

VITA: ‘Carefully Chosen and Weighted Less’ Is Better in Medication Recommendation

Taeri Kim^{1*}, Jiho Heo^{1*}, Hongil Kim^{2*}, Kijung Shin^{3†}, Sang-Wook Kim^{1†}

¹Department of Computer Science, Hanyang University, South Korea

²Department of Artificial Intelligence, Hanyang University, South Korea

³Kim Jaechul Graduate School of AI & School of Electrical Engineering, KAIST, South Korea
{taerik, linda0123, hong0814, wook}@hanyang.ac.kr, kijungs@kaist.ac.kr

Abstract

We address the medication recommendation problem, which aims to recommend effective medications for a patient’s current visit by utilizing information (*e.g.*, diagnoses and procedures) given at the patient’s current and past visits. While there exist a number of recommender systems designed for this problem, we point out that they are challenged in accurately capturing the relation (*spec.*, the degree of relevance) between the current and each of the past visits for the patient when obtaining her *current health status*, which is the basis for recommending medications. To address this limitation, we propose a novel medication recommendation framework, named **VITA**, based on the following two novel ideas: (1) relevant-Visit selectIon; (2) Target-aware Attention. Through extensive experiments using real-world datasets, we demonstrate the superiority of VITA (*spec.*, up to 5.67% higher accuracy, in terms of Jaccard, than the best competitor) and the effectiveness of its two core ideas. The code is available at <https://github.com/jhheo0123/VITA>.

1 Introduction

Medication recommendation aims to help doctors in prescribing effective medications for a patient’s current visit (Zhang et al. 2017; Shang et al. 2019b). When prescribing medications to a patient, doctors should consider the following factors: (1) the diagnoses and procedures given to her at the current visit; (2) her past health records; (3) the relations between medications to be prescribed (*e.g.*, the possibility of adverse effects when taken together). This is a time-consuming and challenging process even for experienced doctors (Wu et al. 2022b). Therefore, the medication recommendation that can alleviate these difficulties of doctors has been an important research area in medical tasks.

Early medication recommender systems recommend medications to a patient by considering *only* the diagnoses and procedures given to her at the current visit (*i.e.*, *current visit information*), which are called *instance-based* methods (Zhang et al. 2017; Wang et al. 2018; Gong et al. 2021). However, they overlooked the fact that, although patients are

given with the same diagnoses at the current visit, the cause of the given diagnoses may be different for each patient, so the medications for them could also be different (Shang et al. 2019b; Yang et al. 2021b); this is because they do not utilize the patients’ past health records (Wu et al. 2022b).

To alleviate this limitation, *longitudinal-based* medication recommender systems have emerged, which utilize not only a patient’s *current* visit information but also her *past* health records (Shang et al. 2019a,b; Bhoi et al. 2021; Wang et al. 2021a,b; Yang et al. 2021b; Wu et al. 2022a,b). These methods usually consist of an *encoder* and a *predictor*. For their encoders, most of them (Shang et al. 2019b; Bhoi et al. 2021; Wang et al. 2021a,b; Yang et al. 2021b; Wu et al. 2022a) learn a *patient representation* indicative of her *current health status* by aggregating her current visit information and the diagnoses and procedures given at her past visits (*i.e.*, *past visit information*) in consideration of the visit order via a Recurrent Neural Network (RNN)-based model. Then, for their predictors, they recommend medications to her by utilizing other past health records (*e.g.*, medications prescribed at her past visits) based on their relevance with her patient representation via various deep-learning models (*e.g.*, an attention network (Bahdanau, Cho, and Bengio 2015)). On the other hand, COGNet (Wu et al. 2022b), another longitudinal-based method recently proposed, encodes *only* a patient’s *current* visit information, unlike the above methods, to obtain her current health status in the encoder by capturing the relation between diagnoses (resp. procedures) at her current visit via a transformer (Vaswani et al. 2017)-based model. Then, in the predictor, COGNet additionally takes into account the similarity between (a) medications that should be prescribed at her current visit by considering only her current visit information and (b) those prescribed at her past visits, to improve upon the predictor of the aforementioned methods. Please refer to the Appendix¹ for more detailed information on these (*i.e.*, related studies).

Although longitudinal-based methods achieve higher accuracy than instance-based methods by utilizing a patient’s past health records, we point out a key limitation of them: when *obtaining the patient’s current health status* (*i.e.*, in

*These authors contributed equally as co-first authors.

†Co-corresponding authors.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The Appendix for VITA can be found at <https://github.com/jhheo0123/VITA>.

the *encoder*), they have a difficulty in accurately capturing the relation (*spec.*, the degree of relevance) of each past visit to the current one. Specifically, they either compute the relevance score by considering simply the *order of visits* via a RNN-based model (Shang et al. 2019b; Bhoi et al. 2021; Wang et al. 2021a,b; Yang et al. 2021b; Wu et al. 2022a), or do not attempt to capture the degree of relevance, not taking into account *any* of the patient’s past visits (Wu et al. 2022b). However, when obtaining the patient’s current health status, accurately capturing the degree of relevance between her current and past visits is highly important. In particular, her past visit information relevant (resp. irrelevant) to her current visit should (resp. should not) be reflected in representing her current health status. Our empirical findings supporting this claim are detailed in Section 2.

In this paper, we propose VITA (relevant-visit selection and target-aware attention), a novel medication recommendation framework, based on a more accurate understanding of her current health status. To achieve this, we learn a patient representation indicative of her current health status by taking into account an accurate relevance score between the current visit and each of the past ones, without relying solely on the visit order. Furthermore, we do *not* consider her past visits *irrelevant* to her current one *at all*, instead of simply assigning lower relevance scores to them, since this approach enhances accurate medication recommendation. Our empirical validation of these ideas can be found in Section 4.

