



Natural Language-Assisted Multi-modal Medication Recommendation

Jie Tan
jtan@se.cuhk.edu.hk
The Chinese University of Hong Kong
Hong Kong, China

Yu Rong*
yu.rong@hotmail.com
DAMO Academy, Alibaba Group
Hupan Lab
Hangzhou, China

Kangfei Zhao*
zkf1105@gmail.com
Beijing Institute of Technology
Beijing, China

Tian Bian
tbian@se.cuhk.edu.hk
The Chinese University of Hong Kong
Hong Kong, China

Tingyang Xu
xuty_007@hotmail.com
DAMO Academy, Alibaba Group
Hupan Lab
Hangzhou, China

Junzhou Huang
jzhuang@uta.edu
University of Texas at Arlington
Arlington, Texas, United States

Hong Cheng
hcheng@se.cuhk.edu.hk
The Chinese University of Hong Kong
Hong Kong, China

Helen Meng
hmmeng@se.cuhk.edu.hk
The Chinese University of Hong Kong
Hong Kong, China

ABSTRACT

Combinatorial medication recommendation (CMR) is a fundamental task of healthcare, which offers opportunities for clinical physicians to provide more precise prescriptions for patients with intricate health conditions, particularly in the scenarios of long-term medical care. Previous research efforts have sought to extract meaningful information from electronic health records (EHRs) to facilitate combinatorial medication recommendations. Existing learning-based approaches further consider the chemical structures of medications, but ignore the textual medication descriptions in which the functionalities are clearly described. Furthermore, the textual knowledge derived from the EHRs of patients remains largely underutilized. To address these issues, we introduce the Natural Language-Assisted Multi-modal Medication Recommendation (*NLA-MMR*), a multi-modal alignment framework designed to learn knowledge from the patient view and medication view jointly. Specifically, *NLA-MMR* formulates CMR as an alignment problem from patient and medication modalities. In this vein, we employ pretrained language models (PLMs) to extract in-domain knowledge regarding patients and medications, serving as the foundational representation for both modalities. In the medication modality, we exploit both chemical structures and textual descriptions to create medication representations. In the patient modality, we generate the patient representations based

on textual descriptions of diagnosis, procedure, and symptom. Extensive experiments conducted on three publicly accessible datasets demonstrate that *NLA-MMR* achieves new state-of-the-art performance, with a notable average improvement of 4.72% in Jaccard score.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; • **Information systems** → **Data mining**; • **Applied computing** → **Health informatics**.

KEYWORDS

combinatorial medication recommendation, multi-modal alignment, pretrained language models

ACM Reference Format:

Jie Tan, Yu Rong, Kangfei Zhao, Tian Bian, Tingyang Xu, Junzhou Huang, Hong Cheng, and Helen Meng. 2024. Natural Language-Assisted Multi-modal Medication Recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3627673.3679529>

1 INTRODUCTION

The implementation of digital patient records in healthcare institutions has boosted the development of a substantial repository of information known as Electronic Health Records (EHRs). This valuable asset holds significant potential for a variety of applications within the medical field, such as predicting mortality rates, forecasting treatment outcomes, and offering medication recommendations. As one of the fundamental tasks of healthcare, combinatorial medication recommendation (CMR) has attracted a lot of attention in recent years, and achieved significant progress in providing personalized and safe medication recommendations for patients. This advancement is especially beneficial for the elderly who have been suffering from chronic illnesses for a prolonged time [27, 30].

*Co-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.
CIKM '24, October 21–25, 2024, Boise, ID, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0436-9/24/10
<https://doi.org/10.1145/3627673.3679529>

A multitude of deep learning models have been designed to solve the CMR task in the literature, which not only use the EHRs (e.g., diagnoses, procedures, and symptoms), but also leverage the chemical structures of drugs. For example, *SafeDrug* [38] developed dual molecular graph encoders to embed global and local molecular structures. *MoleRec* [39] proposed a molecular substructure-aware attentive method for CMR. *DrugRec* [31] proposed to represent drugs with the SMILES, a line-based representation (i.e., string) of molecules. Besides EHRs and chemical structures, we observe there are a vast amount of textual descriptions of drug molecules that are human-understandable and easily accessible on platforms such as PubChem [13] and DrugBank [35]. These texts describe the functionalities of drug molecules and reveal their therapeutic purposes. Unfortunately, none of the existing methods try to exploit the textual medication descriptions, thus failing to establish the semantic relations between the textual descriptions of patients and medications. Consider the example depicted in Figure 1 with two modalities of patient and medication, respectively. In the patient modality, an EHR shows that the patient with the symptoms “heart pain” and “fatigue” is diagnosed with diseases “old myocardial infarction” and “chronic ischemic”, and then is prescribed with a set of drugs “carvedilol”, “aspirin” and “acetaminophen”. In the medication modality, the textual medication description reveals that “Carvedilol is used to treat chronic heart failure, hypertension, and myocardial infarction”. This example illustrates that the textual medication description can supplement the chemical structures to enhance the expressiveness of medication representations, and help establish the semantic relations between the patients and medications. This motivates us to leverage the useful textual medication description and consider patient and medication as two distinct modalities for medication recommendation.

In another direction, recent studies have attempted to incorporate structured domain knowledge, such as biomedical knowledge graphs (BioKG) [29, 36] and drug-drug interactions (DDI) [2, 31, 38] to enhance the CMR performance. However, the quality of such structured domain knowledge may degrade due to potential information loss and biased sampling when it is extracted from raw medical documents. A more effective approach is to employ the pre-trained language models (PLMs), which are pretrained on web-scale textual data from the chemical and biological domain to capture clinical specialist knowledge in raw texts for medication recommendation. Representations learned by PLMs like *PubMedBERT* [6] and *ClinicalBERT* [16] achieve strong performance across many domain-specific applications. Intuitively, using these PLMs to encode the textual knowledge of the patient and medication modalities can provide additional gains over conventional encoding methods in the CMR task because they effectively fuse various types of information, such as diagnoses, procedures, symptoms and medications in a unified global semantic space.

In light of these observations, we propose a novel model, called Natural Language-Assisted Multi-modal Medication Recommendation (*NLA-MMR*), to make medication recommendations by effectively harnessing the wealth of clinical knowledge in the EHRs and drugs. *NLA-MMR* incorporates the textual medication descriptions for molecular representation learning. It also integrates expert knowledge by leveraging PLMs to capture various types of semantic relations between the texts in the EHRs and medication descriptions.

