



# Leave No Patient Behind: Enhancing Medication Recommendation for Rare Disease Patients

Zihao Zhao

University of Science and Technology  
of China  
Hefei, Anhui, China  
zzh1998@mail.ustc.edu.cn

Yi Jing

University of Science and Technology  
of China  
Hefei, Anhui, China  
jingi@mail.ustc.edu.cn

Fuli Feng\*

University of Science and Technology  
of China  
Hefei, Anhui, China  
fulifeng93@gmail.com

Jiancan Wu

University of Science and Technology  
of China  
Hefei, Anhui, China  
wujcan@gmail.com

Chongming Gao

University of Science and Technology  
of China  
Hefei, Anhui, China  
chongming.gao@gmail.com

Xiangnan He\*

University of Science and Technology  
of China  
Hefei, Anhui, China  
xiangnanhe@gmail.com

## ABSTRACT

Medication recommendation systems have gained significant attention in healthcare as a means of providing tailored and effective drug combinations based on patients' clinical information. However, existing approaches often suffer from fairness issues, as recommendations tend to be more accurate for patients with common diseases compared to those with rare conditions. In this paper, we propose a novel model called Robust and Accurate REcommendations for Medication (RAREMed), which leverages the pretrain-finetune learning paradigm to enhance accuracy for rare diseases. RAREMed employs a transformer encoder with a unified input sequence approach to capture complex relationships among disease and procedure codes. Additionally, it introduces two self-supervised pre-training tasks, namely Sequence Matching Prediction (SMP) and Self Reconstruction (SR), to learn specialized medication needs and interrelations among clinical codes. Experimental results on two real-world datasets demonstrate that RAREMed provides accurate drug sets for both rare and common disease patients, thereby mitigating unfairness in medication recommendation systems. The implementation is available via <https://github.com/zzhUSTC2016/RAREMed>

## CCS CONCEPTS

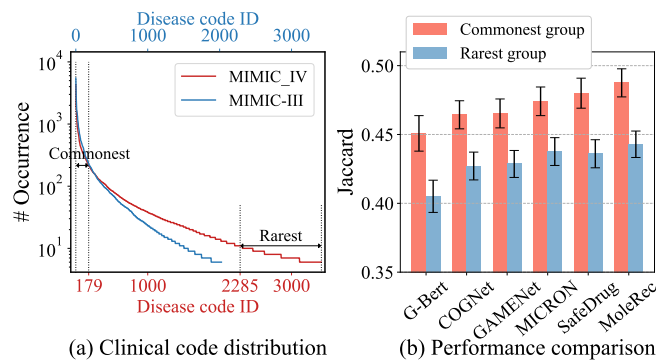
• Information systems → Recommender systems; • Applied computing → Health informatics.

## KEYWORDS

Medication recommendation; Electronic health record; Fairness; Rare disease

\*Corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
SIGIR '24, July 14–18, 2024, Washington, DC, USA.  
© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-0431-4/24/07...\$15.00  
<https://doi.org/10.1145/3626772.3657785>



**Figure 1: (a) The long-tail distribution of disease codes in the MIMIC-III [14] and MIMIC-IV [13] datasets. (b) Patients in MIMIC-IV are divided into five equal-sized groups based on the ranking of their least common disease codes in terms of popularity. The subgraph illustrates the recommendation accuracy, as measured by the Jaccard index, of SOTA methods for both the commonest and the rarest code groups.**

## ACM Reference Format:

Zihao Zhao, Yi Jing, Fuli Feng, Jiancan Wu, Chongming Gao, and Xiangnan He. 2024. Leave No Patient Behind: Enhancing Medication Recommendation for Rare Disease Patients. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3626772.3657785>

## 1 INTRODUCTION

Medication recommendation has garnered growing interest as a means of offering secure and effective drug combinations tailored to patients' clinical information [1]. Through the utilization of medical knowledge and extensive patient data, a medication recommendation system can facilitate healthcare professionals in prescribing more precise and effective drug combinations while reducing the occurrences of medication errors and adverse drug-drug interactions (DDI) among selected medicines [29, 39]. The substantial harm caused by inaccurate clinical decisions has led to a profound impact on numerous patients [44]. Notably, it is estimated that at least 1.5 million individuals [2] are harmed by preventable medication errors

each year, incurring a substantial raises of healthcare expenses [17]. Consequently, it is imperative to explore improved methodologies for medication recommendations.

Current approaches to medication recommendation primarily focus on maximizing overall accuracy by incorporating patients' current medical conditions [34, 44] or clinical information from previous visits [28, 29, 31, 37–39]. While significant progress has been made in comprehensive utilization these data, existing approaches encounter fairness issues in terms of recommendations across patient groups. This stems from the highly skewed distribution of clinical codes in Electronic Health Records (EHR) [16] — a digital compilation of patients' medical conditions, including diseases, procedures, medications and other clinical data during patient visits. As illustrated in Figure 1(a), a small subset of codes exhibit high prevalence, while the long tail comprises low-occurrence diseases in two real-world EHR datasets. Consequently, patients resided in the most common group receive significantly more accurate recommendations compared to those in the rarest group, as depicted in Figure 1(b). This disparity in accuracy across patient groups raises concerns regarding equity and fairness, ultimately undermining the reliability and utility of medication recommendation systems.

The key to promoting equity in medication recommendation is to improve predictive accuracy for underserved patient populations, particularly those with rare conditions. Nevertheless, accurately modeling and recommending medications for rare diseases remains as a challenging task. We summarize three central difficulties:

- **Scarce quality data:** Due to the inherently low prevalence of rare diseases, there is a paucity of clinical data related to these conditions in EHR [21]. This scarcity of high-quality data severely impedes the accurate discernment of medication recommendation patterns for rare diseases.
- **Intricate disease/procedure associations<sup>1</sup>:** Distinct from common diseases, rare diseases often present complex combinations of symptoms and clinical events [10, 22], which confounds efforts to elucidate reliable disease associations. Moreover, the complex relationships between diseases and procedures further obfuscate appropriate medication selection [7, 30] — certain treatments may be commonly applied across multiple conditions while the suitability varies depending on disease severity.
- **Tailored medication needs<sup>1</sup>:** The atypical pathophysiology of rare diseases often necessitates more precisely targeted pharmacological interventions. On the other hand, patients with rare diseases tend to have more complicated therapeutic regimens, warrant tailored medication recommendations.

In this work, we aim to foster fair medication recommendations by uplifting their accuracy on rare diseases. We resort to the pretrain-finetune learning paradigm for addressing the aforementioned challenges, drawing inspiration from its success in natural language processing [8, 24] and computer vision [5, 12], which yet remains largely unexplored in medication recommendation. Compared to traditional learning paradigms, *e.g.*, supervised learning, pretrain-finetune enjoys the merit of exploiting the unlabeled data space through self-supervised pretraining, showing great potential in modeling intricate relations and addressing data scarcity.

<sup>1</sup>We will further discuss the profile of rare disease patients in Section 4.