Our contributions are summarized as follows:

- **Important Discovery:** We discovered that existing medication recommender systems face challenges in accurately capturing the degree of relevance of each past visit to the current one when obtaining a patient’s current health status. In addition, we are the first to demonstrate that using past visits irrelevant to the current one has a negative effect on recommending accurate medications.
- **Novel Framework:** To address this key limitation, we propose a novel medication recommendation framework, named **VITA**, that recommends medications by employing an *enhanced patient representation*, based on the following two novel ideas: (1) **relevant-Visit selection** that automatically excludes past visits irrelevant to the current one; (2) **Target-aware Attention** that accurately captures the relevance score of the past visits to the current one.
- **Extensive Evaluation:** We validate the effectiveness of VITA through extensive experiments using public real-world datasets. Most importantly, VITA surpasses *all* six state-of-the-art competitors, achieving an improvement of up to 5.67% in terms of Jaccard. Also, the two core ideas of VITA can be orthogonally combined with existing medication recommender systems; when combined, they elevate the accuracy beyond that of the original ones.

2 Motivation

In this section, we demonstrate the limitations of existing medication recommender systems via preliminary experiments answering the following preliminary questions (PQs):




Figure 1: Accuracies of GAMENet when varying the use of past visit information of a patient. ‘All’ (resp. ‘No’) means to use all (resp. not to use any) past visit information. ‘Top-1/Mid.-1/Bot.-1’ means to use only one past visit information that is the most/moderately/the least similar to the current visit information, respectively. All differences are statistically significant with a p -value ≤ 0.001 .

- **PQ1:** Does using past visit information in representing a patient’s current health status help accurate medication recommendation?
- **PQ2:** Which past visit information is beneficial to better representing a patient’s current health status?

Experimental Settings. We conducted experiments using the GAMENet model (Shang et al. 2019b)², the MIMIC-III dataset (Johnson et al. 2016), and the accuracy measure Jaccard, which are most commonly used in medication recommendation (Bhoi et al. 2021; Wang et al. 2021a,b; Yang et al. 2021a,b; Wu et al. 2022b). Please refer to Section 4 for more detailed information on experimental settings.

PQ1: Effect of the use of a patient’s past visit information. We first analyze whether it is really helpful to use a patient’s past visit information for representing her current health status (*i.e.*, in her patient representation) for recommending medications at the current visit. To this end, we compared the accuracy of GAMENet that uses *all* past and current visit information as input to the RNN-based model in the encoder (*i.e.*, the original method) with that of its counterpart that uses *only* her *current* visit information.

The results for the above two methods are represented as ‘All’ and ‘No’, respectively, in Figure 1-(PQ1). From the results, we found that using not only the current visit information but also the past ones for patient representation is helpful for accurate medication recommendation.

PQ2: Effect of the use of a patient’s past visit information depending on its degree of relevance to her current visit. Now, we analyze which past visit information of a patient is helpful to obtain better patient representation in medication recommendation. To this end, we first calculated the Jaccard similarity between the current and each of past visit information for each patient. Specifically, we first represented the diagnoses (resp. procedures) given at each visit as a multi-hot vector for each visit per patient. Then, we represented each visit as a single vector by concatenating the two multi-hot vectors of the diagnoses and procedures to calculate the

²The results on other existing medication recommender systems showed similar tendencies to those on the GAMENet. These results would be shown in the Appendix.




Figure 2: Overview of VITA, composed of two components: an encoder based on relevant-visit selection and target-aware attention; a predictor based on (a) current health-aware and (b) current health-relevant past medication representations.

similarity. After that, for each patient, we measured the accuracy of GAMENet using *only* the past visit information *most similar* to the current visit, along with the current one, as an input to the RNN-based model in the encoder. The result of this method is represented as ‘Top-1’ in Figure 1-(PQ2). Also, for comparison, we additionally measured the accuracy using moderately (resp. the least) similar one, *instead of* using the most similar past visit information (the moderately similar past visit is the one with the median similarity). The result for this method is represented as ‘Mid.-1’ (resp. ‘Bot.-1’) in Figure 1-(PQ2).

As shown in Figure 1-(PQ2), the accuracy decreases in the order of using only the most similar, moderately similar, and the least similar past visit information along with the current one. In addition, we found that the accuracy when using *only* the most similar (resp. the least similar) past visit information along with the current one is rather *higher* (resp. *lower*) than that of the original method using *all* past visit information along with the current one (resp. the method using *only* the current visit information) (compare ‘All’ (resp. ‘No’) and ‘Top-1’ (resp. ‘Bot.-1’) accuracies in Figure 1-(PQ1) and - (PQ2), respectively).

Summary. Based on the above results, we draw the following conclusions: (1) referencing past visit information in obtaining the patient representation is *helpful* for medication recommendation; furthermore, (2) using *only* relevant (*i.e.*, similar) past visit information to the current one is *more* helpful; however, (3) using irrelevant (*i.e.*, dissimilar) past visit information to the current one has a *negative* effect.

3 VITA: Proposed Framework

In this section, we detail our proposed framework, VITA, based on relevant-visit selection and target-aware attention.