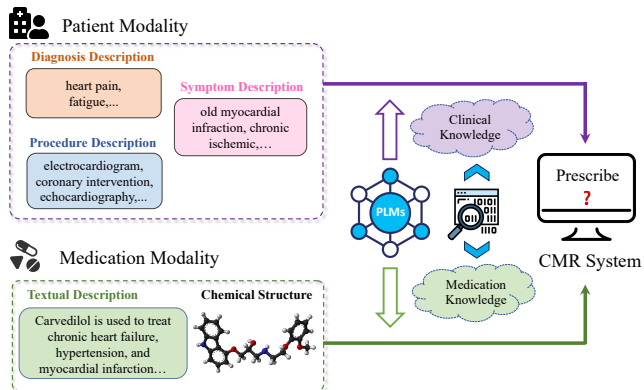


Figure 1: Illustration of combinatorial medication recommendation with help of knowledge extraction from PLMs

Specifically, *NLA-MMR* aligns patient and medication modalities utilizing a cross-modal alignment module based on the patient’s prescriptions in the EHRs. In the prediction phase, *NLA-MMR* transforms a recommendation problem into a similarity retrieval problem that aims to find the most similar medication representation given the patient’s representation. Our contributions in this paper are summarized as follows:

- To the best of our knowledge, we are the first to consider the patient and medication as two distinct modalities and design a multi-modal alignment framework *NLA-MMR* to learn their representations in a unified latent space. *NLA-MMR* converts the recommendation task into a similarity retrieval problem, enabling better generalizability and transferability of the model.
- We propose utilizing the PLMs as a fundamental building block to construct the representations of the patient and medication modalities by extracting expert knowledge from their descriptions. For the medication modality, we incorporate both chemical structures and medication descriptions for the medication representation learning.
- Experimental results on three publicly available healthcare datasets demonstrate that our proposed method is highly effective compared to the state-of-the-art methods. Our model boosts Jaccard by 2.86%, 7.01%, 4.29% on MIMIC-III, MIMIC-IV, and eICU datasets, respectively.

2 PRELIMINARIES

For a patient v , his/her electrical health records (EHRs) are represented as a sequence $R_v = [x_v^{(1)}, x_v^{(2)}, \dots, x_v^{(T_v)}]$, where $x_v^{(t)}$ represents the health record from the t -th clinical visit and T_v is the total number of visits of the patient v . When the context is clear, we omit the subscript v to simplify notations. Then, the t -th visit in the EHR of a patient can be represented by a $x^{(t)} = (d^{(t)}, p^{(t)}, s^{(t)}, m^{(t)})$, where $d^{(t)} \in \{0, 1\}^{|\mathcal{D}|}$, $p^{(t)} \in \{0, 1\}^{|\mathcal{P}|}$, $s^{(t)} \in \{0, 1\}^{|\mathcal{S}|}$, $m^{(t)} \in \{0, 1\}^{|\mathcal{M}|}$ are multi-hot diagnosis, procedure, symptom, and medication vectors, respectively. \mathcal{D} , \mathcal{P} , \mathcal{S} , \mathcal{M} indicate the set of possible diagnoses, procedures, symptoms and medications, respectively, and $|\cdot|$ denotes the cardinality of a set. As shown in Figure 1, our approach treats the patient and medication as two modalities to consider their information.

Patient Modality. Following [31], we model the patient modality as a series of clinical visits containing the diagnosis, procedure, or symptom, which reflect the health state of this patient. Thus, the t -th visit for the patient v is defined as $r_v^{(t)} = (r_d^{(t)}, r_p^{(t)}, r_s^{(t)})$, where $r_d^{(t)}$, $r_p^{(t)}$, and $r_s^{(t)}$ represent the textual description of diagnoses in $d^{(t)}$, procedures in $p^{(t)}$, and symptoms in $s^{(t)}$, respectively. Therefore, the patient modality of v is denoted as $v = [r_v^{(t)}]_{t=1}^{T_v}$.

Medication Modality. To enrich the medication modality, we collect the textual description r for each molecule in \mathcal{M} , which provides a high-level explanation of the molecule’s functionality. Additionally, for each medication in \mathcal{M} , we represent its molecular structure by an undirected graph $g = (\mathcal{A}, \mathcal{B})$. Here, $\mathcal{A} = \{a_1, \dots, a_n\}$ is the set of atoms and $\mathcal{B} \subseteq \mathcal{A} \times \mathcal{A}$ is the set of bonds. The neighborhood of atom a_i is denoted as $\mathcal{N}(a_i) = \{a_j | (a_i, a_j) \in \mathcal{B}\}$. Therefore, the medication modality is represented as (r, g) . Figure 1 depicts an example of medication “Carvedilol”.

Combinatorial Medication Recommendation. Given the EHR of a patient till time $t - 1$, denoted as $[x^{(1)}, x^{(2)}, \dots, x^{(t-1)}]$, containing the textual description of the t -th diagnosis, procedure, and symptom sets, denoted as $r_d^{(t)}$, $r_p^{(t)}$ and $r_s^{(t)}$, respectively, and a candidate medication set \mathcal{M} , the CMR task aims to predict an appropriate combination of medications $\mathcal{M}_t \subseteq \mathcal{M}$ for the patient’s t -th visit using a neural network model $f(\cdot)$.

3 METHODOLOGY

3.1 Overview

Figure 2 illustrates an overview of our proposed drug recommendation approach *NLA-MMR*, which is composed of three modules: 1) the **patient representation module** using a PLM to encode the texts of diagnosis, procedure, and symptom; 2) the **medication representation module** generating medication representations by employing the GNN model and a PLM for incorporating molecular structure-level information and textual knowledge, respectively; 3) the **cross-modal alignment module** responsible for making prescriptions by using two projection layers to calculate the similarity between the patient representation and medication representations on the joint latent space. In particular, in the patient representation module, we conduct cross-attention fusion to fully leverage the different types of textual information present in the patient modality. Besides, we integrate the medication information of historical visits to capture dependencies between medication usage in previous visits and the current visit.

3.2 Patient Representation Module

In this section, we construct the patient representation with the textual components defined in the patient modality. By leveraging the power of PLMs, we are able to extract valuable expert knowledge from these unprocessed texts, thereby contributing to the development of an advanced medication recommendation model. Specifically, we first consider the PLM-based representation of the concatenated text of diagnosis, procedure, and symptom as the basic patient representation. Then, we devise a cross-attention fusion mechanism to merge information from the three types of patient descriptions to obtain the final patient representation.

PLM-based Representation. We employ a PLM encoder [17] to extract the representation for the patient modality. Give a patient’s t -th visit $r^{(t)} = (r_d^{(t)}, r_p^{(t)}, r_s^{(t)})$, we concatenate these descriptions into a synthesis text as $r_{dps}^{(t)}$ and employ the PLM encoder to obtain its representation:

$$\mathbf{h}_{dps}^{(t)} = \text{PLM-Encoder} \left(r_{dps}^{(t)} \right), \quad (1)$$

where $\mathbf{h}_{dps}^{(t)} \in \mathbb{R}^{d_{enc}}$ is the output of the PLM encoder, which is obtained via an average pooling of the embeddings of all tokens. Meanwhile, we individually encode the textual descriptions of diagnosis, procedure, and symptom with the PLM and obtain $\mathbf{h}_d^{(t)}$, $\mathbf{h}_p^{(t)}$, and $\mathbf{h}_s^{(t)}$, which are representations of the corresponding type of textual description in the patient modality.

