M}| \quad (10)$$

where α_i and β_i control time decay. h_j^o leverages the occurrence time, frequency, and time decay of medications to build the representation of medication history. $\mathbf{1}\{\cdot\}$ is an indicator function which will return 1 if the given expression is true and otherwise 0. h_j^N indicates unused medications, aiming to make the module consider potential new treatments. Fi-

nally, we compute the responsive medication representation $\mathbf{E}'_m \in \mathbb{R}^{|\mathcal{D}| \times \dim}$ as follows:

$$\begin{aligned}\mathbf{E}'_m &= \mathbf{E}_m \odot (\mathbf{h}^o \odot \mathbf{q}_o + \mathbf{h}^n \odot \mathbf{q}_n) \\ \mathbf{q}_o &= \text{Softmax}\left(\frac{\mathbf{E}_m(\mathbf{h}_m^o)^T \mathbf{W}_o}{\sqrt{\dim}}\right) \\ \mathbf{q}_n &= \text{Softmax}\left(\frac{\mathbf{E}_m(\mathbf{h}_m^n)^T \mathbf{W}_n}{\sqrt{\dim}}\right)\end{aligned}\quad (11)$$

where $\mathbf{W}_o, \mathbf{W}_n \in \mathbb{R}^{2N \times \dim}$ are learned weights, and $\mathbf{q}_o$ and $\mathbf{q}_n$ combine the base medication embedding $\mathbf{E}_m$ with the time-sensitive representations $\mathbf{h}_m^o$ and $\mathbf{h}_m^n$ respectively, so that we can get a dynamic representation of each medication.

4.4 Patient Representation Learning

The patient representation learning module aims to create comprehensive representations of patients' health status based on their historical diagnosis and procedure codes. This module processes the visit records $\mathcal{V} = \{\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_T\}$ to encode both diagnosis and procedure information for each visit. At first, we utilize an embedding table $\mathbf{E}_d \in \mathbb{R}^{|\mathcal{D}| \times \dim}$, where each row represents a diagnosis code embedding. For a multi-hot diagnosis vector $\mathbf{d}^{(t)} \in \{0, 1\}^{|\mathcal{D}|}$ at the t -th visit, we compute the diagnosis representation $\mathbf{d}_e^{(t)}$ as follows:

$$\mathbf{d}_e^{(t)} = \mathbf{d}^{(t)} \mathbf{E}_d \quad (12)$$

This operation sums the embeddings of all active diagnosis codes for the visit. Similarly, we use a procedure embedding table $\mathbf{E}_p \in \mathbb{R}^{|\mathcal{P}| \times \dim}$ to encode the procedure vector $\mathbf{p}^{(t)}$ as follows:

$$\mathbf{p}_e^{(t)} = \mathbf{p}^{(t)} \mathbf{E}_p \quad (13)$$

The resulting embedding vectors $\mathbf{d}_e^{(t)}, \mathbf{p}_e^{(t)} \in \mathbb{R}^{\dim}$ encode the patient's health condition for each visit. To capture longitudinal information, we apply two PTL components PTL_d and PTL_p to the sequences of diagnosis and procedure embeddings respectively as follows:

$$\begin{aligned}\mathbf{d}_h &= \text{PTL}_d(\{\mathbf{d}_e^{(t)} | 1 \leq t \leq T\}), \\ \mathbf{p}_h &= \text{PTL}_p(\{\mathbf{p}_e^{(t)} | 1 \leq t \leq T\})\end{aligned}\quad (14)$$

where $\mathbf{d}_h, \mathbf{p}_h \in \mathbb{R}^{2N \times \dim}$. This step allows the method to learn temporal patterns in both diagnosis and procedure histories. Finally, We combine the diagnosis and procedure information to create the final patient representation as follows:

$$\mathbf{h}_{patient} = \text{CONCAT}(\mathbf{d}_h, \mathbf{p}_h) \quad (15)$$

where $\mathbf{h}_{patient} \in \mathbb{R}^{4N \times \dim}$. By integrating both the content (diagnoses and procedures) and the temporal aspects of a patient's medical history, this module creates a representation that can be used in the subsequent drug recommendation.