3.1 Problem Definition

Definition 1: Patient Health Records. In Electronic Health Records (EHR) data (*e.g.*, MIMIC-III dataset), the health records of each patient x consist of sequential visits $\mathcal{V}_x = [\mathbf{V}_x^1, \dots, \mathbf{V}_x^{(T-1)}, \mathbf{V}_x^T]$, where $\mathbf{V}_x^{(T-1)}$ denotes the $(T-1)$ -th visit of patient x . For simplicity, following (Zhang et al. 2017; Shang et al. 2019a,b; Wu et al. 2022b), we omit the subscript indicating a patient (*i.e.*, x) and describe VITA

with a single patient. Therefore, the patient’s visits are denoted by $\mathcal{V} = [\mathbf{V}^1, \dots, \mathbf{V}^{(T-1)}, \mathbf{V}^T]$. Each visit $\mathbf{V}^{(T-1)}$ of a patient \mathcal{P} consists of three subsets of all diagnoses \mathcal{D} , all procedures \mathcal{P} , and all medications \mathcal{M} ; the subsets are represented as multi-hot vectors (*e.g.*, $\mathbf{d}^{(T-1)} \in \mathbb{R}^{|\mathcal{D}|}$, $\mathbf{p}^{(T-1)} \in \mathbb{R}^{|\mathcal{P}|}$, and $\mathbf{m}^{(T-1)} \in \mathbb{R}^{|\mathcal{M}|}$).

Definition 2: EHR and DDI Graphs. The EHR and the Drug-Drug Interactions (DDI) graphs are denoted by $\mathcal{G}_{EHR} = (\mathcal{M}, \mathcal{E}_{EHR})$ and $\mathcal{G}_{DDI} = (\mathcal{M}, \mathcal{E}_{DDI})$, respectively, where \mathcal{E}_{EHR} denotes the set of edges between medications, each of which indicates that two medications have been prescribed together at any visit of any patient, and \mathcal{E}_{DDI} denotes the set of edges between medications, each of which indicates that two medications may cause adverse effects if taken together; their adjacency matrices \mathbf{A}_{EHR} , $\mathbf{A}_{DDI} \in \mathbb{R}^{|\mathcal{M}| \times |\mathcal{M}|}$ satisfy: $\mathbf{A}_{EHR}[i, j] = 1$ if and only if the medications i and j have been prescribed together at any visit of any patient, $\mathbf{A}_{DDI}[i, j] = 1$ if and only if the medication i and j can be harmful when they are taken together. The same EHR and DDI graphs are used for all patients.

Medication Recommendation Problem. Given past health records $[\mathbf{V}^1, \dots, \mathbf{V}^{(T-1)}]$ and current visit information (*i.e.*, diagnoses \mathbf{d}^T and procedures \mathbf{p}^T given at the current visit T) of a patient, and EHR and DDI graphs, the goal is to recommend the medications $\hat{\mathcal{M}}^T$ for her current visit T . Key notations used in this paper can be found in the Appendix.

3.2 Key Components in VITA

VITA consists of two components, an encoder and a predictor, and its schematic overview is presented in Figure 2. In the following, we delve into the details of each component.

Encoder. In its encoder, VITA aims to obtain an enhanced patient representation $\mathbf{q}^T \in \mathbb{R}^{dim}$, which denotes her current health status (*i.e.*, at the T -th visit), by employing relevant-visit selection and target-aware attention. To achieve this, for each visit, VITA first concatenates visit t ’s diagnoses \mathbf{d}^t and procedures \mathbf{p}^t , and then feeds it into an embedding layer to obtain a dense representation \mathbf{v}^t of visit t , as follows: $\forall t \in \{1, \dots, (T-1), T\}$,

$$\mathbf{v}^t = \text{concat}(\mathbf{d}^t, \mathbf{p}^t) \mathbf{W}_e, \quad (1)$$

where $\mathbf{W}_e \in \mathbb{R}^{(|\mathcal{D}|+|\mathcal{P}|) \times dim}$ denotes the learnable weight matrix of the embedding layer.

Relevant-Visit Selection. Then, VITA aggregates the dense representations \mathbf{v}^t for all visits to obtain her patient representation \mathbf{q}^T . Recall that, in Section 2, we showed that using *only* the past visit information *relevant* to the current one is helpful for accurate medication recommendation. Therefore, we design a novel *relevant-visit selection* module to find only the past visits relevant to the current visit. However, here we encounter the following challenge: the number of past visits relevant to the current visit may *differ per patient*; for some patients, it is possible that *no* past visit is relevant to the current visit. To address this challenge, we predict per patient whether each of the past visits is relevant to her current one or not (*i.e.*, whether to select or not).

Specifically, VITA obtains a set \mathcal{V}^{rel} of past visits relevant to the current visit via the relevant-visit selection module. To this end, for each visit, VITA first concatenates the dense representations \mathbf{v}^t of the past visit t and \mathbf{v}^T of the current visit T to provide the relevant-visit selection module with the context for deciding whether to select the past visit t , and then feeds them into a Multi-Layer Perceptron (MLP) layer to obtain the probability s^t of the past visit t being selected, as follows: $\forall t \in \{1, \dots, (T-1)\}$,

$$s^t = \text{sigmoid}(\text{concat}(\mathbf{v}^t, \mathbf{v}^T)\mathbf{W}_s + b_s), \quad (2)$$

where $\mathbf{W}_s \in \mathbb{R}^{2dim \times 1}$ and $b_s \in \mathbb{R}$ denote a weight matrix and a bias value of the MLP layer, respectively. Finally, VITA employs the Gumbel-softmax (Jang, Gu, and Poole 2017), which transforms the discrete sampling problem into a differentiable continuous problem, thereby allowing for the flow of gradients, to select the past visits relevant to the current visit (*i.e.*, to obtain the set \mathcal{V}^{rel}), as follows:

$$\mathcal{V}^{rel} = \{\mathbf{v}^t : t \in \{1, \dots, (T-1)\} \text{ and } [\sigma_1^t + 0.5] = 1\}, \quad (3)$$

where $\sigma_\gamma^t = \frac{\exp((\log \pi_\gamma^t + z_\gamma^t)/\tau_g)}{\sum_{\delta=1}^2 \exp((\log \pi_\delta^t + z_\delta^t)/\tau_g)}$, and $\pi^t \in \mathbb{R}^{2 \times 1}$ denotes a vector obtained by concatenating the probability s^t and the probability $(1 - s^t)$ of each past visit t not being selected; z_γ^t and τ_g denote the γ -th sample from the $Gumbel(0, 1)$ distribution for the past visit t and the temperature hyperparameter, respectively.

Target-Aware Attention. Given the set \mathcal{V}^{rel} , VITA aggregates the dense representations \mathbf{v}^t in \mathcal{V}^{rel} and \mathbf{v}^T of the current visit T to obtain the ultimate patient representation \mathbf{q}^T . Here, to the best of our knowledge, we are first to employ an attention network in the *encoder* of medication recommender systems; this aims to accurately capture the relevance score α^t between the current visit T and each relevant past visit t , without relying solely on the visit order³. Moreover, to obtain an accurate relevance score α^t even when *all* the past visits are *weakly* relevant to the current visit, we design a novel *target-aware attention* module that assigns the relevance score α^t of each past visit t relatively not only to the other past visit but also to the current visit.

Specifically, using the dense representation \mathbf{v}^T of the current visit T as a query and all dense representations $\mathbf{v}^t \in \mathcal{V}^{rel}$ for the relevant past visits t and \mathbf{v}^T of the current visit

³Nevertheless, we experimented with a variant of VITA that incorporates positional encoding into our target-aware attention, but the current version of VITA showed higher accuracy.

T as keys and values, VITA computes the relevance score α^t between the current visit and each of relevant past visits, as follows: $\forall \mathbf{v}^t \in \mathcal{V}^{rel} \cup \{\mathbf{v}^T\}$,

$$\alpha^t = \frac{\exp((\mathbf{v}^T \mathbf{W}_\alpha \mathbf{v}^{t\top})/\sqrt{dim})}{\sum_{\mathbf{v}^f \in \mathcal{V}^{rel} \cup \{\mathbf{v}^T\}} \exp((\mathbf{v}^T \mathbf{W}_\alpha \mathbf{v}^{f\top})/\sqrt{dim})}, \quad (4)$$

where $\mathbf{W}_\alpha \in \mathbb{R}^{dim \times dim}$ denotes a learnable weight matrix of the target-aware attention module. It is worth noting that the relevance score α^T for the current visit T is also calculated; thus, the sum of the softmax values for all selected past visits (except for the current visit) is not necessarily fixed to 1; it is also possible for the sum to be much smaller than 1 if *all* selected visits are *weakly relevant* to the current visit.

Finally, VITA obtains the patient representation \mathbf{q}^T by aggregating the dense representation \mathbf{v}^T of the current visit T and all dense representations $\mathbf{v}^t \in \mathcal{V}^{rel}$ of the relevant past visits t based on their relevance scores α^t , as follows:

$$\mathbf{q}^T = \sum_{\mathbf{v}^t \in \mathcal{V}^{rel} \cup \{\mathbf{v}^T\}} \alpha^t \mathbf{v}^t. \quad (5)$$

We highlight that VITA’s two core ideas (*i.e.*, relevant-visit selection and target-aware attention modules) can be orthogonally combined with the encoder of *any* medication recommender systems, thereby improving their accuracy; this claim will be empirically validated in Section 4.

Predictor. In its predictor, VITA aims to recommend the set $\hat{\mathcal{M}}^T$ of necessary medications for a patient based on her patient representation \mathbf{q}^T obtained. To this end, VITA first obtains, using the patient representation \mathbf{q}^T , the two medication representations (a) $\mathbf{p}_k^T \in \mathbb{R}^{dim}$, which captures features of the k -th medication that are necessary at her current visit T when considering her current health status, and (b) $\bar{\mathbf{p}}_k^T \in \mathbb{R}^{|\mathcal{M}|}$, which captures the probabilities that all medications prescribed at her *past* visits will be recommended as the k -th medication at her current visit T when considering her current health status. Then, VITA fuses them to obtain the k -th medication of the necessary medications for her current visit. The above process is iteratively performed until all necessary medications for her current visit are obtained.⁴

(a) Current Health-Aware Medication Representation. To obtain a current health-aware medication representation \mathbf{p}_k^T , VITA first obtains enriched medication representations \mathbf{e}_i via the relations between medications, by applying an independent two-layer Graph Convolutional Network (GCN) to each of the EHR and DDI graphs with randomly initialized medication representations \mathbf{e}_i , and then fusing outputs of the two GCNs per medication (Shang et al. 2019b).

Subsequently, VITA begins to obtain the current health-aware medication representation \mathbf{p}_k^T using the patient representation \mathbf{q}^T and the medication representations \mathbf{e}_i . In this process, especially, VITA focuses on the medications already predicted (*i.e.*, medications up to the $(k-1)$ -th⁵) rather

⁴In practice, VITA predicts the <END> token (class), indicating that the medication prediction at the patient’s current visit is completed. Also, we compared the accuracies of the one-by-one approach and the all-at-once approach on our framework. The results showed that the one-by-one approach was more accurate.