Cross-Attention Fusion. To enhance the basic patient representation in Eq. (1), we aggregate the features derived from three aspects, $\mathbf{h}_d^{(t)}$, $\mathbf{h}_p^{(t)}$, and $\mathbf{h}_s^{(t)}$, by utilizing a cross-attention mechanism, weighted by a set of learnable weights $w_c \in \mathbb{R}^{|C|}$, as $\mathbf{h}_{attn}^{(t)} = \sum_{c \in C} w_c \mathbf{h}_c^{(t)}$, $C = \{s, d, p\}$. Let $\mathbf{H}^{(t)} \in \mathbb{R}^{|C| \times d_{enc}}$ be the matrix stacked by the embeddings in $\{\mathbf{h}_d^{(t)}, \mathbf{h}_p^{(t)}, \mathbf{h}_s^{(t)}\}$. In Eq. (2), $\mathbf{H}_1 \in \mathbb{R}^{|C| \times d_K}$, $\mathbf{H}_2 \in \mathbb{R}^{1 \times d_K}$ are transformed by linear weight matrices $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{d_{enc} \times d_K}$, respectively. w_c is computed by the multiplication of \mathbf{H}_1 and the transpose of \mathbf{H}_2 followed by a softmax function that normalizes the weights, as Eq. (3) shows.

$$\mathbf{H}_1 = \mathbf{H}^{(t)} \mathbf{W}_1, \quad \mathbf{H}_2 = \mathbf{h}_{dps}^{(t)} \mathbf{W}_2, \quad (2)$$

$$\{w_c\}_{c \in C} = \text{softmax} \left(\frac{\mathbf{H}_1 \mathbf{H}_2^T}{\sqrt{d_K}} \right). \quad (3)$$

Next, we concatenate the aggregated representation $\mathbf{h}_{attn}^{(t)} \in \mathbb{R}^{d_{enc}}$ and the basic patient feature $\mathbf{h}_{dps}^{(t)} \in \mathbb{R}^{d_{enc}}$ as the patient representation, which can be formulated as:

$$\mathbf{p}^{(t)} = \left[\mathbf{h}_{attn}^{(t)} \parallel \mathbf{h}_{dps}^{(t)} \right] \mathbf{W}_r + \mathbf{b}_r, \quad (4)$$

where \parallel denotes the concatenation operation and $\mathbf{W}_r \in \mathbb{R}^{2d_{enc} \times d_{enc}}$ and $\mathbf{b}_r \in \mathbb{R}^{d_{enc}}$ are trainable parameters.

3.3 Medication Representation Module

The medication modality encompasses two essential components, namely textual knowledge and structural information. As depicted in Figure 2, our module for representing medications is composed of two branches: *the textual description branch* and *the chemical structure branch*. These branches are designed to address the external domain knowledge of molecules through textual descriptions and their intrinsic property through chemical structures, respectively. To harness the domain-specific knowledge contained within the textual descriptions, we employ a text encoder known as the PLM. Additionally, we employ a GNN model to extract representations at the structure level for drugs based on their chemical structures. **PLM-based Functional Representation.** The textual description of a drug provides a high-level overview of the molecule’s functionality, which offers valuable insights into the therapeutic potential of the drug and can aid in drug predictions. In this paper, we extract therapeutic descriptions of the drugs from DrugBank [35], which illustrates drug effects on the target organism. Given the




Figure 2: The architecture of NLA-MMR. NLA-MMR is composed of three modules: (a) Patient Representation Module employs PLMs as the base encoder to construct the representation of patient modality from the diagnosis, procedure, and symptom. (b) Medication Representation Module incorporates the embedding derived from textual medication descriptions and chemical structures to build the representation of medication modality. (c) Cross-Modal Alignment Module takes the representation from patient and medication modality as input and aligns them in the same latent space. We further consider historical medication usage information to model the patient’s clinical history, which can enhance the representation of patient modality.

textual medication descriptions, we obtain its PLM-based functional representations by conducting a similar process as in Eq. (1). We collect the textual embeddings of all drugs into a drug matrix $E_f \in \mathbb{R}^{|\mathcal{M}| \times d_{enc}}$, where each row corresponds to the textual embedding of a drug.

Structure-level Representation. We incorporate chemical structures to learn the molecule representation. Specifically, we employ the GNN to model the interactions between all the atoms across the single molecule graph g and obtain the structure-level drug representation. The aggregation and combination paradigm is conducted on the graph g of a drug molecule whose nodes are atoms and edges are chemical bonds. Given the features of the n atoms $E_a^{(k)} = \{e_{a_1}^{(k)}, \dots, e_{a_n}^{(k)}\} \in \mathbb{R}^{n \times d_s}$, for each atom $a_i \in \mathcal{A}$, the messages of all the other atoms in the molecule are aggregated to $e_i^{(k)} \in \mathbb{R}^{d_s}$ (Eq. (5)). Afterward, the aggregated message $e_i^{(k)}$ is combined with the feature of the atom $e_{a_i}^{(k)}$ by a summation operation, followed by a multilayer perceptron (MLP), as the combination function of GIN [28], generating new atom feature $e_{a_i}^{(k+1)}$ for the next layer (Eq. (6)).

$$e_i^{(k)} = \text{AGG}^{(k)} \left(\left\{ e_{a_j}^{(k)} \mid \forall a_j \in \mathcal{N}(a_i) \right\} \right), \quad (5)$$

$$e_{a_i}^{(k+1)} = \text{MLP}^{(k)} \left(\left(1 + \epsilon^{(k+1)} \right) e_{a_i}^{(k)} + e_i^{(k)} \right), \forall i \in \{1, \dots, n\}, \quad (6)$$

where ϵ can be a learnable parameter or a fixed scalar, it determines the importance of the target node compared to its neighbors, $e_{a_j}^{(k)}$ is the feature vector of atom a_j at the k -th layer, $\mathcal{N}(a_i)$ denotes the neighbors of the atom a_i . After applying message passing for

L layers, the embeddings of the atoms in the molecule graph are aggregated into a global structural-level representation as Eq. (7),

$$e_g = \text{Pooling} \left(\left\{ e_{a_i}^{(L)} \mid \forall i \in \{1, \dots, n\} \right\} \right), \quad (7)$$

where the pooling function is the average function. We employ the same GNN encoder with shared parameters for each drug molecule (with a total of $|\mathcal{M}|$ different drugs). We collect the GNN embeddings of all drugs into a drug memory $E_s \in \mathbb{R}^{|\mathcal{M}| \times d_s}$, where each row corresponds to the drug’s structural representation, d_s represents the dimension of the structure-level drug representation.