4.5 Joint Recommendation

At first, we employ a direct mapping from the patient representation to drug recommendations. This is achieved through

a specialized neural network FF_{o_1} , which transforms the patient encoding $\mathbf{h}_{patient}$ into a preliminary recommendation vector $\mathbf{o}_1$ as follows:

$$\mathbf{o}_1 = \text{FF}_{o_1}(\mathbf{h}_{patient}), \quad \mathbf{o}_1 \in \mathbb{R}^{|\mathcal{M}|} \quad (16)$$

Concurrently, we utilize a similarity-based method to capture nuanced relationships between patient characteristics and drug properties to obtain $\mathbf{o}_2$. This involves projecting the patient representation into a drug-compatible space using another neural network FF_{o_2} , followed by a similarity computation with drug embeddings $\mathbf{E}'_m$ as follows:

$$\mathbf{o}_2 = \text{sim}(\text{FF}_{o_2}(\mathbf{h}_{patient}), \mathbf{E}'_m), \quad \mathbf{o}_2 \in \mathbb{R}^{|\mathcal{M}|} \quad (17)$$

where $\text{sim}(\cdot, \cdot)$ represents the cosine similarity function as follows:

$$\text{sim}(\mathbf{x}, \mathbf{y}) = \frac{\mathbf{x} \cdot \mathbf{y}}{|\mathbf{x}| |\mathbf{y}|} \quad (18)$$

The final recommendation $\hat{\mathbf{o}}$ is derived by combining these two outputs $\mathbf{o}_1$ and $\mathbf{o}_2$ through a weighted sum, followed by a sigmoid activation to ensure output values are constrained between 0 and 1 as follows:

$$\hat{\mathbf{o}} = \sigma(\mathbf{w}_1 \mathbf{o}_1 + \mathbf{w}_2 \mathbf{o}_2) \quad (19)$$

where $\mathbf{w}_1$ and $\mathbf{w}_2$ are learnable parameters that determine the relative importance of each recommendation, and $\sigma(\cdot)$ denotes the sigmoid function.

4.6 Model Training and Inference

This study frames the drug recommendation task as a multi-label binary classification problem. Let $\mathbf{m}^{(t)}$ represent the ground truth drug recommendation vector, and $\hat{\mathbf{o}}^{(t)}$ is a vector denoting the method's output probabilities for each drug (as opposed to the final binary predictions $\hat{\mathbf{m}}^{(t)}$ obtained after thresholding). Our training strategy employs two complementary loss functions. First, we utilize the binary cross-entropy loss $\mathcal{L}_{bce}$, treating each drug prediction as an independent sub-problem as follows:

$$\mathcal{L}_{bce} = - \sum_{i=1}^{|\mathcal{M}|} \mathbf{m}_i^{(t)} \log(\hat{\mathbf{o}}_i^{(t)}) + (1 - \mathbf{m}_i^{(t)}) \log(1 - \hat{\mathbf{o}}_i^{(t)}) \quad (20)$$

where $\mathbf{m}_i^{(t)}$ and $\hat{\mathbf{o}}_i^{(t)}$ are the i -th entries. To enhance the robustness of our predictions, we incorporate a multi-label hinge loss $\mathcal{L}_{multi}$, which ensures that the scores for true positive labels maintain a minimum margin above those of true negative labels as follows:

$$\mathcal{L}_{multi} = \sum_{i,j: \mathbf{m}_i^{(t)=1, \mathbf{m}_j^{(t)=0}} \frac{\max(0, 1 - (\hat{\mathbf{o}}_i^{(t)} - \hat{\mathbf{o}}_j^{(t)}))}{|\mathcal{M}|} \quad (21)$$

The overall loss is computed as a weighted combination [Dosovitskiy and Djolonga, 2019] of the above two losses as follows:

$$\mathcal{L} = \alpha \mathcal{L}_{bce} + (1 - \alpha) \mathcal{L}_{multi} \quad (22)$$

where $\alpha \in [0, 1]$ controls the relative importance of each loss term, and we empirically set $\alpha = 0.7$. During training, we