However, implementing pretrain-finetune in fair medication recommendations presents challenges, which we summarize in two key questions: (1) **How to build an expressive encoder to fully capture patients' clinical information?** This encoder should be able to handle the complexity of rare diseases, ensuring that the medication recommendations are based on a thorough understanding of the patient's condition. (2) **How to design optimal pretraining objectives to learn rare diseases' specialized medication needs and interrelations among diseases, procedures, and medications?**

These two questions motivate us to propose a novel model called Robust and Accurate REcommendations for Medication (RAREMed), which brings the superiority of pretrain-finetune to representation learning for fair medication recommendations. RAREMed utilizes a transformer encoder architecture, which is effective at capturing complex relationships. To comprehensively represent patient information, unlike prior studies that overlook the intricate connections between diseases and procedures [29, 31, 37, 39, 41], we introduce a unified sequence approach to jointly encode both components. This allows modeling potential associations, *e.g.*, certain procedures are applied differently depending on disease type and severity<sup>2</sup> [7, 30], better capturing rare disease complexity for accurate recommendations. Additionally, we enrich the standard token embedding layer with *segment embeddings* and *relevance embeddings*, where segment embeddings help differentiate between disease and procedure codes, while relevance embeddings capture their varying importance. To learn rare diseases' tailored needs and maximize the utilization of available data, we design two self-supervised pretraining tasks: Sequence Matching Prediction (SMP) and Self Reconstruction (SR). SMP predicts whether a disease-procedure sequence pair belongs to the same patient, enabling the learning of contextual dependencies and underlying connections among clinical codes. While SR reconstructs the input sequence from the patient representation, promoting a comprehensive comprehension of all codes, especially the rare ones. By pretraining on these objectives, RAREMed acquires robust encoding abilities, even for limited clinical codes. Through experiments on two public real-world datasets, we demonstrate that RAREMed provides accurate drug sets for both rare disease and common disease patients, mitigating the unfairness in the medication recommendation system.

In summary, our work makes the following main contributions:

- To our knowledge, this is the first work to address the unfairness issue in medication recommendation, whereby patients with rare diseases cannot obtain accurate recommendations.
- We propose a novel medication recommendation model, RAREMed, which combines pre-training techniques to learn robust representations for rare diseases, thereby improving the accuracy and mitigating unfairness in medication recommendation.
- Extensive experiments on two benchmark datasets demonstrate RAREMed's superiority over a range of state-of-the-art methods. We will release the source code to facilitate future research.

<sup>2</sup>For example, the typical choice of non-invasive shock wave lithotripsy for kidney stones suggests a better overall health condition for the patient, whereas endoscopic surgery is considered a more severe option. [30]. In contrast, prostate enlargement is commonly treated with endoscopic surgery for mild cases and shock wave lithotripsy for severe cases involving prostate stones [7].

## 2 RELATED WORKS

**Medication Recommendation.** Methods for medication recommendation can be broadly categorized into two groups: longitudinal and instance-based approaches. Longitudinal methods make use of patients' longitudinal medical history, such as the approach taken by Choi et al. [6], who employ a two-level temporal attention mechanism. Similarly, Shang et al. [28] pre-train their model on single-visit data<sup>3</sup> and fine-tune it using multi-visit data<sup>4</sup>. Additionally, Shang et al. [29] incorporate a graph-augmented memory module and DDI graph to reduce adverse drug-drug interactions, while Yang et al. [39] consider drug molecule structures for medication security. Wu et al. [37] introduce a copy-or-predict mechanism, and Yang et al. [38] propose a residual-based recurrent network. Sun et al. [31] present a causal model to address recommendation bias, and Yang et al. [41] leverage a molecular substructure-aware attentive method. More recently, Bhoi et al. [3] propose a fine-grained medication recommendation approach. However, many of these methods heavily rely on historical records, making them less suitable for single-visit patients [3, 6, 28, 31], or exhibit reduced accuracy in such cases [29, 37–39, 41]. Instance-based methods focus on the current medical condition of patients. For instance, Zhang et al. [44] utilize an RNN model based on disease codes. Wang et al. [34] employ Deep Q Learning to capture correlations and adverse interactions between medicines. However, these models may not adequately capture patient-specific information, potentially leading to reduced accuracy. Despite the initial success of these methods, a common issue among them is fairness [4, 35]. Patients with common diseases tend to receive more accurate medication predictions compared to those with rare diseases. This discrepancy undermines the overall performance and impairs the reliability and usefulness of medication recommendation systems.

**Fairness in Recommendation.** Fairness concerns in recommender systems have garnered significant attention due to observed systematic and unfair discrimination against certain individuals or groups [4, 35]. Approaches to achieve fairness in machine learning can be categorized into three groups: Pre-processing, In-processing, and Post-processing [19]. Pre-processing methods modify the training data before the learning process to eliminate biases. For example, Ekstrand et al. [9] address gender bias in movie/music recommendation by creating gender-balanced data through random resampling. Rastegarpanah et al. [25] propose adding "antidote" data to improve the social desirability of recommender system outputs. However, in medication recommendation, it is ethically unfeasible to modify individual patients' Electronic Health Record (EHR) data, limiting the effectiveness of these pre-processing methods. In-processing methods focus on removing sensitive attribute information through regularization or adversarial learning. For instance, Yao et al. [42] propose fairness metrics as regularization in collaborative filtering. Li et al. [18] introduce a text-based reconstruction loss to ensure more balanced recommendation utility across all users. However, these methods often involve non-convex optimization and the fairness issues in this research is irrelevant to sensitive attributes. Post-processing methods involve re-ranking the output list of base ranking models to improve fairness. For instance, Geyik

et al. [11] employ interval constrained ranking to achieve multiple group fairness. However, in medication recommendation, the focus is on recommending a set of drugs rather than a ranking of individual drugs, making the relevance of re-ranking less significant in this task. In short, while methods have been proposed to tackle unfairness in recommender systems, their applicability to the specific fairness challenges in medication recommendation is limited.

**Pre-training in Recommendation.** Pre-training techniques are widely used in feature-based recommendation systems to enhance user or item representations [43]. Qiu et al. [23] propose review encoder pre-training to complement user representations, while Wong et al. [36] utilize pre-training on a large-scale knowledge graph for conversational recommender systems. However, these methods are not applicable to medication recommendation. Shang et al. [28] introduce a pre-training method to address selection bias, but it lacks procedure coding and recommendations for patients without historical records, and fairness concerns remain unresolved.

## 3 PROBLEM FORMULATION

**Electronic Health Record (EHR).** An EHR, denoted as  $\mathcal{R} = \{\mathcal{V}^{(j)}\}_{j=1}^N$ , represents a collection of medical records for  $N$  patients. In this context,  $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$  refers to the set of diseases,  $\mathcal{P} = \{p_1, p_2, \dots, p_{|\mathcal{P}|}\}$  denotes the set of procedures, and  $\mathcal{M} = \{m_1, m_2, \dots, m_{|\mathcal{M}|}\}$  represents the set of medications. The notation  $|\cdot|$  indicates the cardinality of a set.

The record of patient  $j$ , denoted as  $\mathcal{V}^{(j)} = \{\mathbf{d}^{(j)}, \mathbf{p}^{(j)}, \mathbf{m}^{(j)}\}$ , comprises three components. Firstly,  $\mathbf{d}^{(j)} = [d_1, d_2, \dots, d_x] \in \mathcal{D}$  represents the sequence of diseases experienced by the patient, ordered by their priority or significance<sup>5</sup>, reflecting their relative importance in the patient's healthcare journey. Similarly,  $\mathbf{p}^{(j)} = [p_1, p_2, \dots, p_y] \in \mathcal{P}$  captures the sequence of procedures, also ordered by priority. Here,  $x$  and  $y$  denote the lengths of the disease and procedure sequences, respectively. Lastly,  $\mathbf{m}^{(j)} \in \{0, 1\}^{|\mathcal{M}|}$  is a multi-hot vector indicating the medications prescribed to the patient. Each element in  $\mathbf{m}^{(j)}$  corresponds to a specific medication in  $\mathcal{M}$ , and a value of 1 indicates that the medication was prescribed, while 0 indicates its absence. We will simplify the notation by omitting the superscript  $j$  whenever there is no ambiguity.