The textual representation and structural representation are combined to form the final medication embedding matrix $E_m \in \mathbb{R}^{|\mathcal{M}| \times (d_s + d_{enc})}$ as follows:

$$E_m = [E_s \parallel E_f]. \quad (8)$$

3.4 Cross-Modal Alignment Module

In the cross-modal alignment module, the patient representation $p^{(t)}$ and medication representations E_m are fed into two MLPs which map the representations extracted from different modalities to a joint latent space as Eq. (9):

$$\tilde{p}^{(t)} = \text{MLP}(p^{(t)}), \tilde{E}_m = \text{MLP}(E_m). \quad (9)$$

Subsequently, we use inner product operation \odot to compute the similarity of the patient feature $\tilde{p}^{(t)}$, and the medication features \tilde{E}_m in the joint space as Eq. (10):

$$\hat{m}_i^{(t)} = \sigma \left(\tilde{p}^{(t)} \odot \left(\tilde{E}_m \right)_i \right), \forall i \in \{1, \dots, |\mathcal{M}|\}. \quad (10)$$

The predicted recommendation probabilities $\hat{m}^{(t)} \in [0, 1]^{|M|}$ are obtained by a sigmoid function σ .

3.5 Historical Information Integration

Within the realm of medication recommendation, it is of utmost importance to effectively model the historical clinical information of patients, particularly those suffering from chronic ailments that necessitate long-term medication. Prior studies commonly employ the extraction of historical features, such as diagnosis, procedure, and symptom, to capture the patient’s clinical trajectory [38, 39]. Nevertheless, the observed patient data often suffers from incompleteness and inadequacy, subsequently resulting in imprecise predictions in the CMR task [31]. To overcome this limitation, we propose a historical information integration mechanism that directly assimilates the patient’s medication history. This module generates a comprehensive representation of the patient’s historical clinical condition by duly considering the temporal patterns of medication usage. By leveraging this integrated information, we aim to enhance the accuracy and precision of medication recommendations.

Firstly, we denote $r_m^{(t)}$ as the textual description of the medications used in the t -th visit and obtain the medication representations of all the past visits $\mathbf{h}_m^{(i)} \in \mathbb{R}^{d_{enc}}$ for the visit $i \in \{1, \dots, t-1\}$ through the similar process as in Eq. (1):

$$\mathbf{h}_m^{(i)} = \text{PLM-Encoder}\left(r_m^{(i)}\right), \forall i \in \{1, \dots, t-1\}. \quad (11)$$

Then, we use the patient representation in Eq. (4), $\mathbf{p}^{(t)}$, as the query vector to calculate the selection score $\{w_m^{(i)} | i \in 1, \dots, t-1\}$ along the historical medication combination features. Let $\mathbf{H}_m \in \mathbb{R}^{(t-1) \times d_{enc}}$ be the matrix stacked by the historical medication embeddings $\{\mathbf{h}_m^{(i)} | i \in \{1, \dots, t-1\}\}$, we calculate w_m as follows:

$$\mathbf{H}_3 = \mathbf{H}_m \mathbf{W}_3, \quad \mathbf{H}_4 = \mathbf{p}^{(t)} \mathbf{W}_4, \quad (12)$$

$$\{w_m^{(i)}\}_{i \in \{1, \dots, t-1\}} = \text{softmax}\left(\frac{\mathbf{H}_3 \mathbf{H}_4^T}{\sqrt{d_K}}\right), \quad (13)$$

where $\mathbf{H}_3 \in \mathbb{R}^{(t-1) \times d_K}$, $\mathbf{H}_4 \in \mathbb{R}^{1 \times d_K}$, and $\mathbf{W}_3, \mathbf{W}_4 \in \mathbb{R}^{d_{enc} \times d_K}$ are learnable parameters. Finally, the historical medication representation of the patient can be computed as follows:

$$\mathbf{p}_m^{(t)} = \left(\sum_{i \in \{1, \dots, t-1\}} w_m^{(i)} \mathbf{h}_m^{(i)} \right) \mathbf{W}_m + \mathbf{b}_m, \quad (14)$$

where $\mathbf{W}_m \in \mathbb{R}^{d_{enc} \times d_K}$ and $\mathbf{b}_m \in \mathbb{R}^{d_K}$ are trainable parameters. The embedding $\tilde{\mathbf{p}}^{(t)} + \mathbf{p}_m^{(t)}$ can be considered as the patient feature for the alignment in Eq. (10).

3.6 Training Objective

For each patient in the training set, we use the bce loss in Eq. (15) as the optimization objective, which is a summation for all the T times visits. Here, $m^{(t)} \in \{0, 1\}^{|M|}$ and $\hat{m}^{(t)} \in [0, 1]^{|M|}$ are the ground truth label and recommendation probabilities for a drug set \mathcal{M} , respectively.

$$\mathcal{L}_{bce} = - \sum_{t=1}^T \sum_{i=1}^{|M|} \left(m_i^{(t)} \log \hat{m}_i^{(t)} + (1 - m_i^{(t)}) \log(1 - \hat{m}_i^{(t)}) \right) \quad (15)$$

We utilize the margin loss \mathcal{L}_{mar} in Eq. (16), aiming to enlarge the discrepancy between the predicted probability of positive and negative cases in the prescription.

$$\mathcal{L}_{mar} = \sum_{t=1}^T \sum_{\{i|m_i^{(t)}=1\}} \sum_{\{j|m_j^{(t)}=0\}} \frac{\max\{1 - (\hat{m}_i^{(t)} - \hat{m}_j^{(t)}), 0\}}{|\mathcal{M}|}. \quad (16)$$

The final training objective of the recommendation prediction can be formulated as:

$$\mathcal{L} = \alpha \mathcal{L}_{bce} + (1 - \alpha) \mathcal{L}_{mar}, \quad (17)$$

where α is a hyperparameter to balance these two loss terms.

3.7 Prediction of \mathcal{M}_t

In the prediction stage, given the medication set \mathcal{M} and the model $f(\cdot)$, we first obtain the features for all the medications as $\tilde{\mathbf{E}}_m$. For the t -th visit of the patient v , we acquire its representation $\mathbf{p}_v = \tilde{\mathbf{p}}^{(t)} + \mathbf{p}_m^{(t)}$, by employing $f(\cdot)$ as the query embedding. Subsequently, we compute the similarities between \mathbf{p}_v and all medications in $\tilde{\mathbf{E}}_m$ and then sort the medications based on similarity in ascending order. To produce the final recommendation set \mathcal{M}_t , we apply the predefined threshold value δ to select medications with a similarity greater than δ .