⁵ $k \geq 1$, when $k = 1$, a randomly initialized <START> token (representation) is used.

than all medications, since this approach helps ensure that the relations between recommended medications are considered when recommending the next one. This process can be formally expressed as follows:

$$\mathbf{p}_k^T = \text{softmax}\left(\frac{\text{transformer}(\mathbf{e}_*^{T,1}, \dots, \mathbf{e}_*^{T,(k-1)}) \odot \mathbf{q}^T}{\sqrt{\text{dim}}}\right) \mathbf{q}^T, \quad (6)$$

where $\text{transformer}(\cdot)$ denotes a transformer-based model as used in (Wu et al. 2022b), which aggregates inputs considering the relation between inputs; $\mathbf{e}_*^{T,(k-1)}$ denotes a medication representation of the $(k-1)$ -th predicted medication at the current visit T ; \odot denotes a dot-product.

(b) Current Health-Relevant Past Medication Representation. Next, VITA obtains a current health-relevant past medication representation $\bar{\mathbf{p}}_k^T$. To achieve this, VITA evaluates two levels (*spec.*, medication-level $r_{m,i}^t \in \mathbb{R}$ and visit-level $r_v^t \in \mathbb{R}$) of relevance between the current visit and each past visit (Wu et al. 2022b). Specifically, the relevance $r_{m,i}^t$ at the medication-level indicates how much each medication i prescribed at each of the patient’s past visits t is related to the k -th medication that is necessary at her current visit T (i.e., \mathbf{p}_k^T), as follows: $\forall i \in \mathbf{m}^t; \forall t \in \{1, \dots, (T-1)\}$,

$$r_{m,i}^t = \frac{\exp((\mathbf{e}_i^t \odot \mathbf{p}_k^T) / \sqrt{\text{dim}})}{\sum_{j=1}^{|\mathbf{m}^t|} \exp((\mathbf{e}_j^t \odot \mathbf{p}_k^T) / \sqrt{\text{dim}})}, \quad (7)$$

where \mathbf{e}_i^t denotes a medication representation of medication i prescribed at the t -th visit; \mathbf{m}^t is defined in Section 3.1. While the relevance r_v^t at the visit-level indicates how much the patient’s health status at each past visit t is related to her current health status. So, VITA first obtains the patient representation \mathbf{q}^t at the t -th past visit for all her past visits in the same way as in its encoder. Then, VITA employs the target-aware attention module, using the patient representation \mathbf{q}^T as a query and the patient representations \mathbf{q}^t and \mathbf{q}^t for all her visits as keys and values (refer to Eq. (4)).

Given the two levels $r_{m,i}^t$ and r_v^t of relevance between the current and each past visit, VITA obtains the current health-relevant past medication representation $\bar{\mathbf{p}}_k^T$, as follows:

$$\bar{\mathbf{p}}_k^T = \sum_{t=1}^{(T-1)} r_v^t \mathbf{r}_m^t \quad (\mathbf{r}_m^t \in \mathbb{R}^{|\mathcal{M}|}), \quad (8)$$

$$\text{where } \forall i \in \mathcal{M}, \mathbf{r}_m^t = \begin{cases} r_{m,i}^t, & i \in \mathbf{m}^t, \\ 0, & \text{otherwise.} \end{cases}$$

Finally, VITA fuses the two medication representations \mathbf{p}_k^T and $\bar{\mathbf{p}}_k^T$ to obtain the k -th medication of the necessary medications for the patient’s current visit T , as follows:

$$\hat{\mathcal{M}}^T = \hat{\mathcal{M}}^T \cup \text{argmax}_{k \in \{1, \dots, |\mathcal{M}|\}} (\hat{\mathbf{p}}_k^T), \quad (9)$$

$$\text{where } \hat{\mathbf{p}}_k^T = \lambda_k \{\text{softmax}(\mathbf{p}_k^T \mathbf{W}_p + \mathbf{b}_p)\} + (1 - \lambda_k) \bar{\mathbf{p}}_k^T,$$

where the set $\hat{\mathcal{M}}^T$ begins as an empty set; $\hat{\mathbf{p}}_k^T \in \mathbb{R}^{|\mathcal{M}|}$ denotes the probabilities that all medications will be recommended as the k -th medication at her current visit T ; $\mathbf{W}_p \in \mathbb{R}^{\text{dim} \times |\mathcal{M}|}$ and $\mathbf{b}_p \in \mathbb{R}^{|\mathcal{M}|}$ denote a learnable weight matrix and a bias vector, respectively; $\lambda_k \in \mathbb{R}$ denotes a learnable parameter. In other words, based on such a $\hat{\mathbf{p}}_k^T$, VITA predicts the medication with the highest probability as the k -th medication.

3.3 Training

For the medications predicted by VITA, we employ the cross-entropy loss as the objective function, same as in (Wu et al. 2022b), to learn the medication representations \mathbf{e}_i and other learnable parameters of VITA, as follows:

$$\mathcal{L} = - \sum_{t=1}^T \sum_{i=1}^{|\mathcal{M}|} \mathbf{m}_i^t \log(\hat{\mathbf{m}}_i^t). \quad (10)$$

4 Evaluation

In this section, we conducted extensive experiments aiming at answering the following key research questions (RQs):

- **RQ1:** Does VITA provide more accurate recommendation than state-of-the-art medication recommenders?
- **RQ2:** Is each of VITA’s two core ideas effective for medication recommendation?
- **RQ3:** Does equipping existing methods with VITA’s two core ideas consistently improve their accuracies?
- **RQ4:** Which past visit information does VITA select from all of the patients’ past visit information?