Item	MIMIC-III	MIMIC-IV
# of patients	6,350	9,036
# of visits	15,032	20,616
avg. # of visit per patient	2.37	2.28
max. # of visit	29	28
# of diagnosis codes	1,958	1,892
# of procedure codes	1,430	4,939
# of medication codes	131	131

Table 2: The statistics of the processed dataset.

aggregate the loss across all visits for each patient before performing backpropagation. This patient-level optimization allows the method to capture long-term dependencies in medication histories. For inference, we follow a similar process in the training phase. The final drug recommendations are determined by applying a threshold δ to the output probabilities $\hat{o}^{(t)}$. Specifically, we recommend drugs corresponding to the entries where $\hat{o}^{(t)} > \delta$, and δ is empirically set as 0.5.

5 Experiments

This section presents our experimental results, demonstrating the efficacy of our proposed ARMR method.

5.1 Experimental Setting

Datasets. We conducted experiments using the MIMIC-III [Johnson *et al.*, 2016] and MIMIC-IV [Johnson *et al.*, 2018] datasets. Following data preprocessing procedures outlined in [Chen *et al.*, 2023], we identified 131 medications for recommendation, using ATC Third Level codes as target labels. Each ATC Third Level code encompasses one or more medications, with each medication corresponding to a single ATC Third Level code. Table 2 provides key statistics of the processed datasets.

Evaluation Metrics. According to the previous studies in medication recommendation [Yang *et al.*, 2021b; Wu *et al.*, 2022; Chen *et al.*, 2023], we employed three primary metrics: 1) Jaccard Similarity Score (Jaccard) quantifies the degree of overlap between two sets by computing the ratio of their intersection to their union; 2) F1-score (F1) is the harmonic mean of precision and recall; 3) PRAUC refers to precision recall area under curve. These metrics were calculated as the mean across all patients in the dataset.

Baselines. We compared our method with different baseline methods: 1) Traditional machine learning methods like Logistic Regression (LR); 2) Deep learning baselines including RETAIN [Choi *et al.*, 2016b], LEAP [Zhang *et al.*, 2017], GAMENet [Shang *et al.*, 2019b], LEADER [Liu *et al.*, 2024], Carmen [Chen *et al.*, 2023], MoleRec [Yang *et al.*, 2023] and COGNet [Wu *et al.*, 2022]. Additionally, we evaluated several variants of our ARMR method which are detailed in the section of ablation study.

5.2 Performance Comparison

Table 3 presents a comprehensive comparison of all baseline methods and ARMR variants across both MIMIC-III and MIMIC-IV datasets.

Overall Performance. Our ARMR method consistently outperforms all baseline methods in the metrics of Jaccard, F1 and PRAUC. The simpler methods (LR, ECC, LEAP) show limited effectiveness, likely due to their focus on current visit data without considering longitudinal patient history. More sophisticated methods like RETAIN, GAMENet, MoleRec, Carmen and COGNet demonstrate improved performance through various methods of incorporating historical patient data. However, these methods still fall short in certain aspects. RETAIN relies solely on bidirectional RNN for historical encoding. While GAMENet introduces graph-based information but lacks in temporal dynamics and MoleRec incorporates drug molecule structures, enhancing performance but missing broader contextual factors. COGNet employs positional encoding self-attention, recognizing the potential for medication reuse, but overlooks the dynamic nature of health conditions. In contrast, ARMR’s superior performance can be attributed to its comprehensive strategy in integrating historical data with current health dynamics.

5.3 Ablation Study

To evaluate the contribution of each component in our proposed ARMR method, we conducted a series of ablation experiments. We designed the following variant methods:

- ARMR *w/o PTL*: We used standard RNN instead of PTL.
- ARMR *w/o PTL(L)*: We maintained the PTL module but removes the distant history processing part, which means simply changing Eq.(2) to an empty sequence.
- ARMR *w/o PTL(N)*: We maintained the PTL module but remove the recent history processing part by changing the input of Eq.(2) to the whole sequence and removing recent history processing part.
- ARMR *w/o ARM*: We removed the ARM module by using PTL on the $\mathbf{m}_i$ sequence directly and get the final medication representations similarly to Eq.(11)