3.8 Complexity Analysis

Given N patients, $|\mathcal{M}|$ medications, d_K as the hidden dimension of MLP layer, d_{enc} as the output dimension of PLMs, d_s as the output dimension of GNN encoder, T as the maximum number of historical visits for all patients, and $iter$ as the number of iterations, we analyze the complexity of *NLA-MMR* as follows.

Time complexity. The time complexity of linear feature transformation (Eq.(4)) is $O(T \cdot d_{enc}^2)$, and that of the attention between patient textual descriptions (Eq.(2)-(3)) is $O(T \cdot d_{enc} \cdot d_K)$. For the GNN encoder (Eq.(5)-(7)), the time complexity is $O(|\mathcal{M}| \cdot L \cdot (e_0 \cdot d_s + n \cdot d_s^2))$, where L is the number of GNN layers, e_0 is the number of edges in a molecular graph, and n denotes the number of atoms. The time complexity of MLP layers in Eq.(9) is $O(T \cdot d_{enc} \cdot d_K)$ for patients and $O(|\mathcal{M}| \cdot (d_{enc} + d_s) \cdot d_K)$ for medications, respectively. The complexity of the output layer in Eq.(10) for medication prediction is $O(T \cdot |\mathcal{M}| \cdot d_K)$. The complexity of attention between all past visits (Eq.(12)-(14)) is $O(T \cdot d_{enc} \cdot d_K)$. With an average degree of 2.1 in molecular graphs for MIMIC-III and MIMIC-IV, and $d_s < d_K < d_{enc}$, the total time complexity is $O(iter \cdot N \cdot (T \cdot d_{enc}^2 + |\mathcal{M}| \cdot d_{enc} \cdot d_K + |\mathcal{M}| \cdot L \cdot n \cdot d_s^2))$.

Space complexity. The space of linear feature transformation in Eq.(4) is $O(d_{enc}^2)$. For attention in the patient representation and historical medication integration, the space is $O(d_{enc} \cdot d_K)$. The GNN encoder (Eq.(5)-(7)) costs $O(n \cdot d_s)$ space. The matrix in Eq.(10) costs $O(|\mathcal{M}| \cdot d_K)$ space for the medication prediction task. Assume $|\mathcal{M}| < d_K < d_{enc}$, the total space complexity is $O(d_{enc}^2 + n \cdot d_s)$.

4 EXPERIMENTAL STUDY

In this section, we describe the experimental setting (§4.1) and report the experiments in the following facets: (1) Compare the modeling effectiveness of *NLA-MMR* with state-of-the-art drug recommendation approaches (§4.2). (2) Conduct ablation studies to test the core components of *NLA-MMR* (§4.3). (3) Test the effect

of PLMs on *NLA-MMR* (§4.4). (4) Compare the impact of different numbers of visits on *NLA-MMR* (§4.5) and two strong baselines. (5) Conduct a visualization study on the effectiveness of the alignment module proposed in our approach (§4.6). (6) Investigate the sensitivity of *NLA-MMR* regarding different hyper-parameter configurations (§4.7). (7) Compare the efficiency of *NLA-MMR* with the deep-learning based baselines (§4.8).

4.1 Experimental Setup

Baselines. To comprehensively evaluate *NLA-MMR*, we compare with 3 instance-based methods: *LR*, *ECC* [25], *LEAP* [40], and 7 longitudinal-based methods: *RETAIN* [3], *GAMENet* [20], *MICRON* [37], *SafeDrug* [38], *COGNet* [36], *DrugRec* [31] and *MoleRec* [39].

Datasets. We use 3 public datasets: MIMIC-III [11], MIMIC-IV [12] and eICU [24]. For the MIMIC-III and MIMIC-IV datasets, we follow the data processing of *DrugRec* [31], which collates the diagnosis, procedure, and medication records of patients. For the eICU dataset, we only collect the diagnosis and medication records of the patients since the eICU dataset itself does not include procedure records. Because the DDI between the medications in the eICU dataset is not provided, we do not evaluate the DDI rate on eICU. Table 1 lists the profiles of the processed data. For the MIMIC datasets, we choose the ATC Third Level code as the target label and extract its corresponding textual description from DrugBank [35] as the textual medication description. As the ATC code is not available in the eICU dataset, we choose the generic therapeutic class (GTC) as the target label and obtain the textual description of the GTC code by engaging professionals to annotate its corresponding ATC Third Level code. The textual description of the ATC Third Level code is then adopted as the textual description of the GTC code. We split train, validation, and test by 2/3, 1/6, and 1/6, respectively.

Table 1: Profile of datasets

Items	MIMIC-III	MIMIC-IV	eICU
# of patients	5,208	6,136	9,539
# of visits	13,490	17,813	20,912
# of diagnoses	1,895	1,851	1,371
# of procedures	1,378	4,001	-
# of symptoms	428	163	-
# of medications	112	121	31
Avg # of visits	2.59	2.90	2.19
Avg # of diagnoses per visit	10.24	11.78	4.70
Avg # of procedures per visit	3.85	2.18	-
Avg # of symptoms per visit	7.67	1.09	-
Avg # of medications per visit	11.30	6.68	5.98

Evaluation Metrics. We use the following metrics: (a) F1 score (F1) - the harmonic mean of precision and recall; (b) Jaccard score (Jaccard) - the ratio of the intersection to the union of predicted and true labels; (c) Precision Recall AUC (PRAUC) - the area under the recall-precision curve, measuring the performance of a model across various recall levels.

Implementation and Reproducibility. Our experiments are performed using PyTorch 2.0.1. For drug molecule encoding, we adopt a 4-layer GIN [28] with a hidden embedding size of $d_s = 64$. For each molecule graph, the 9-dimensional initial node features contain the

atomic number and chirality, as well as other additional atom features, while the 3-dimensional edge features contain bond type, bond stereochemistry and whether the bond is conjugated. The size of text embeddings, denoted as d_{enc} , is 768. The projection layers are implemented as the MLP with 3 hidden layers of sizes [768, 256, 256]. We set the dropout ratio as 0.2. The threshold value δ is set to 0.5. The loss weight α is set to 0.95. We use Adam optimizer with a decaying learning rate to train the model for 110 epochs and the initial learning rate is set to 0.0005. We use NVIDIA 3090 GPUs. We consider five different PLMs and use public releases of *BioBERT* [17]¹, *SciBERT* [1]², *ClinicalBERT* [16]³, *PubMedBERT* [6]⁴ and *BlueBERT* [23]⁵, and PLMs are frozen during the experiments.