**DDI Graph.** The Drug-Drug Interaction (DDI) Graph [26, 29, 39] is a tool used to assess and mitigate the risk of adverse drug-drug interactions within recommended medication combinations. It is represented by a symmetric binary adjacency matrix  $\mathbf{A} \in \{0, 1\}^{|\mathcal{M}| \times |\mathcal{M}|}$ , where each element  $\mathbf{A}_{uv}$  is assigned a value of 1 only if there exists a harmful interaction between the  $u$ -th and  $v$ -th drugs.

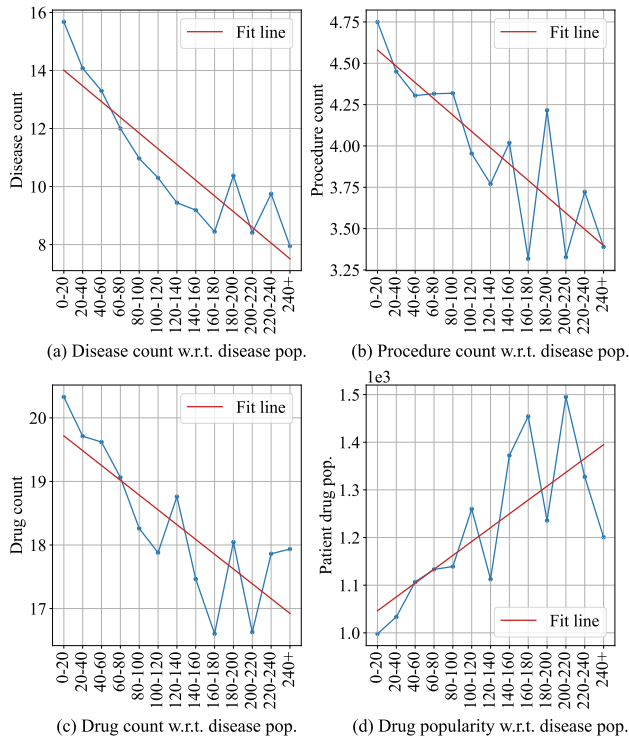
**Medication Recommendation Problem.** Given a patient's current disease sequence  $\mathbf{d}$  and procedure sequence  $\mathbf{p}$ , along with the DDI graph  $\mathbf{A}$ , the objective is to generate an appropriate combination of medications  $\hat{\mathcal{Y}}$  for the patient that maximizes prediction accuracy while minimizing the risk of adverse drug interactions.

**Fair Medication Recommendation Problem.** In addition to the conventional focus on overall accuracy, our research emphasizes the fairness of medication recommendation. Fairness in this context refers to providing accurate recommendations for both patients with common diseases and those with rare diseases.

<sup>3</sup>Single-visit patients have no historical visits, i.e., have only one visit recorded in EHR.

<sup>4</sup>Multi-visit patients have at least two visits in EHR.

<sup>5</sup>The sequence numbers of priority are provided in the EHR datasets [13, 14], which are labeled by qualified physicians.



**Figure 2: We categorize patients into 13 distinct groups based on the frequency of their least common disease codes. Sub-plots (a), (b), (c), and (d) depict the relationship between the average disease count, procedure count, drug count, and drug popularity, respectively, w.r.t. the rarest disease popularity.**

#### 4 PROFILE OF RARE DISEASE PATIENTS

In this section, we conduct empirical analysis on real-world EHR dataset MIMIC-IV [13] to explore the profile of patients with rare diseases. Specifically, we categorize patients into 13 groups based on the popularity of their rarest disease codes, using uniform segmentation across popularity intervals. We then calculate various metrics for each group. From the results presented in Figure 2, we make two important observations:

**Observation 1:** The data presented in Figure 2(a) and 2(b) clearly demonstrates a negative correlation between disease count and procedure count with patient disease popularity. This suggests that as the rarity of a patient’s disease increases, their clinical condition becomes more complex. Consequently, patients with rare diseases exhibit more intricate combinations of clinical conditions.

**Observation 2:** Figure 2(c) and 2(d) demonstrate that patients with rarer diseases tend to have a higher number of prescribed medications and a lower popularity of these medications. This indicates that rare diseases patients require more specifically tailored medication options, resulting in more complex therapeutic regimens.

In summary, our analysis highlights the increased complexity of both clinical conditions and therapeutic regimens in patients with rarer diseases. These findings underscore the need for more advanced techniques, such as expressive encoder and pre-training techniques, to effectively address the unique challenges faced by rare disease patients.

## 5 OUR METHOD

The model framework of our RAREMed is illustrated in Figure 3. We will commence by providing a detailed description of our encoder, which is specifically designed to effectively model patient clinical information. Subsequently, we will delve into the utilization the pre-training techniques to enhance the representation of input clinical codes, thereby ensuring equitable recommendation for patients with both common and rare diseases. Lastly, we will elaborate on the fine-tuning process of the pretrained model, customized for the medication recommendation task.

### 5.1 Patient Representation

In the context of medication recommendation, it is crucial to obtain a comprehensive representation of a patient’s clinical condition. Traditional approaches commonly utilize disease and procedure codes to represent patient information. However, these methods have limitations. Some overlook certain codes [28, 44], while others treat them as distinct elements and simply concatenate them without considering their intricate associations [29, 37–39, 41]. Consequently, these approaches result in models that are less expressive in capturing the complexity of patient data.

To address this limitation, our approach considers disease and procedure codes as a unified sequence and utilizes a transformer encoder to generate a comprehensive representation of the patient’s clinical condition. Formally, for a patient  $\mathcal{V} = \{\mathbf{d}, \mathbf{p}, \mathbf{m}\}$ , we construct the input sequence as follows:

$$\text{input} = [\text{CLS}] \oplus \mathbf{d} \oplus [\text{SEP}] \oplus \mathbf{p}, \quad (1)$$

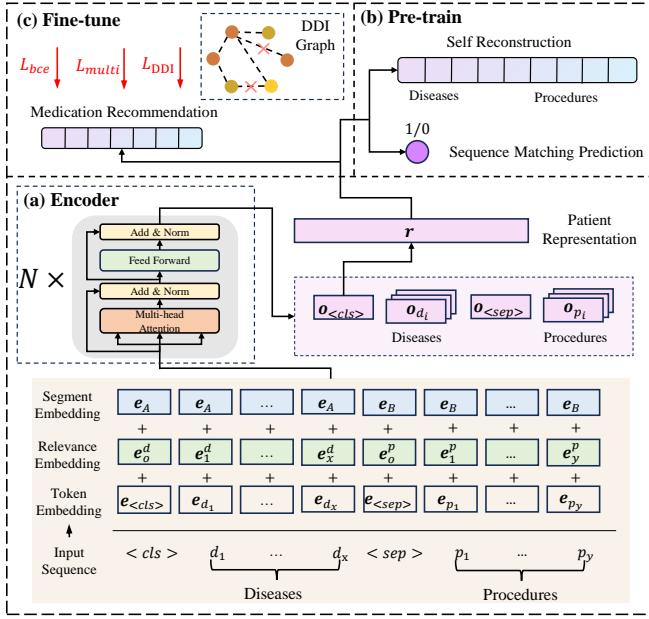
where [CLS] represents a special token typically placed at the beginning of the sequence, and its representation can be utilized as the patient’s overall representation. [SEP] indicates another special token that signifies the separation between the disease and procedure code sequences. The symbol  $\oplus$  denotes the concatenation operation between sequences.

To augment the standard token embedding layer, we incorporate two additional embeddings: segment embedding and relevance embedding. The segment embedding serves to differentiate between the two categories of input codes, specifically disease and procedure, allowing the model to discern context and distinguish between various types of medical information. While the relevance embedding is utilized to capture the varying significance of different diseases and procedures. Rather than treating all input codes equally, we recognize that certain diseases and procedures may exert greater influence on the patient’s clinical condition. To address this, we sort the input codes based on their relevance to the patient. Meanwhile, we introduce two learnable relevance embedding matrices to capture the priority information, in which  $\mathbf{e}_i^d$  and  $\mathbf{e}_j^p$  represent the relevance embedding of the  $i$ -th disease and  $j$ -th procedure codes, respectively.