4.1 Experimental Settings

Datasets. We used the MIMIC-III dataset (Johnson et al. 2016), widely used in medication recommendation studies (Shang et al. 2019b; Yang et al. 2021a,b; Wu et al. 2022b), and the MIMIC-IV dataset (Johnson et al. 2023), which is a follow-up dataset to MIMIC-III dataset. Following (Shang et al. 2019b; Yang et al. 2021a,b; Wu et al. 2022b), we also filtered out the patients who visited only once and the medical records related only to them. Please refer to the Appendix for detailed statistics of the two datasets.

Competitors. We compared VITA with six competitors: one basic classifier (Nearest as used in (Shang et al. 2019b)) and five state-of-the-art medication recommender systems (LEAP (Zhang et al. 2017), GAMENet (Shang et al. 2019b), MRSC (Wang et al. 2021b), SafeDrug (Yang et al. 2021b), and COGNet (Wu et al. 2022b)).

Evaluation Protocols. We randomly split the patients in each dataset into training (4/6), validation (1/6), and test (1/6) sets as in (Shang et al. 2019b; Yang et al. 2021b; Wu et al. 2022b). To measure accuracy, as in (Shang et al. 2019b; Wang et al. 2021a,b; Yang et al. 2021b; Wu et al. 2022b), we used the following three measures: Jaccard, PRAUC, and F1. In addition, following (Shang et al. 2019b; Wang et al. 2021a; Yang et al. 2021b; Wu et al. 2022b), we used the DDI rate (Shang et al. 2019b) to measure the possibility of adverse effects between recommended medications. However, due to space limitations, we present only the results in terms of accuracy here. Please refer to the Appendix for the results from the DDI rate. Furthermore, for each measure, we report the average values over five independent runs for most of our experiments; in Tables 1, 2 and Figure 3, all improvements are statistically significant with a p -value ≤ 0.001 .

4.2 Results and Analysis

Due to space limitations, for RQ3 and RQ4, we present the results on the MIMIC-III dataset only. Please refer to the Appendix for the results from the MIMIC-IV dataset.

Datasets	MIMIC-III / MIMIC-IV		
	Jaccard	PRAUC	F1
Nearest	0.392 / 0.452	0.381 / 0.446	0.547 / 0.605
LEAP	0.431 / 0.427	0.638 / 0.595	0.594 / 0.584
GAMENet	0.435 / 0.460	0.675 / 0.703	0.590 / 0.611
MRSC	0.486 / 0.449	0.741 / 0.691	0.645 / 0.606
SafeDrug	0.506 / 0.473	0.754 / 0.698	0.663 / 0.625
COGNet	0.509 / 0.494	0.757 / 0.707	0.667 / 0.647
VITA	0.528 / 0.522	0.767 / 0.715	0.682 / 0.669

Table 1: Accuracies of six competitors and VITA. The best and the second-best results in each column (*i.e.*, each measure) are in bold and underlined, respectively.

RQ1: Comparison with Six Competitors. To demonstrate the superiority of VITA, we compared the accuracy of VITA and those of the six competitors⁶.

As shown in Table 1, VITA outperforms *all* competitors on *all* datasets for *all* measures consistently. Specifically, on the MIMIC-III and -IV datasets, VITA outperforms the best competitor (*i.e.*, COGNet) by up to 3.73% and 5.67% in terms of Jaccard. The improvement gap increased from 3.73% in the MIMIC-III dataset to 5.67% in the MIMIC-IV dataset in terms of Jaccard. This is interpreted as our idea of selecting and using only the past visits relevant to the patient’s current visit becoming more important as the number of visits increases (the MIMIC-IV dataset, compared to the MIMIC-III dataset, contains about 2.15 times as many patients and about 2.53 times as many visits, with the average number of visits per patient increasing from 2.59 to 3.05).

Also, in the Appendix, we will analyze the training time of VITA and competitors for a more comprehensive understanding of the efficiency of VITA’s encoder.

RQ2: Ablation Study. In order to verify the effectiveness of our two novel ideas of VITA (relevant-visit selection and target-aware attention), we conducted comparative experiments using VITA and its variants.

VITA’s variants for confirming the effectiveness of relevant-visit selection. (1) VITA-R does *not* use the relevant-visit selection module in VITA. (2) VITA-R_{T-1} uses only one past visit, which is most similar to the current visit of a patient in terms of Jaccard similarity, and the current visit information as input for its encoder, instead of using the relevant-visit selection module. Also, we note that, when the value of temperature of the attention network is close to zero, the difference between the attention weights increases, sharpening their distribution (Hinton, Vinyals, and Dean 2015; Nguyen, Pernkopf, and Kosmider 2020). It will make our target-aware attention module have an effect similar to using only the past visits relevant to the current visit; in other words, this will be able to play a role similar to the relevant-visit selection module. Therefore, we measured the accuracies of VITA-R, while varying the temperature τ_a of its target-aware attention module from 1 to 0.2 by 0.2 (please refer to the Appendix for detailed results of this experiment);

⁶We carefully tuned the hyperparameters of all methods. Please refer to the Appendix for the specific hyperparameter values.