4.2 Overall Performance

We compare *NLA-MMR* with the baselines in terms of Jaccard, F1, and PRAUC on MIMIC-III, MIMIC-IV and eICU using *BioBERT* [17] as the PLM. The comparison results in Table 2 show that *NLA-MMR* outperforms previous approaches regarding all the evaluation metrics, by a significant margin. The main reason for this is our effective integration of expert knowledge through the introduction of PLMs to extract patient and medication representations, followed by the alignment of the representations. Additionally, our medication representations, which take into account both the chemical structures and textual descriptions, exhibit greater expressiveness compared to the representations employed in *MoleRec* and *DrugRec*, which fail to leverage the textual medication descriptions for learning medication representations. The instance-based approaches such as *LR*, *ECC*, and *LEAP* perform poorly and this is because these approaches fail to utilize longitudinal information in the EHR. Among the longitudinal-based methods, *GAMENet* incorporates additional graph information, while *SafeDrug* and *MoleRec* utilize drug molecule structures and substructures, respectively. These enhancements contribute to further performance improvements.

The average DDI rate of the ground truth medication combinations is 0.08238 and 0.07423 in MIMIC-III and MIMIC-IV, respectively. We compute the DDI rate from the recommendation results of each method. Then we report the difference between the DDI of each method and that of the ground truth (i.e., ΔDDI) in Figure 3. We can see that the DDI rate from the predicted results by *NLA-MMR* has the smallest difference from that of the ground truth. This observation demonstrates that our proposed *NLA-MMR* model proficiently simulates physicians’ medication prescribing patterns.

4.3 Ablation Study

In this section, we conduct ablation studies to investigate the effectiveness of different components and settings of *NLA-MMR*, including the drug representation module, cross-attention fusion, historical information integration mechanism, and using different types of textual descriptions in the EHRs for cross-attention fusion. First, we ablate the structure-level representation and PLM-based functional representation in *NLA-MMR*. From Table 3 (a)-(b), we can observe that conducting the recommendation solely based on

¹<https://huggingface.co/dmis-lab/biobert-base-cased-v1.2>

²https://huggingface.co/allenai/scibert_scivocab_uncased


³https://huggingface.co/emilyalsentzer/Bio_ClinicalBERT

⁴<https://huggingface.co/microsoft/BiomedNLP-BiomedBERT-base-uncased-abstract-fulltext>

⁵https://huggingface.co/bionlp/bluebert_pubmed_mimic_uncased_L-12_H-768_A-12

Table 2: Result comparison of different methods on three public datasets: MIMIC-III, MIMIC-IV and eICU

Method	MIMIC-III			MIMIC-IV			eICU		
	Jaccard \uparrow	F1 \uparrow	PRAUC \uparrow	Jaccard \uparrow	F1 \uparrow	PRAUC \uparrow	Jaccard \uparrow	F1 \uparrow	PRAUC \uparrow
<i>LR</i>	0.4896 \pm 0.0025	0.6491 \pm 0.0024	0.7568 \pm 0.0025	0.3844 \pm 0.0028	0.5379 \pm 0.0031	0.6568 \pm 0.0036	0.4199 \pm 0.0027	0.5723 \pm 0.0031	0.7515 \pm 0.0032
<i>ECC</i>	0.4799 \pm 0.0022	0.6390 \pm 0.0022	0.7572 \pm 0.0026	0.3680 \pm 0.0041	0.5173 \pm 0.0047	0.6541 \pm 0.0030	0.3786 \pm 0.0036	0.5240 \pm 0.0039	0.7498 \pm 0.0036
<i>LEAP</i>	0.4465 \pm 0.0037	0.6097 \pm 0.0036	0.6490 \pm 0.0033	0.3653 \pm 0.0028	0.5201 \pm 0.0033	0.5314 \pm 0.0038	0.4514 \pm 0.0035	0.5986 \pm 0.0036	0.6645 \pm 0.0048
<i>RETAIN</i>	0.4780 \pm 0.0036	0.6397 \pm 0.0036	0.7601 \pm 0.0035	0.3903 \pm 0.0038	0.5471 \pm 0.0040	0.6563 \pm 0.0055	0.4240 \pm 0.0033	0.5723 \pm 0.0034	0.7199 \pm 0.0036
<i>GAMENet</i>	0.5039 \pm 0.0021	0.6609 \pm 0.0020	0.7632 \pm 0.0027	0.3957 \pm 0.0035	0.5525 \pm 0.0041	0.6479 \pm 0.0055	0.4590 \pm 0.0014	0.6073 \pm 0.0012	0.7557 \pm 0.0026
<i>MICRON</i>	0.5076 \pm 0.0037	0.6634 \pm 0.0035	0.7685 \pm 0.0038	0.4009 \pm 0.0044	0.5545 \pm 0.0048	0.6584 \pm 0.0043	0.3741 \pm 0.0042	0.5192 \pm 0.0044	0.7551 \pm 0.0039
<i>SafeDrug</i>	0.5090 \pm 0.0038	0.6664 \pm 0.0033	0.7647 \pm 0.0020	0.4082 \pm 0.0026	0.5651 \pm 0.0028	0.6495 \pm 0.0036	0.4474 \pm 0.0038	0.5964 \pm 0.0032	0.7565 \pm 0.0037
<i>COGNet</i>	0.5134 \pm 0.0027	0.6706 \pm 0.0043	0.7677 \pm 0.0013	0.4131 \pm 0.0020	0.5660 \pm 0.0019	0.6460 \pm 0.0017	0.4588 \pm 0.0020	0.5963 \pm 0.0024	0.7263 \pm 0.0020
<i>DrugRec</i>	0.5220 \pm 0.0034	0.6771 \pm 0.0031	0.7720 \pm 0.0036	0.4194 \pm 0.0020	0.5713 \pm 0.0022	0.6558 \pm 0.0026	0.4508 \pm 0.0030	0.5999 \pm 0.0027	0.7496 \pm 0.0030
<i>MoleRec</i>	0.5278 \pm 0.0033	0.6825 \pm 0.0030	0.7736 \pm 0.0036	0.4197 \pm 0.0030	0.5772 \pm 0.0032	0.6644 \pm 0.0046	0.4476 \pm 0.0031	0.5958 \pm 0.0032	0.7575 \pm 0.0033
<i>NLA-MMR</i>	0.5429\pm0.0027	0.6958\pm0.0024	0.7890\pm0.0016	0.4491\pm0.0044	0.6051\pm0.0042	0.6921\pm0.0045	0.4787\pm0.0025	0.6262\pm0.0025	0.7650\pm0.0039

**Figure 3: Δ DDI comparison of different methods on MIMIC-III and MIMIC-IV****Table 3: Ablation study on *NLA-MMR***

Model Variants	MIMIC-III		
	Jaccard	F1	PRAUC
(a) w/o E_s	0.5368	0.6909	0.7872
(b) w/o E_f	0.4747	0.6352	0.7394
(c) w/o cross attention fusion	0.5348	0.6892	0.7815
(d) w/o historical information	0.5303	0.6848	0.7794
(e) cross attention fusion w/o h_d	0.5413	0.6946	0.7867
(f) cross attention fusion w/o h_p	0.5425	0.6955	0.7884
(g) cross attention fusion w/o h_s	0.5415	0.6946	0.7883
Full version	0.5429	0.6958	0.7890

either the structural or textual medication representation results in performance degradation, which confirms the importance of incorporating both structural information and textual knowledge in the medication modality.