Finally, the transformer encoder processes the embedded input sequence, yielding the final patient representation  $\mathbf{r}$ :

$$\mathbf{r} = \text{Encoder}(E_{\text{tok}}(\text{input}) + E_{\text{seg}}(\text{input}) + E_{\text{rel}}(\text{input}))[0], \quad (2)$$

where  $E_{\text{tok}}$ ,  $E_{\text{seg}}$ , and  $E_{\text{rel}}$  represent the token, segment, and relevance embedding layers respectively. We take the representation of [CLS] token as the patient representation.



**Figure 3: The illustration of our proposed RAREMed. We first model the patient clinical information with three embedding layers and a transformer encoder to get the patient representation. Then we pre-train our model on Sequence Matching Prediction task and Self Reconstruction task. Finally, the model is fine-tuned on medication recommendation task to get the recommended medication combination.**

## 5.2 Pre-training

In order to enhance the representation of clinical codes, especially for rare ones, we introduce two self-supervised pre-training tasks for the patient encoder. These tasks are specifically designed to leverage the inherent patterns and relationships within the clinical codes, enabling the model to better capture the nuances and complexities of various clinical conditions. Through pre-training on these tasks, our goal is to facilitate the acquisition of contextualized representations that can more effectively capture the intricacies of diverse clinical conditions, including rare cases.

**Task #1: Sequence Matching Prediction (SMP):** The objective of this task is to capture the intricate associations between disease and procedure codes more effectively. Specifically, we aim to train the model to discern whether the disease and procedure sequences belong to the same patient or not. This task is crucial for RAREMed to better understand the contextual dependencies between different clinical codes and improve its ability to capture the underlying connections within the patient’s clinical condition.

To achieve this objective, we create an unpaired example for each input sequence pair  $(d_i, p_i)$  by randomly substituting either the disease  $d_i$  or procedure sequence  $p_i$  with a corresponding sample from a different patient. Subsequently, our model is trained using Binary Cross-Entropy (BCE) loss to distinguish between the paired and unpaired inputs:

$$L_{SMP} = -\log(\hat{y}_i) + \log(1 - \hat{y}_j), \quad (3)$$

where  $\hat{y}_i = \sigma(W_1 r_i + b_1) \in \mathcal{R}$  represents the predicted probability for the paired input and  $\hat{y}_j$  for the unpaired input. Here,  $\sigma$  denotes the sigmoid function.  $W_1 \in \mathbb{R}^{dim}$  and  $b_1 \in \mathbb{R}$  are trainable parameters. **Task #2: Self Reconstruction (SR):** Considering the tailored medication needs in this domain, it is essential for the patient representation to retain knowledge of all the components in the input sequence, especially rare clinical codes. Specifically, in this task, we train RAREMed to reconstruct the input clinical code sequence from the modeled patient representation  $r$ . This encourages RAREMed to capture and preserve the essential information within the clinical codes, ensuring a comprehensive representation of the patient’s clinical condition.

The reconstruction loss, denoted by  $L_{SR}$ , is defined as:

$$L_{SR} = - \sum_{j=1}^{|\mathcal{D}|+|\mathcal{P}|} [c_j \log(\hat{c}_j) + (1 - c_j) \log(1 - \hat{c}_j)], \quad (4)$$

where  $\hat{c} = \sigma(W_2 r + b_2) \in [0, 1]^{|\mathcal{D}|+|\mathcal{P}|}$  represents the probabilities of all diseases and procedures reconstructed by our model,  $W_2 \in \mathbb{R}^{(|\mathcal{D}|+|\mathcal{P}|) \times dim}$  and  $b_2 \in \mathbb{R}^{|\mathcal{D}|+|\mathcal{P}|}$  are learnable parameters. Here,  $c \in \{0, 1\}^{|\mathcal{D}|+|\mathcal{P}|}$  represents the ground truth labels.  $c_j$  is set to 1 only if the corresponding label is present in the input sequence.

## 5.3 Fine-tune and Inference

After pre-training RAREMed on two tasks, we fine-tune the model to achieve accurate and fair medication recommendations. To predict the medication, we integrate a multi-label classification layer and utilize the patient representation as input:

$$\hat{o} = \sigma(W_3 r + b_3), \quad (5)$$

where  $\hat{o} \in [0, 1]^{|\mathcal{M}|}$  is the probability of medications being recommended.  $W_3 \in \mathbb{R}^{|\mathcal{M}| \times dim}$  and  $b_3 \in \mathbb{R}^{|\mathcal{M}|}$  are learnable parameters.

Following previous research works [29, 31, 38, 39], we fine-tune the entire model using the subsequent objective functions. First, we treat the prediction of each medication as an independent task and use the BCE loss for optimization:

$$L_{bce} = - \sum_{i=1}^{|\mathcal{M}|} [m_i \log(\hat{o}_i) + (1 - m_i) \log(1 - \hat{o}_i)]. \quad (6)$$

Additionally, we employ the multi-label margin loss to ensure the model assigns higher scores to the correct medications compared to the incorrect ones:

$$L_{multi} = \sum_{i,j:m_i=1,m_j=0} \frac{\max(0, 1 - (\hat{o}_i - \hat{o}_j))}{|\mathcal{M}|}. \quad (7)$$

To ensure medication safety, we use the DDI loss to penalize drug pairs with adverse interactions as described in [29, 39]:

$$L_{ddi} = \sum_{i=1}^{|\mathcal{M}|} \sum_{j=1}^{|\mathcal{M}|} A_{ij} \cdot \hat{o}_i \cdot \hat{o}_j. \quad (8)$$

The balance between accuracy and safety is achieved by combining the losses using a weighted sum:

$$L = (1 - \beta)((1 - \alpha)L_{bce} + \alpha L_{multi}) + \beta L_{ddi}, \quad (9)$$

where  $\alpha$  and  $\beta$  are hyperparameters balancing the contributions of the different losses.

During inference, medications with a probability greater than the threshold  $\delta = 0.5$  are recommended to the patient. Therefore, the final set of recommended medications  $\hat{\mathcal{Y}}$  can be defined as:

$$\hat{\mathcal{Y}} = \{i | \hat{o}_i > 0.5, 1 \leq i \leq |\mathcal{M}|\}. \quad (10)$$

## 6 EXPERIMENTS

In this section, we conduct comprehensive experiments to answer the following four questions:

- **RQ1:** How does the performance of the proposed RAREMed compare to that of existing medication recommendation methods?
- **RQ2:** Does the RAREMed effectively mitigate unfairness in medication recommendations?
- **RQ3:** How do the individual components of the RAREMed influence its performance with respect to accuracy and fairness?
- **RQ4:** What are the influential factors, e.g., the DDI Graph, that markedly impact recommendation performance of RAREMed?

### 6.1 Experimental Setup

**Datasets.** We used EHR data from two authentic EHR datasets, namely MIMIC-III [14], MIMIC-IV [13] and DDI data extracted from the TWOSIDES database [33]. Following previous works [29, 31, 39], we processed the datasets and randomly divided them into training, validation, and testing sets in a ratio of 4:1:1. The statistics of processed datasets are detailed in Table 1.

**Evaluation Protocol.** We evaluate the overall performance of all methods using widely accepted metrics [29, 31, 37, 39, 41], including Jaccard coefficient (Jaccard), Precision-Recall Area Under Curve (PRAUC), F1-score (F1), DDI rate (DDI), and the average number of recommended medicines (#MED), providing a comprehensive assessment of overall recommendation performance. It is worth noting that higher values of Jaccard, PRAUC, and F1 indicate improved accuracy, while a lower DDI value suggests enhanced medication safety. Furthermore, we believe that a successful medication recommender should resemble the behavior of doctors by recommending a similar number of medications, as reflected in the metric #MED.