Datasets	MIMIC-III / MIMIC-IV		
	Jaccard	PRAUC	F1
VITA	0.528 / 0.522	0.767 / 0.715	0.682 / 0.669
VITA-R	0.516 / 0.520	0.756 / 0.709	0.671 / 0.659
VITA-R _{T-1}	0.514 / 0.480	0.741 / 0.694	0.670 / 0.635
VITA-R _{sha.}	0.519 / 0.520	0.762 / 0.709	0.672 / 0.664
VITA-T _{avg.}	0.512 / 0.513	0.751 / 0.707	0.665 / 0.660
VITA-T _{rnn}	0.515 / 0.517	0.754 / 0.701	0.668 / 0.663
VITA-T _{att.}	0.518 / 0.520	0.759 / 0.709	0.670 / 0.662

Table 2: The effects of VITA’s two core ideas (relevant-visit selection and target-aware attention). The best result in each column (*i.e.*, each measure) is in bold.

we considered (3) VITA-R_{sha.}, which is VITA-R with the best-performing temperature τ_a .

VITA’s variants for confirming the effectiveness of target-aware attention. (1) VITA-T_{avg.} (resp. (2) VITA-T_{rnn}) employs the mean pooling (resp. a RNN-based model (*spec.*, GRU)) instead of the target-aware attention module in VITA when fusing the past visit information relevant to the current visit. (3) VITA-T_{att.} employs a typical attention network that uses the current visit information as a query and the relevant past visit information as keys and values, instead of the target-aware attention module, when fusing the past visit information relevant to the current visit.

Table 2 shows the accuracies of VITA and its all variants.

Results regarding the effectiveness of relevant-visit selection. We observed that VITA outperforms VITA-R. This result indicates that using only the past visits relevant to the current visit is important in recommending effective medications to patients; in other words, the past visits irrelevant to the current visit should not be considered at all when recommending medications. However, even though VITA-R_{T-1} employs *only one* past visit most *relevant* to the current visit of a patient in terms of Jaccard similarity, this shows lower accuracy than VITA-R. This is because, although the number of past visits relevant to the current visit may differ per patient, this employs a fixed number of past visits for all patients, which has adversely affected learning by including (resp. excluding) the past visits that are actually irrelevant (resp. relevant) to the current visit. This supports the design of our relevant-visit selection module, which does not have such a restriction. We also observed that VITA-R_{sha.} showed lower accuracy than VITA. Note that it is difficult for VITA-R_{sha.} to flexibly select any number of past visits relevant to the current visit for each patient, because it operates similarly to the argmax function as the distribution of attention weights continues to sharpen.

Results regarding the effectiveness of target-aware attention. We observed that VITA outperforms all its variants related to target-aware attention (*i.e.*, VITA-T_{avg.}, VITA-T_{rnn}, and VITA-T_{att.}), confirming the effectiveness of the target-aware attention module. Also, we observed that the accuracies for most measures improve in the order of VITA-T_{avg.}, VITA-T_{rnn}, and VITA-T_{att.}. The results indicate that (1) capturing the degree of relevance between the current




Figure 3: Accuracies of four longitudinal-based methods and their variants equipped with two core ideas of VITA.

visit and each past visit plays an important role in obtaining higher accuracy (this is from the fact the accuracy of VITA- T_{avg} is worse than those of VITA- T_{rnn} , VITA- T_{att} , and VITA in most measures); (2) such relevance may not be consistent with the order of the visits (this is from the fact the accuracy of VITA- T_{rnn} is worse than those of VITA- T_{att} and VITA in most measures); (3) it is important to accurately capture the degree of relevance between the current visit and each past visit in providing better accuracy (this is from the fact the accuracy of VITA- T_{att} is worse than that of VITA in all measures). These findings again support the limitation of existing works that we previously pointed out.

Also, it bears mentioning that all variants of VITA, except for VITA- R_{T-1} , outperform the *best competitor* (i.e., COGNet) for most measures (compare Tables 1 and 2), even though they incorporate *only* one of VITA’s two core ideas; this validates that VITA’s single core idea alone is beneficial in improving the accuracy of medication recommendation.

RQ3: Compatibility of VITA’s Two Core Ideas. In Section 3.2, we claimed that VITA’s two core ideas can be applied orthogonally to most medication recommenders, thereby improving their accuracy. To validate this claim, we compare the accuracy of four longitudinal-based methods (GAMNet, MRSC, SafeDrug, and COGNet) and their variants, each of which is equipped with the two core ideas of VITA (GAMNet(VITA), MRSC(VITA), SafeDrug(VITA), and COGNet(VITA)); the results of the methods are represented as the first three letters of the methods in Figure 3.

As shown in Figure 3, the variants of existing methods equipped with the two core ideas of VITA outperform the original methods in Jaccard (please refer to the Appendix for the results from PRAUC and F1); that is, the result shows that our two core ideas are equipped orthogonally to most medication recommenders, improving their accuracy.

RQ4: Analysis of Selected Past Visits. One of the two core ideas of VITA, the relevant-visit selection, automatically selects only the past visit relevant to the current visit. To investigate which past visit of patients was selected by the relevant-visit selection module, we analyze the visit information of patients in the test set. We first calculated the Jaccard similarity between the current and each past visit information of a patient in the test set as in Section 2. Then, we divided the past visits of patients into the following two groups: (i) (resp. (ii)) the past visits selected (resp. not selected) by the relevant-visit selection module; the statistics of the Jaccard similarity for each group, along with those of all past visits for comparison, are represented in Figure 4.