Second, to exploit the effectiveness of the cross-attention fusion in the patient representation module and historical information integration technique, we remove these two components, respectively. As shown in Table 3 (c)-(d), only leveraging the feature in Eq. (1) as the patient representation will degrade the performance of our proposed model to a large extent, showing our model can benefit from fusing different types of texts in the EHRs. Besides, we notice that disabling the historical information hurts performance significantly on all evaluation metrics.

Third, we investigate the importance of different types of text from the patient modality used in the cross-attention fusion module in Table 3 (e)-(g). For each ablation, we ignore one type of textual description for conducting cross-attention fusion in Eq. (4). We observe that implementing the cross-attention with two types of textual description is also beneficial and the best performance is achieved with all the types of textual description.

4.4 Effect of Pretrained Language Models

To assess the impact of different PLMs on *NLA-MMR*, we use five public releases of PLMs, including *BioBERT* [17], *SciBERT* [1], *ClinicalBERT* [16], *PubMedBERT* [6] and *BlueBERT* [23]. Among these PLMs, *BioBERT* is initialized with the standard BERT model and then continues pretraining using PubMed abstracts, while *PubMedBERT* is pretrained from scratch using these texts. *ClinicalBERT* conducts continual pretraining from *BioBERT* with clinical notes from MIMIC-III [11]. *BlueBERT* mixes PubMed and MIMIC-III [11] to conduct continual pretraining from BERT. *SciBERT* is pretrained from scratch using biomedicine and computer science articles.

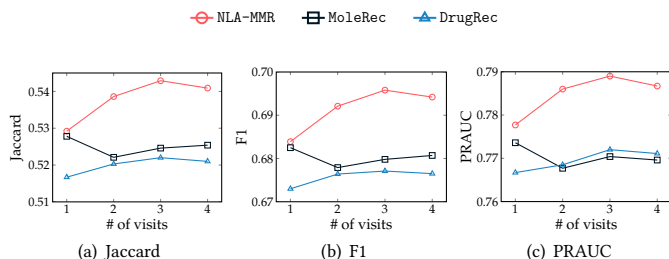
Table 4 compares the performance of *NLA-MMR* using five different PLMs. The evaluation measures for the *NLA-MMR* implemented with all five PLMs are much higher than those for the best baseline *MoleRec*, which confirms the general applicability and superiority of our proposed approach. We observe that the model achieves the best performance when using *BioBERT* as the PLM. In particular, even though clinical notes in MIMIC-III are relevant to the domain of diagnosis and procedure texts in our drug recommendation task, adding them to the pretraining corpus does not bring any advantage, as is evident by the results of *ClinicalBERT* and *BlueBERT*. This reveals that mixed-domain pre-training of the texts related to the patients and drugs is not beneficial to the CMR task, which supports the assumption that textual descriptions of the patient and drugs are different from the two modalities. However, the PLM such as *SciBERT* mixing out-of-domain texts from science articles during the pretraining procedure leads to worse performance.

4.5 Effect of Historical Visits

To evaluate the impact of the number of historical visits on the model performance, we test our approach with different numbers of visits on MIMIC-III using *BioBERT*. As a comparison, we include the two strongest baselines *MoleRec* and *DrugRec*, which

Table 4: The influence of using various PLMs on *NLA-MMR*


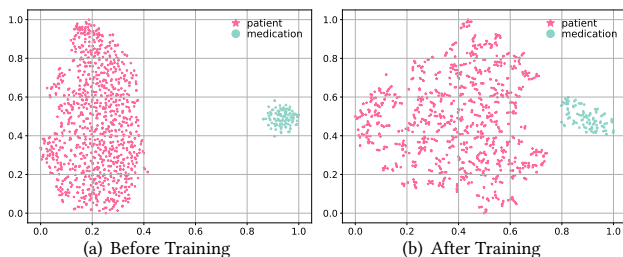
PLM	MIMIC-III			MIMIC-IV			eICU		
	Jaccard \uparrow	F1 \uparrow	PRAUC \uparrow	Jaccard \uparrow	F1 \uparrow	PRAUC \uparrow	Jaccard \uparrow	F1 \uparrow	PRAUC \uparrow
<i>BioBERT</i>	0.5429 \pm 0.0027	0.6958 \pm 0.0024	0.7890 \pm 0.0016	0.4491 \pm 0.0044	0.6051 \pm 0.0042	0.6921 \pm 0.0045	0.4787 \pm 0.0025	0.6262 \pm 0.0025	0.7650 \pm 0.0039
<i>PubMedBERT</i>	0.5375 \pm 0.0024	0.6909 \pm 0.0020	0.7849 \pm 0.0017	0.4471 \pm 0.0044	0.6028 \pm 0.0043	0.6891 \pm 0.0043	0.4762 \pm 0.0023	0.6240 \pm 0.0022	0.7662 \pm 0.0042
<i>ClinicalBERT</i>	0.5382 \pm 0.0025	0.6917 \pm 0.0021	0.7838 \pm 0.0016	0.4483 \pm 0.0044	0.6037 \pm 0.0042	0.6880 \pm 0.0043	0.4755 \pm 0.0018	0.6227 \pm 0.0017	0.7649 \pm 0.0038
<i>BlueBERT</i>	0.5296 \pm 0.0026	0.6842 \pm 0.0022	0.7772 \pm 0.0016	0.4377 \pm 0.0042	0.5946 \pm 0.0042	0.6785 \pm 0.0042	0.4712 \pm 0.0026	0.6186 \pm 0.0026	0.7606 \pm 0.0042
<i>SciBERT</i>	0.5336 \pm 0.0032	0.6878 \pm 0.0027	0.7779 \pm 0.0019	0.4399 \pm 0.0053	0.5961 \pm 0.0050	0.6810 \pm 0.0050	0.4719 \pm 0.0023	0.6195 \pm 0.0022	0.7632 \pm 0.0030

**Figure 4: The impact of # of visits on MIMIC-III for two strong baselines and *NLA-MMR* using *BioBERT* as PLM**

also incorporate historical information. Figure 4 shows the Jaccard, F1, and PRAUC of *NLA-MMR* and the two baselines when setting the number of visits in $\{1, 2, 3, 4\}$. We observe that *NLA-MMR* achieves relatively better performance with more visits, while the performance of *MoleRec* almost stays flat and *DrugRec* consistently exhibits poor performance. This indicates that the design of the historical information integration mechanism can effectively incorporate prescription information of historical visits into the CMR task, thereby directly enhancing the accuracy of drug predictions. On the other hand, *MoleRec* relies on an RNN-based mechanism, which may not be as adept at capturing and utilizing the relevant information from past visits, leading to its relatively stagnant performance. Note that both *NLA-MMR* and *DrugRec* achieve the best performance by utilizing the historical information in the last two visits. It suggests that outdated information may not provide useful guidance and could potentially lead to misleading prescriptions in current drug predictions.