The superior performance of ARMR over ARMR *w/o PTL* and ARMR *w/o ARM* indicates that the PTL component and ARM module bring significant improvement to the basic method without PTL and ARM. The PTL module creates a more accurate and context-aware representation of a patient’s health history, potentially leading to more informed drug recommendations. The ARM module dynamically adjusts drug recommendations based on the patient’s medication history and current health condition. By distinguishing between existing and new medications, and considering their temporal patterns, the module can provide a more personalized and context-aware recommendation strategy. The difference between ARMR *w/o PTL* and ARMR *w/o PTL(L)* is that they are using different base sequence learning architectures, i.e., standard RNN and Mamba. The superior performance of the latter highlights the advantages of using Mamba for medical sequence modeling, particularly in capturing long-range dependencies while maintaining computational efficiency. Both ARMR *w/o PTL(N)* and ARMR *w/o PTL(L)* show suboptimal performance, emphasizing the importance of our temporal segmentation strategy in the PTL component. This suggests that balancing the influence of recent and distant history is

Method	MIMIC-III			MIMIC-IV		
	Jaccard (%)	PRAUC (%)	F1 (%)	Jaccard (%)	PRAUC (%)	F1 (%)
LR	49.40±0.14	75.88±0.19	65.08±0.11	47.32±0.18	73.80±0.16	63.03±0.17
RETAIN	48.55±0.15	75.68±0.12	64.67±0.10	44.03±0.10	71.40±0.11	60.13±0.08
LEAP	45.44±0.18	65.71±0.12	60.57±0.12	43.50±0.10	71.61±0.40	60.13±0.23
GAMENet	51.59±0.09	76.11±0.07	66.74±0.10	48.84±0.08	73.95±0.05	64.41±0.17
LEADER	51.75±0.22	77.95±0.25	67.37±0.19	47.84±0.08	74.95±0.05	63.32±0.17
Carmen	52.67±0.21	76.52±0.36	68.12±0.19	50.06±0.12	74.62±0.30	65.69±0.07
MoleRec	53.05±0.33	77.36±0.27	68.43±0.29	49.92±0.21	74.73±0.31	65.45±0.19
COGNet	53.36±0.11	77.39±0.09	68.69±0.10	50.23±0.12	75.21±0.21	65.82±0.14
ARMR	54.51±0.13	79.63±0.28	69.59±0.29	51.74±0.14	76.95±0.21	67.12±0.18
ARMR w/o PTL	52.45±0.12	77.49±0.09	66.88±0.24	49.82±0.21	74.51±0.34	65.47±0.23
ARMR w/o PTL(L)	54.01±0.08	78.99±0.18	69.32±0.16	51.32±0.17	76.46±0.21	66.58±0.15
ARMR w/o PTL(N)	53.55±0.17	78.49±0.15	68.72±0.14	50.51±0.17	75.35±0.33	65.98±0.09
ARMR w/o ARM	53.21±0.14	78.28±0.34	68.69±0.30	50.22±0.10	75.20±0.11	65.66±0.14

Table 3: The comparison results on MIMIC-III and MIMIC-IV datasets.

critical for accurate medication recommendation. In general, the complete ARMR outperforms all variants removing some part, confirming that each key design contributes synergistically to the overall performance.

5.4 Case Study

To demonstrate the efficacy of ARMR in balancing the continuation of previous treatments and the introduction of new medications, we conducted an in-depth analysis of a representative patient from the MIMIC-IV dataset. This individual’s medical history comprises five distinct hospital visits, which are characterized by the evolving diagnoses and corresponding adjustments to their medication prescriptions.




Figure 4: An example of case study. The figure shows four methods’ predicted medications for a patient with 5 visits. The y-axis shows the number of recommend drugs for each method per visit, which are divided into historical (upper) and new (lower) drugs split by the $y=0$ line. The bright color means true positives, and the dark color means false positives.