As shown in Figure 4, the relevant-visit selection module tended to select the past visit information with a high Jaccard similarity to the current visit information. Numerically, the




Figure 4: The statistics of the Jaccard similarity between patients’ current visit and ‘all’ past visits, (i) (resp. (ii)) the past visits ‘selected’ (resp. ‘not selected’) by the relevant-visit selection module.

average of Jaccard similarity between the past visit information selected by the relevant-visit selection module and the current visit information is 0.2164, which is 22.81% higher than that for the past visit not selected by the relevant-visit selection module (0.1762). Note that the relevant-visit selection module did not always select the past visit information similar to the current one of the patients in Jaccard similarity. This implies that, even if the past visit information is not similar to the current one in Jaccard similarity (i.e., on the face of it), it can be useful for accurate medication recommendation (e.g., due to its possible inherent relevance with the current one), as inferred from the superior performance of VITA, which is equipped with the relevant-visit selection. Additionally, a case study on past visits dissimilar to the current visit in Jaccard similarity but selected by the relevant-visit selection module would be shown in the Appendix.

5 Conclusions

In this paper, we demonstrated the following two important points w.r.t. representing a patient’s current health status: (1) it is required to accurately capture the relevance between current visit information and each of past visit information; (2) using past visit information irrelevant to the current one is harmful; in other words, only the past visits relevant to a patient’s current visit should be ‘carefully chosen’, and the chosen past visits should be ‘weighted less’ when aggregated according to the degree of relevance. Considering these points, we proposed a novel medication recommendation framework, named VITA, based on the following two core ideas: *relevant-visit selection*, which allows for flexible selection of only the past visits relevant to the current visit per patient (even allowing for the selection of *all* past visits or selecting *none* of them, if necessary); *target-aware attention*, which accurately captures the relevance score between a patient’s current visit and each of past visits (even covering when *all* the past visits are *weakly* relevant to the current visit). Our two novel ideas are effective in improving the accuracy, thereby making VITA, which is the final version equipped with all the ideas, consistently and significantly more accurate than its six competitors on real-world datasets. Furthermore, our two core ideas can be applied to the various medication recommender systems (even in the recommender systems of another domain) orthogonally.

Acknowledgments

This work was supported by Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2023 (Project Name: Development of Intelligent Personalized Rehabilitation Service Technology, Project Number: SR202104001, Contribution Rate: 33.34%) and by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00352 and No. RS-2022-00155586).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Bhoi, S.; Lee, M. L.; Hsu, W.; Fang, A. H. S.; and Tan, N. C. 2021. Personalizing Medication Recommendation with a Graph-Based Approach. *ACM Transactions on Information Systems*, 40: 1–23.
- Gong, F.; Wang, M.; Wang, H.; Wang, S.; and Liu, M. 2021. SMR: Medical Knowledge Graph Embedding for Safe Medicine Recommendation. *Big Data Research*, 23: 100174.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Johnson, A.; Bulgarelli, L.; Pollard, T.; Horng, S.; Celi, L. A.; and Mark, R. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, 10: 1.
- Johnson, A. E.; Pollard, T. J.; Shen, L.; Lehman, L.-w. H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3: 1–9.
- Nguyen, T.; Pernkopf, F.; and Kosmider, M. 2020. Acoustic Scene Classification for Mismatched Recording Devices Using Heated-Up Softmax and Spectrum Correction. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 126–130.
- Shang, J.; Ma, T.; Xiao, C.; and Sun, J. 2019a. Pre-training of Graph Augmented Transformers for Medication Recommendation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 5953–5959.
- Shang, J.; Xiao, C.; Ma, T.; Li, H.; and Sun, J. 2019b. GAMENet: Graph Augmented MEMory Networks for Recommending Medication Combination. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, 1126–1133.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, volume 30, 5998–6008.
- Wang, L.; Zhang, W.; He, X.; and Zha, H. 2018. Personalized Prescription for Comorbidity. In *Proceedings of the International Conference on Database Systems for Advanced Applications (DASFAA)*, 3–19.
- Wang, Y.; Chen, W.; Pi, D.; Yue, L.; Wang, S.; and Xu, M. 2021a. Self-Supervised Adversarial Distribution Regularization for Medication Recommendation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3134–3140.
- Wang, Y.; Chen, W.; Pi, D.; Yue, L.; Xu, M.; and Li, X. 2021b. Multi-hop Reading on Memory Neural Network with Selective Coverage for Medication Recommendation. In *Proceedings of the International ACM Conference on Information and Knowledge Management (CIKM)*, 2020–2029.
- Wu, J.; Qian, B.; Li, Y.; Gao, Z.; Ju, M.; Yang, Y.; Zheng, Y.; Gong, T.; Li, C.; and Zhang, X. 2022a. Leveraging multiple types of domain knowledge for safe and effective drug recommendation. In *Proceedings of the International ACM Conference on Information and Knowledge Management (CIKM)*, 2169–2178.
- Wu, R.; Qiu, Z.; Jiang, J.; Qi, G.; and Wu, X. 2022b. Conditional Generation Net for Medication Recommendation. In *Proceedings of the international ACM conference on World Wide Web (WWW)*, 935–945.
- Yang, C.; Xiao, C.; Glass, L.; and Sun, J. 2021a. Change Matters: Medication Change Prediction with Recurrent Residual Networks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3728–3734.
- Yang, C.; Xiao, C.; Ma, F.; Glass, L.; and Sun, J. 2021b. SafeDrug: Dual Molecular Graph Encoders for Recommending Effective and Safe Drug Combinations. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3735–3741.
- Zhang, Y.; Chen, R.; Tang, J.; Stewart, W. F.; and Sun, J. 2017. LEAP: Learning to Prescribe Effective and Safe Treatment Combinations for Multimorbidity. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 1315–1324.