Details of metrics are listed below. Building upon the formulation detailed in Section 3,  $\mathcal{M}$  represents the set of all medications in the dataset, while  $\mathbf{m}, \hat{\mathbf{m}} \in \{0, 1\}^{|\mathcal{M}|}$  signify the prescribed and predicted medications to a patient, respectively.

$$\text{Jaccard} = \frac{\{i : \mathbf{m}_i = 1\} \cap \{j : \hat{\mathbf{m}}_j = 1\}}{\{i : \mathbf{m}_i = 1\} \cup \{j : \hat{\mathbf{m}}_j = 1\}}.$$

$$F_1 = \frac{2R \times P}{R + P},$$

where the recall and precision are formulated as

$$R = \frac{\{i : \mathbf{m}_i = 1\} \cap \{j : \hat{\mathbf{m}}_j = 1\}}{\{i : \mathbf{m}_i = 1\}}, P = \frac{\{i : \mathbf{m}_i = 1\} \cap \{j : \hat{\mathbf{m}}_j = 1\}}{\{j : \hat{\mathbf{m}}_j = 1\}}.$$

$$\text{PRAUC} = \sum_{k=1}^{|\mathcal{M}|} P_k (R_k - R_{k-1}),$$

where  $P_k, R_k$  represent the precision and recall at cut-off  $k$ .

**Table 1: Statistics of processed data.**

Item	MIMIC-III	MIMIC-IV
# of visits / # of patients	14949/6344	16117/6352
dis. / proc. / med. space size	2026/630/112	3415/1072/123
avg. / max # of visits	4.96/29	7.30/47
avg. / max # of dis. per visit	13.76/39	12.66/39
avg. / max # of proc. per visit	4.45/28	2.53/28
avg. / max # of med. per visit	19.60/52	12.79/53

$$\text{DDI} = \frac{\sum_{l,k \in \{i: \hat{\mathbf{m}}_i = 1\}} A_{lk}}{\sum_{l,k \in \{i: \hat{\mathbf{m}}_i = 1\}} 1},$$

where  $A$  is the DDI graph defined in Section 3.

To examine the fairness dimension, we conduct a group-wise comparison by categorizing all patients into five groups, denoted as  $\{\mathcal{G}_1, \dots, \mathcal{G}_5\}$ , according to the occurrence frequency of the least common disease codes they exhibit.  $\mathcal{G}_1$  represents patients with the most commonly observed clinical codes, while  $\mathcal{G}_5$  comprises patients with the rarest codes. We calculate the average Jaccard coefficients for each group to assess the accuracy of recommendations for patients with different code occurrences. Additionally, we compute the standard deviation (denoted as  $\sigma$ ) of these five Jaccard coefficients to evaluate whether the models perform similarly across various groups, indicating fairness in the recommendations.

**Compared Methods.** We compare RAREMed against the following baseline algorithms:

- **LR** is a standard logistic regression technique, where inputs are represented as a multi-hot vector of length  $|\mathcal{D}| + |\mathcal{P}|$ .
- **LEAP** [44] is an instance-based method, which employs the LSTM model to generate medication sequence.
- **RETAIN** [6] is a longitudinal model that utilizes two-level neural attention mechanism to predict patients' future condition.
- **G-Bert** [28] integrates the GNN representation into transformer-based visit encoders, which is pre-trained on single-visit data.
- **GAMENet** [29] uses memory neural networks and graph convolutional networks to encode historical EHR data and DDI graph.
- **SafeDrug** [39] leverages the drug molecular graph and DDI graph to ensure the safety of medication recommendations.
- **COGNet** [37] uses a novel copy-or-predict mechanism which frames drug recommendation as a sequence generation problem.
- **MICRON** [38] is a recurrent residual learning model that focuses on the change of medications.
- **MoleRec** [41] models the dependencies between patient's health condition and molecular substructures.

In addition to the medication recommendation methods mentioned above, we also implement a baseline approach specifically designed for fair recommendation in group-wise comparison:

- **Rebalancing:** We employ a resampling technique [9, 15] to address the issue of data imbalance. Specifically, we calculate the Inverse Propensity Score (IPS) score [27, 40] for each disease code. Subsequently, we assign each patient a weight score based on the IPS score associated with the rarest disease. In each training epoch, we perform random sampling with replacement from the training data, the probability of selecting a particular patient is proportional to their assigned patient weight. We conduct experiment for this approach on our RAREMed without pre-training.



**Table 2: Overall performance comparison in the single-visit setting. RETAIN and G-Bert are excluded as they cannot generate recommendations for single-visit patients. Note that DDI rates of ground truth on MIMIC-III and MIMIC-IV are 0.082 and 0.081 respectively. The best and the runner-up results in each column are highlighted in bold and underlined, respectively.**

Method	MIMIC-III					MIMIC-IV				
	Jaccard	PRAUC	F1	DDI	#MED	Jaccard	PRAUC	F1	DDI	#MED
LR	0.4933	0.7639	0.6519	0.0786	16.85	0.4150	0.6782	0.5650	0.0736	9.85
LEAP	0.4526	0.6583	0.6154	0.0722	19.04	0.3907	0.5540	0.5437	0.0548	12.66
GAMENet	0.5210	0.7780	0.6762	0.0781	19.78	0.4401	0.6833	0.5933	0.0718	12.89
SafeDrug	0.5255	0.7732	0.6804	0.0688	20.84	<u>0.4560</u>	0.6858	<u>0.6098</u>	0.0689	14.07
COGNet	0.5109	0.6465	0.6676	0.0737	25.05	0.4313	0.6112	0.5850	0.0866	15.46
MICRON	0.5119	0.5190	0.6676	<u>0.0610</u>	20.94	0.4495	0.4353	0.6033	<b>0.0502</b>	14.50
MoleRec	<u>0.5303</u>	<u>0.7795</u>	<u>0.6844</u>	0.0692	21.30	0.4502	<u>0.6867</u>	0.6040	0.0699	14.05
<b>RAREMed</b>	<b>0.5414</b>	<b>0.7922</b>	<b>0.6942</b>	<b>0.0529</b>	19.66	<b>0.4625</b>	<b>0.7008</b>	<b>0.6158</b>	<u>0.0508</u>	12.51

**Table 3: Overall performance comparison under the multi-visit setting.**

Method	MIMIC-III					MIMIC-IV				
	Jaccard	PRAUC	F1	DDI	#MED	Jaccard	PRAUC	F1	DDI	#MED
LR	0.4933	0.7639	0.6519	0.0786	16.85	0.4150	0.6782	0.5650	0.0736	9.85
LEAP	0.4526	0.6583	0.6154	0.0722	19.04	0.3907	0.5540	0.5437	0.0548	12.66
RETAIN	0.4922	0.7560	0.6517	0.0792	24.76	0.4113	0.6543	0.5674	0.0839	16.95
G-Bert	0.5037	0.7631	0.6617	0.0832	24.46	0.4292	0.6746	0.5822	0.0794	16.19
GAMENet	0.5205	0.7767	0.6754	0.0774	19.52	0.4462	0.6923	0.5979	0.0807	12.60
SafeDrug	0.5232	0.7699	0.6781	0.0698	20.67	0.4566	0.6867	0.6080	0.0644	14.00
COGNet	0.5177	0.6360	0.6715	0.0757	24.95	0.4518	0.6359	0.6026	0.0846	15.70
MICRON	0.5188	0.6398	0.6735	<u>0.0658</u>	19.51	0.4555	0.5350	0.6081	<u>0.0528</u>	13.24
MoleRec	0.5293	<u>0.7786</u>	<u>0.6838</u>	0.0698	21.27	<u>0.4576</u>	0.6933	0.6097	<u>0.0683</u>	14.19
<b>RAREMed</b>	<b>0.5414</b>	<b>0.7922</b>	<b>0.6942</b>	<b>0.0529</b>	19.66	<b>0.4625</b>	<b>0.7008</b>	<b>0.6158</b>	<b>0.0508</b>	12.51

We omit the evaluation of the 4SDrug, DrugRec, and REFINE methods from our study due to their reliance on additional symptom information [31, 32] or unavailability of the source codes [3].