Figure 4 presents the medication recommendation results for the same patient generated by our method and three state-of-the-art baselines (i.e., COGNet, MoleRec and Carmen)

which are good in general. A comprehensive evaluation of the results reveals that ARMR consistently outperforms its counterparts in terms of overall accuracy of medication recommendations. Notably, ARMR recommends more correct historical drugs and relatively less unseen drugs, exhibiting comparable proficiency to COGNet in identifying appropriate medications for continued use, but COGNet tends to recommend more unseen drugs. Besides, ARMR predicts the most correct new drugs in all visits, which truly distinguishes itself for its superior ability to suggest new, previously unprescribed medications that align with the patient’s evolving medical needs. These findings highlight the efficacy of our proposed ARM, which effectively balances continuity in treatment while addresses emergent patient needs.

6 Conclusion

In this paper, we propose a new medication recommendation method ARMR which leverages a piecewise temporal learning component to extract changing patient conditions, and utilizes an adaptively responsive module to determine how to use existing and new medications. Experiment results on the publicly available MIMIC-III and MIMIC-IV dataset demonstrate that ARMR outperforms existing baselines. In the future, we plan to 1) design quantified experiments to evaluate the specific ability of all methods on introducing correct new drugs, and 2) enhance the explainability of recommendations to further increase clinicians’ trust to apply ARMR in real-world scenarios.

Acknowledgements

This work is supported by the NSFC (Grant No. 62376058), the National Key Research and Development Program (Grant No. 2022YEB2703300, 2022YFB2703301), the Key R&D Program of Jiangsu (Grant No. BE2022798), the Southeast University Interdisciplinary Research Program for Young Scholars, and the Big Data Computing Center of Southeast University.

References

- [Ali *et al.*, 2023] Zafar Ali, Yi Huang, Irfan Ullah, Junlan Feng, Chao Deng, Nimbeshaho Thierry, Asad Khan, Asim Ullah Jan, Xiaoli Shen, Wu Rui, et al. Deep Learning for Medication Recommendation: A Systematic Survey. *Data Intelligence*, 5(2):303–354, 2023.
- [Almirall *et al.*, 2012] Daniel Almirall, Scott N Compton, Meredith Gunlicks-Stoessel, Naihua Duan, and Susan A Murphy. Designing a pilot sequential multiple assignment randomized trial for developing an adaptive treatment strategy. *Statistics in Medicine*, 31(17):1887–1902, 2012.
- [Ba *et al.*, 2016] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [Bajor and Lasko, 2022] Jacek M Bajor and Thomas A Lasko. Predicting medications from diagnostic codes with recurrent neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.
- [Bhoi *et al.*, 2021] Suman Bhoi, Mong Li Lee, Wynne Hsu, Hao Sen Andrew Fang, and Ngiap Chuan Tan. Personalizing medication recommendation with a graph-based approach. *ACM Transactions on Information Systems (TOIS)*, 40(3):1–23, 2021.
- [Bhoi *et al.*, 2023] Suman Bhoi, Mong Li Lee, Wynne Hsu, and Ngiap Chuan Tan. Refine: A fine-grained medication recommendation system using deep learning and personalized drug interaction modeling. 36:24013–24024, 2023.
- [Chen *et al.*, 2016] Zhuo Chen, Kyle Marple, Elmer Salazar, Gopal Gupta, and Lakshman Tamil. A physician advisory system for chronic heart failure management based on knowledge patterns. *Theory and Practice of Logic Programming*, 16(5-6):604–618, 2016.
- [Chen *et al.*, 2023] Qianyu Chen, Xin Li, Kunan Geng, and Mingzhong Wang. Context-aware safe medication recommendations with molecular graph and ddi graph embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 7053–7060, 2023.
- [Choi *et al.*, 2016a] Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F Stewart, and Jimeng Sun. Doctor ai: Predicting clinical events via recurrent neural networks. In *Machine Learning for Healthcare Conference*, pages 301–318, 2016.
- [Choi *et al.*, 2016b] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, 2016.
- [Dosovitskiy and Djolonga, 2019] Alexey Dosovitskiy and Josip Djolonga. You only train once: Loss-conditional training of deep networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [Gong *et al.*, 2021] Fan Gong, Meng Wang, Haofen Wang, Sen Wang, and Mengyue Liu. Smr: medical knowledge graph embedding for safe medicine recommendation. *Big Data Research*, 23:100174, 2021.
- [Gu and Dao, 2023] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.
- [He *et al.*, 2020] Yong He, Cheng Wang, Nan Li, and Zhenyu Zeng. Attention and memory-augmented networks for dual-view sequential learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 125–134, 2020.
- [Johnson *et al.*, 2016] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1):1–9, 2016.
- [Johnson *et al.*, 2018] Alistair EW Johnson, David J Stone, Leo A Celi, and Tom J Pollard. The mimic code repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association*, 25(1):32–39, 2018.
- [Lakkaraju and Rudin, 2017] Himabindu Lakkaraju and Cynthia Rudin. Learning cost-effective and interpretable treatment regimes. In *Proceedings of the Artificial Intelligence and Statistics (AISTATS)*, pages 166–175. PMLR, 2017.
- [Le *et al.*, 2018] Hung Le, Truyen Tran, and Svetha Venkatesh. Dual memory neural computer for asynchronous two-view sequential learning. In *Proceedings of the 24th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1637–1645, 2018.
- [Liu *et al.*, 2024] Qidong Liu, Xian Wu, Xiangyu Zhao, Yuanshao Zhu, Zijian Zhang, Feng Tian, and Yefeng Zheng. Large language model distilling medication recommendation model. *arXiv preprint arXiv:2402.02803*, 2024.
- [Shang *et al.*, 2019a] Junyuan Shang, Tengfei Ma, Cao Xiao, and Jimeng Sun. Pre-training of graph augmented transformers for medication recommendation. pages 5953–5959, 7 2019.
- [Shang *et al.*, 2019b] Junyuan Shang, Cao Xiao, Tengfei Ma, Hongyan Li, and Jimeng Sun. Gamenet: Graph augmented memory networks for recommending medication combination. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, pages 1126–1133, 2019.
- [Solt and Tikk, 2009] Illés Solt and Domonkos Tikk. Yet another rule-based approach for extracting medication information from discharge summaries. In *Proceedings of the Third I2B2 Workshop on Challenges in Natural Language Processing for Clinical Data*, 2009.