**Implementation Details.** The hyper-parameters of all models are determined based on their performance on the validation set. RAREMed is pre-trained on the training set in a sequential manner, first on the SMP task and then on the SP task, for 30 epochs respectively. The transformer encoder in our model consists of 3 layers with 4 attention heads, and the embedding dimension is set to 512. The weights of the loss function, denoted as  $\alpha$  and  $\beta$ , are set to 0.03 and 0.7, respectively. The parameters are trained using the AdamW optimizer [20] with a learning rate of  $1e-5$  and weight decay of 0.1.

## 6.2 Overall Performance Comparison (RQ1)

We conduct experiments on two settings for a thorough comparison: (1) **Single-visit setting**, where each visit is regarded as a separate patient without historical records as described in Section 3; (2) **Multi-visit setting**, where a patient may have multiple visits, and the historical records are available when recommending medications for the second and subsequent visits. This setting is designed for a fair comparison with longitudinal methods [29, 37–39, 41].

Based on the results presented in Table 2 and Table 3, we have the following observations:

- **Baseline Comparison.** LR and RETAIN, which were not originally designed for medication recommendation tasks, exhibit

poor performance. LR overlooks the interaction among clinical codes, while RETAIN fails to adequately capture the current patient condition. Similarly, LEAP and G-Bert, which do not consider the procedure sequence, also underperform in this domain. Furthermore, COGNet and MICRON utilize historical information in a specific manner, but their performance significantly declines in the single-visit setting. By integrating external knowledge such as EHR graph and drug molecule structures, GAMENet, SafeDrug and MoleRec achieve better performance. Interestingly, these three methods perform even worse in multi-visit setting compared to single-visit setting with MIMIC-III, highlighting the complexity of leveraging historical information.

- **Accuracy of RAREMed.** Our RAREMed surpasses all the baseline methods in terms of accuracy, exhibiting higher Jaccard coefficient, F1 score, and PRAUC on both the MIMIC-III and MIMIC-IV dataset. This superiority is observed not only in the single-visit scenario (Table 2) but also in the multi-visit scenario (Table 3). Remarkably, RAREMed achieves this improvement without relying on historical records. This highlights the efficacy of our medication recommendation framework, which can be attributed to the meticulous design of the expressive encoder and the effectiveness of the pre-training tasks employed.
- **Security of RAREMed.** RAREMed ensures the lowest DDI rate compared to the ground-truth level and baseline methods, with few exceptions. This showcases the system’s capability to effectively balance between accuracy and safety, a topic that will be further explored in Section 6.5.

**Table 4: Ablation study on MIMIC-IV dataset. Here,  $\sigma$  represents the standard deviation of Jaccard across five groups.**

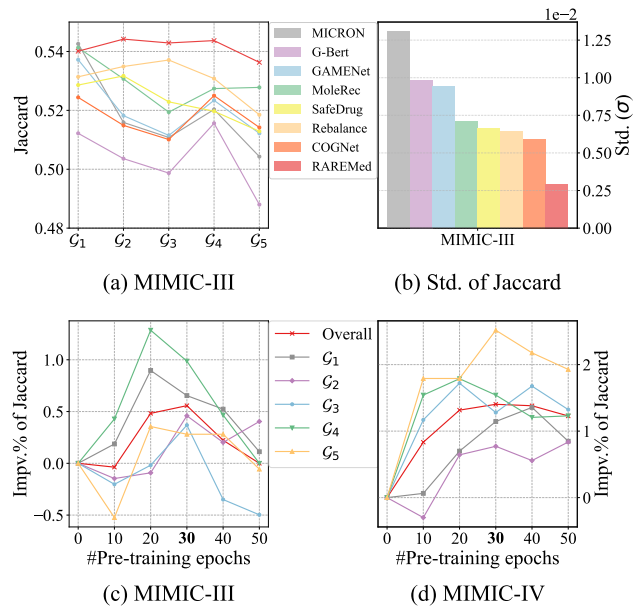
Model	Overall Performance					Group-wise Performance					
	Jaccard	PRAUC	F1	DDI	#MED	$\mathcal{G}_1$	$\mathcal{G}_2$	$\mathcal{G}_3$	$\mathcal{G}_4$	$\mathcal{G}_5$	$\sigma$
w/o P	0.4561	0.6939	0.6092	0.0509	12.94	0.4719	0.4666	0.4534	0.4475	0.4410	0.01157
w/o U	0.4589	0.6965	0.6122	0.0545	13.13	<b>0.4773</b>	<b>0.4707</b>	0.4558	0.4485	0.4423	0.01320
w/o s&r	0.4528	0.6892	0.6062	0.0554	13.15	0.4695	0.4646	0.4508	0.4451	0.4342	0.01286
<b>RAREMed</b>	<b>0.4625</b>	<b>0.7008</b>	<b>0.6158</b>	<b>0.0508</b>	12.51	<b>0.4773</b>	0.4702	<b>0.4592</b>	<b>0.4544</b>	<b>0.4521</b>	<b>0.00962</b>

### 6.3 Group-wise Performance Comparison (RQ2)

In this subsection, our objective is to examine model performance among patient groups with varying levels of disease prevalence. Figure 4 depicts the performance of RAREMed and baseline methods, from which the following observations can be made:

- **Baseline Methods Yield Unfair Recommendations.** Conventional medication recommendation methods exhibit a notable decrease in accuracy as the frequency of disease codes decreases, as illustrated in Figure 4(a). This leads to a higher standard deviation of Jaccard across different groups, as depicted in Figure 4(b). Notably, G-Bert [28], which utilizes pre-training techniques, struggles to achieve satisfactory accuracy for patients with rare diseases, potentially due to its disregard for procedure codes and subsequent failure to capture a comprehensive representation of patients' clinical information during pre-training. This underscores the superiority of our model structure and pre-training task design. Additionally, Rebalancing, tailored to address fairness concerns, also underperforms compared to RAREMed in terms of both fairness and overall performance.
- **RAREMed Effectively Addresses Fairness Concerns.** As depicted in Figure 4(a), RAREMed demonstrates superior performance for patients with rare diseases, while maintaining efficacy for the common group. Additionally, RAREMed showcases consistent performance across patient groups, as evidenced by minimal variance in Jaccard compared to the baseline methods. This trend highlights the strength of RAREMed in acquiring robust representations, particularly for rare clinical codes, and its capability to deliver fair medication recommendations for patients with rare diseases.
- **Pre-training Demonstrates Effectiveness on Biased Datasets.** The results presented in Figure 4(c) and (d) highlight the overall positive impact of our pre-training strategy on both common group and rare group, with few exceptions. Specifically, on MIMIC-III, where the disease distribution is less skewed, our pre-training strategy consistently enhances performance across various groups, as evidenced by the uniform improvement ratio observed after 30 pre-training epochs. In contrast, on MIMIC-IV, characterized by a more biased disease distribution and presenting greater challenges<sup>6</sup>, RAREMed notably improves the performance of the rarest group, effectively mitigating unfairness. The substantial improvement in overall accuracy on such a more biased dataset is also noteworthy.