- [Tan *et al.*, 2022] Yanchao Tan, Chengjun Kong, Leisheng Yu, Pan Li, Chaochao Chen, Xiaolin Zheng, Vicki S Hertzberg, and Carl Yang. 4sdrug: Symptom-based set-to-set small and safe drug recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 3970–3980, 2022.
- [Wu *et al.*, 2022] Rui Wu, Zhaopeng Qiu, Jiacheng Jiang, Guilin Qi, and Xian Wu. Conditional generation net for medication recommendation. In *Proceedings of the ACM Web Conference (WWW)*, pages 935–945, 2022.
- [Xiao *et al.*, 2018] Cao Xiao, Edward Choi, and Jimeng Sun. Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review. *Journal of the American Medical Informatics Association*, 25(10):1419–1428, 2018.
- [Yang *et al.*, 2021a] Chaoqi Yang, Cao Xiao, Lucas Glass, and Jimeng Sun. Change matters: Medication change prediction with recurrent residual networks. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3728–3734, 2021.
- [Yang *et al.*, 2021b] Chaoqi Yang, Cao Xiao, Fenglong Ma, Lucas Glass, and Jimeng Sun. Safedrug: Dual molecular graph encoders for recommending effective and safe drug combinations. pages 3735–3741, 8 2021.
- [Yang *et al.*, 2023] Nianzu Yang, Kaipeng Zeng, Qitian Wu, and Junchi Yan. Molerec: Combinatorial drug recommendation with substructure-aware molecular representation learning. In *Proceedings of the ACM Web Conference (WWW)*, pages 4075–4085, 2023.
- [Zhang *et al.*, 2017] Yutao Zhang, Ruqi Chen, Jie Tang, Walter F Stewart, and Jimeng Sun. Leap: learning to prescribe effective and safe treatment combinations for multimorbidity. In *Proceedings of the 23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1315–1324, 2017.
- [Zhang *et al.*, 2023] Yingying Zhang, Xian Wu, Quan Fang, Shengsheng Qian, and Changsheng Xu. Knowledge-enhanced attributed multi-task learning for medicine recommendation. *ACM Transactions on Information Systems (TOIS)*, 41(1):1–24, 2023.