<sup>6</sup>This is supported by the extensive distribution of long-tail disease codes in MIMIC-IV, as illustrated in Figure 1, and the greater variance of Jaccard scores across groups, as presented in Table 4, in comparison to the findings in Figure 4(b).



**Figure 4: (a) Jaccard coefficients on various patient groups. (b) Standard deviation of Jaccard coefficients across groups. (c) and (d) The improvement ratio of Jaccard across groups w.r.t. the number of pre-training epochs. The bold number indicates the selected hyper-parameter.**

### 6.4 Ablation Study (RQ3)

To evaluate the effectiveness of each component in RAREMed, we conduct ablation study with the following ablation models:

- **“w/o P”**, which removes two pre-training tasks, resulting in initializing parameters of encoder and embedding layers randomly.
- **“w/o U”**, which disables the unified encoder and learns two encoders for diseases and procedures separately and then concatenates the output. This setup retains two pre-training tasks.
- **“w/o s&r”**, it disables *segment* and *relevance* embedding layers, and patient representation relies solely on the *token* embedding.

We omit results on MIMIC-III, as they align with the same conclusion. We have the following observations from results in Table 4:

- **“w/o P”**. As anticipated, the diminished performance observed in the “w/o P” variant underscores the importance of our pre-training tasks. Notably, RAREMed outperforms “w/o P” across all five groups, affirming that our pre-training strategy not only enhances performance on rare diseases but also preserves the efficacy of the original model on other diseases. The enhancement seen in RAREMed can be attributed to the augmented encoder



obtained through the pre-training phase, enabling RAREMed to grasp intricate disease/procedure associations and produce more comprehensive patient representations.

- “w/o U”. RAREMed surpasses the “w/o U” variant, particularly in terms of fairness, further emphasizing the importance of capturing and comprehending the associations between diseases and procedures for fair medication recommendations.
- “w/o s&r”. RAREMed also outperforms the “w/o s&r” variant. Notably, the “w/o s&r” variant also exhibits significantly fluctuating performance across different patient groups. These results highlight the crucial role played by the segment and relevance embedding layers in capturing the complex relationships among diseases and procedures, particularly for rare ones.

## 6.5 Hyperparameter Studies (RQ4)

We conducted an investigation to understand the influence of hyperparameters on the efficacy of RAREMed on the MIMIC-IV. Specifically, we considered four hyperparameters: the multi-label margin loss weight ( $\alpha$ ), the DDI loss weight ( $\beta$ ), the number of pre-training epochs (#Epochs), and the embedding space dimension ( $dim$ ).

**Loss Weight.** Figure 5(a) demonstrates that the value of  $\alpha$  significantly affects the size of the recommended medicine set. Additionally, we observed that the recommendation accuracy initially increases and then declines as the number of recommended medications increases. This suggests that in addition to the ranking task commonly considered in classic recommendation systems, the size of recommended medication set is also an important factor in the field of medication recommendation.

On the other hand, Figure 5(b) illustrates the impact of  $\beta$  on the DDI rate and Jaccard coefficient. As  $\beta$  increases, the DDI rate consistently decreases, while the Jaccard increases briefly and then drops. This indicates that, overall, there is a trade-off relationship between accuracy and security, while uncontrollable DDI (when  $\beta = 0$ ) also harms accuracy. Notably, even when the DDI rate is restricted to less than 5% (when  $\beta = 0.75$ ), the accuracy of RAREMed remains relatively high, demonstrating its robustness.

**Number of Pre-training Epochs.** Figure 5(c) illustrates the impact of the number of pre-training epochs on the pre-training tasks and the medication recommendation task. It is evident that as the number of pre-training epochs increases, the model’s performance in pre-training tasks (SR task and SMP task) improves. This observation indicates that the model acquires knowledge about clinical code associations and gains comprehensive patient representations during pre-training, which are crucial for medication recommendation, particularly for patients with rare diseases. However, the improvement in the downstream task (Jaccard) starts to diminish after approximately 30 epochs. This phenomenon may be attributed to overfitting in the pre-training task and the inherent differences between the pre-training and downstream tasks. These findings highlight the significance of striking an optimal balance between the performance on pre-training tasks and the downstream medication recommendation task.

**Embedding Dimension.** As depicted in Figure 5(d), the increase in embedding dimension initially enhances RAREMed’s performance on both pre-training and downstream tasks, indicating an augmented expressive capacity. However, beyond a certain threshold,

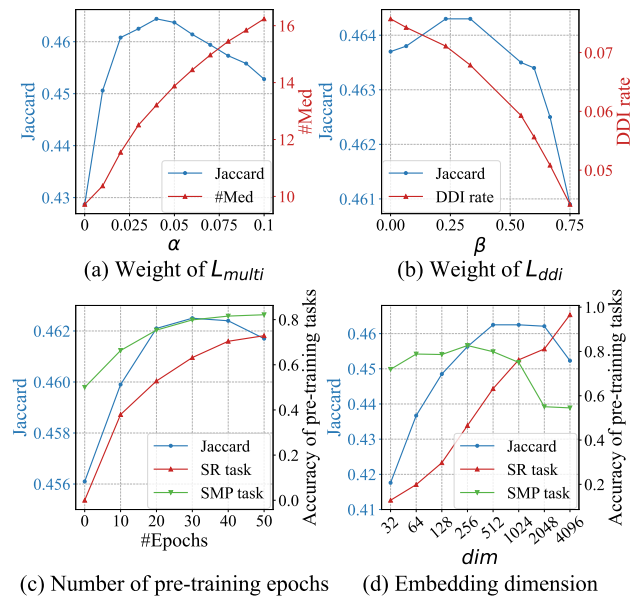


Figure 5: Hyperparameter effects on model performance.

the model’s performance on SMP and downstream tasks begins to decline, note that RAREMed is pre-trained first on the SMP task and then on the SR task, excluding the influence of the SR task on the SMP task. This decline in performance can be attributed to training challenges such as gradient vanishing or explosion, which disrupt convergence and hinder the model’s stability.

## 7 CONCLUSION

In this paper, we propose RAREMed, a novel medication recommendation model designed to address the fairness issue in medication recommendation systems. By focusing on enhancing accuracy for patients with rare diseases, RAREMed leverages two self-supervised pre-training tasks to learn specialized medication needs and interrelations. The model also utilizes a unified input sequence approach to capture complex relationships among disease and procedure codes. Experimental results on real-world datasets demonstrate the effectiveness of RAREMed in providing accurate drug recommendations for both rare and common disease patients, thereby mitigating unfairness.

In future research, we intend to further explore the utilization of patients’ historical records and additional clinical information, such as demographics, lab data, notes, and imaging, to improve the accuracy of medication recommendations. Additionally, incorporating more external knowledge during the pre-training phase and exploring transfer learning across different electronic health record (EHR) datasets are also promising avenues for further development.

## ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China (62272437, 62302321) and the University Synergy Innovation Program of Anhui Province (GXXT-2023-071).

## REFERENCES

- [1] Zafar Ali, Yi Huang, Irfan Ullah, Junlan Feng, Chao Deng, Nimbeshaho Thierry, Asad Khan, Asim Ullah Jan, Xiaoli Shen, Wu Rui, et al. 2023. Deep learning for medication recommendation: a systematic survey. *Data Intelligence* 5, 2 (2023), 303–354.
- [2] Philip Aspden and Philip Aspden. 2007. *Preventing medication errors*. National Acad. Press.
- [3] Suman Bhoi, Mong-Li Lee, Wynne Hsu, and Ngiap Chuan Tan. 2023. REFINE: A Fine-Grained Medication Recommendation System Using Deep Learning and Personalized Drug Interaction Modeling. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [4] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems* 41, 3 (2023), 1–39.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [6] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. 2016. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems* 29 (2016).
- [7] Yon Cui, Wenzhou Cao, Hua Shen, Jianjun Xie, Tamara S Adams, Yuanjuan Zhang, and Qiang Shao. 2014. Comparison of ESWL and ureteroscopic holmium laser lithotripsy in management of ureteral stones. *PloS one* 9, 2 (2014), e87634.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 4171–4186.
- [9] Michael D Ekstrand, Mucun Tian, Ion Madrazo Azpiazu, Jennifer D Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All the cool kids, how do they fit in?: Popularity and demographic biases in recommender evaluation and effectiveness. In *Conference on fairness, accountability and transparency*. PMLR, 172–186.
- [10] Marilyn J Field, Thomas F Boat, et al. 2010. Profile of rare diseases. *Rare diseases and orphan products: Accelerating research and development* (2010).
- [11] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th acm sigkdd international conference on knowledge discovery & data mining*. 2221–2231.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), 770–778.
- [13] Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data* 10, 1 (2023), 1.
- [14] 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. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3, 1 (2016), 1–9.
- [15] Faisal Kamiran and Toon Calders. 2009. Classifying without discriminating. In *2009 2nd international conference on computer, control and communication*. IEEE, 1–6.
- [16] Ellen Kim, Samuel M Rubinstein, Kevin T Nead, Andrzej P Wojcieszynski, Peter E Gabriel, and Jeremy L Warner. 2019. The evolving use of electronic health records (EHR) for research. In *Seminars in radiation oncology*, Vol. 29. Elsevier, 354–361.
- [17] Jason Lazarou, Bruce H Pomeranz, and Paul N Corey. 1998. Incidence of adverse drug reactions in hospitalized patients: a meta-analysis of prospective studies. *Jama* 279, 15 (1998), 1200–1205.
- [18] Roger Zhe Li, Julián Urbano, and Alan Hanjalic. 2021. Leave no user behind: Towards improving the utility of recommender systems for non-mainstream users. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 103–111.
- [19] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. 2022. Fairness in recommendation: A survey. *arXiv preprint arXiv:2205.13619* (2022).
- [20] Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).
- [21] Aya A Mitani and Sebastian Haneuse. 2020. Small data challenges of studying rare diseases. *JAMA network open* 3, 3 (2020), e201965–e201965.
- [22] Tobias Mueller, Andreas Jerrentrup, Max Jakob Bauer, Hans Walter Fritsch, and Juergen Rolf Schaefer. 2016. Characteristics of patients contacting a center for undiagnosed and rare diseases. *Orphanet Journal of Rare Diseases* 11, 1 (2016), 1–9.
- [23] Zhaopeng Qiu, Xian Wu, Jingyue Gao, and Wei Fan. 2021. U-BERT: Pre-training user representations for improved recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4320–4327.
- [24] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [25] Bashir Rastegarpanah, Krishna P Gummadi, and Mark Crovella. 2019. Fighting fire with fire: Using antidote data to improve polarization and fairness of recommender systems. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 231–239.
- [26] Tina Roblek, Tomaz Vaupotic, Ales Mrhar, and Mitja Lainscak. 2015. Drug-drug interaction software in clinical practice: a systematic review. *European journal of clinical pharmacology* 71 (2015), 131–142.
- [27] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations as treatments: Debiasing learning and evaluation. In *international conference on machine learning*. PMLR, 1670–1679.
- [28] Junyuan Shang, Tengfei Ma, Cao Xiao, and Jimeng Sun. 2019. Pre-training of graph augmented transformers for medication recommendation. In *28th International Joint Conference on Artificial Intelligence, IJCAI 2019*. International Joint Conferences on Artificial Intelligence.
- [29] Junyuan Shang, Cao Xiao, Tengfei Ma, Hongyan Li, and Jimeng Sun. 2019. Gamenet: Graph augmented memory networks for recommending medication combination. In *proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1126–1133.
- [30] Attasit Srisubat, Somkiat Potisat, Bannakij Lojanapiwat, Vasun Setthawong, and Malinee Laopaiboon. 2014. Extracorporeal shock wave lithotripsy (ESWL) versus percutaneous nephrolithotomy (PCNL) or retrograde intrarenal surgery (RIRS) for kidney stones. *Cochrane database of systematic reviews* 11 (2014).
- [31] Hongda Sun, Shufang Xie, Shuqi Li, Yuhang Chen, Ji-Rong Wen, and Rui Yan. 2022. Debaised, Longitudinal and Coordinated Drug Recommendation through Multi-Visit Clinic Records. *Advances in Neural Information Processing Systems* 35 (2022), 27837–27849.
- [32] Yanhao Tan, Chengjun Kong, Leisheng Yu, Pan Li, Chaochao Chen, Xiaolin Zheng, Vicki S Hertzberg, and Carl Yang. 2022. 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*. 3970–3980.
- [33] Nicholas P Tatonetti, Patrick P Ye, Roxana Daneshjou, and Russ B Altman. 2012. Data-driven prediction of drug effects and interactions. *Science translational medicine* 4, 125 (2012), 125ra31–125ra31.
- [34] Shanshan Wang, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Jun Ma, and Maarten de Rijke. 2019. Order-free medicine combination prediction with graph convolutional reinforcement learning. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1623–1632.
- [35] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A survey on the fairness of recommender systems. *ACM Transactions on Information Systems* 41, 3 (2023), 1–43.
- [36] Chi-Man Wong, Fan Feng, Wen Zhang, Chi-Man Vong, Hui Chen, Yichi Zhang, Peng He, Huan Chen, Kun Zhao, and Huajun Chen. 2021. Improving conversational recommender system by pretraining billion-scale knowledge graph. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2607–2612.
- [37] Rui Wu, Zhaopeng Qiu, Jiacheng Jiang, Guilin Qi, and Xian Wu. 2022. Conditional generation net for medication recommendation. In *Proceedings of the ACM Web Conference 2022*. 935–945.
- [38] Chaoqi Yang, Cao Xiao, Lucas Glass, and Jimeng Sun. 2021. Change Matters: Medication Change Prediction with Recurrent Residual Networks. In *30th International Joint Conference on Artificial Intelligence, IJCAI 2021*. International Joint Conferences on Artificial Intelligence, 3728–3734.
- [39] Chaoqi Yang, Cao Xiao, Fenglong Ma, Lucas Glass, and Jimeng Sun. 2021. SafeDrug: Dual Molecular Graph Encoders for Recommending Effective and Safe Drug Combinations. In *30th International Joint Conference on Artificial Intelligence, IJCAI 2021*. International Joint Conferences on Artificial Intelligence, 3735–3741.
- [40] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge Belongie, and Deborah Estrin. 2018. Unbiased offline recommender evaluation for missing-not-at-random implicit feedback. In *Proceedings of the 12th ACM conference on recommender systems*. 279–287.
- [41] Nianzu Yang, Kaipeng Zeng, Qitian Wu, and Junchi Yan. 2023. Molerec: Combinatorial drug recommendation with substructure-aware molecular representation learning. In *Proceedings of the ACM Web Conference 2023*. 4075–4085.
- [42] Sirui Yao and Bert Huang. 2017. Beyond parity: Fairness objectives for collaborative filtering. *Advances in neural information processing systems* 30 (2017).
- [43] Zheni Zeng, Chaojun Xiao, Yuan Yao, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, and Maosong Sun. 2021. Knowledge transfer via pre-training for recommendation: A review and prospect. *Frontiers in big Data* 4 (2021), 602071.
- [44] Yutao Zhang, Robert Chen, Jie Tang, Walter F Stewart, and Jimeng Sun. 2017. LEAP: learning to prescribe effective and safe treatment combinations for multi-morbidity. In *proceedings of the 23rd ACM SIGKDD international conference on knowledge Discovery and data Mining*. 1315–1324.