4.6 Visualization

Recall that in the medication representation module, we incorporate the PLM-based functional representation with GNN-based structure-level representation as the medication representation (Eq. (8)). In this section, we conduct a case study to explore whether the structure-level representation can capture distinct information beyond the textual medication descriptions, and thus providing extra assistance for the CMR task. First, we compare the two drug representations, the pure textural representation E_f , and the representation E_m incorporated extra structural feature by GNN. Figure 5 shows the visualization of the two representations, i.e., w/o GNN and w/ GNN, on the MIMIC-III and MIMIC-IV by t-SNE [32], where *BioBERT* is used as the PLM. We observe that the distributions of the representations w/o and w/ GNN are different. It is noticed that the *NLA-MMR* model can learn the chemical properties of drugs because the incorporation of structural information. The drug representations derived from chemical structures and

**Figure 5: Visualization of GNN-enhanced and text-only based medication representations using t-SNE****Figure 6: Visualization of alignment between the representations from the patients and medications**

textual descriptions exhibit noticeable distinctions when compared to text-based drug representations in the embedding space.

Furthermore, to gain a better understanding of the multi-modal alignment module in *NLA-MMR*, we visualize the representation of 1000 patient samples and the representation of all the 112 drugs, generated by *NLA-MMR* before and after training. Figure 6 presents the two representations of MIMIC-III by t-SNE, where *BioBERT* is used as the PLM. We observe that before training the representations of patient and drug are two separate manifolds (Figure 6(a)), while after the training, the two manifolds are closer in the latent space and present a trend of integration (Figure 6(b)). This visualization confirms that our multi-modal alignment module is capable of aligning the representations of patients and drugs.

4.7 Hyper-parameter Sensitivity

In this section, we test the parameter sensitivity of *NLA-MMR* on the MIMIC-III and MIMIC-IV datasets using *BioBERT* as the PLM. **Impact of batch.** We conduct experiments with $lr = 5e-4$ and $\alpha = 0.95$. In Figure 7 (a)-(c), we vary *batch* from 2 to 64 and show the performance of *NLA-MMR*. Overall, the performance of *NLA-MMR* remains stable across different batch sizes. The performance gap of *NLA-MMR* with varied batch size is 0.0123 and 0.0104 for MIMIC-III and MIMIC-IV, respectively. Notably, the optimal batch size for MIMIC-III and MIMIC-IV are 8 and 4, respectively.




Figure 7: The impact of *batch*, *lr*, and α on MIMIC-III and MIMIC-IV using BioBERT as PLM

Impact of *lr*. We conduct experiments with $\alpha = 0.95$ and *batch* = 8. In Figure 7 (d)-(f), we vary *lr* in {1e-5, 5e-5, 1e-4, 5e-4, 1e-3, 5e-3, 1e-2} and show the performance of *NLA-MMR*. The results indicate that our model is sensitive to the learning rate and the optimal learning rate is 5e-4 across two datasets.

Impact of α . We conduct experiments with *lr* = 5e-4 and *batch* = 8. In Figure 7 (g)-(i), we set $\alpha=0$ and the significant performance degradation confirms the greater importance of the bce loss. Another observation is that our model is stable with varied α from 0.90 to 1.0 and the optimal value for α is 0.95 across two datasets.

4.8 Efficiency

We compare the training and inference time of *NLA-MMR* with 8 deep learning-based baselines for MIMIC-III and MIMIC-IV datasets on a Quadro RTX 800 GPU in Table 5. The results demonstrate that our *NLA-MMR* consistently outperforms most of the baselines in terms of training time across two datasets. Additionally, its inference time is comparable to that of *RETAIN*, *SafeDrug*, and *GAMENet*, and better than the other baselines across two datasets. Therefore, we conclude that *NLA-MMR* is a viable solution for real deployment.

5 RELATED WORK

5.1 Molecular Representation Learning

Molecular representation learning aims to utilize deep learning models to encode the input molecules as numerical vectors, which preserve the structural and property information about the molecules and serve as feature vectors for downstream chemical applications. Existing medication recommendation approaches [31, 38, 39] mainly model the chemical structure of molecules through one-dimensional description [8, 14], two-dimensional molecular graphs [4, 18, 19, 21, 26] and three-dimensional geometric graphs [7, 9, 10]. Besides the chemical structural information, we have textual medication descriptions that are available on PubChem [13]

Table 5: Efficiency comparison on MIMIC-III and MIMIC-IV

Model	Training time (s/Epoch)		Testing time (s)	
	MIMIC-III	MIMIC-IV	MIMIC-III	MIMIC-IV
<i>LEAP</i>	317.24 \pm 9.27	422.83 \pm 14.70	26.55 \pm 0.44	28.26 \pm 0.16
<i>RETAIN</i>	45.81 \pm 2.82	54.44 \pm 2.10	3.78 \pm 0.10	5.62 \pm 0.27
<i>GAMENet</i>	140.11 \pm 9.23	185.05 \pm 0.89	7.50 \pm 0.37	10.69 \pm 0.04
<i>MICRON</i>	87.63 \pm 1.61	105.24 \pm 1.14	11.55 \pm 0.49	14.93 \pm 0.03
<i>SafeDrug</i>	194.47 \pm 21.19	303.54 \pm 13.03	7.64 \pm 1.06	9.76 \pm 0.47
<i>COGNet</i>	123.77 \pm 1.86	135.35 \pm 1.70	92.22 \pm 1.57	105.80 \pm 0.22
<i>DrugRec</i>	461.86 \pm 9.37	605.77 \pm 16.15	26.78 \pm 0.17	33.93 \pm 1.76
<i>MoleRec</i>	479.22 \pm 10.89	551.38 \pm 2.78	27.54 \pm 1.95	28.81 \pm 0.07
<i>NLA-MMR</i>	50.83 \pm 3.25	62.98 \pm 2.34	7.67 \pm 0.21	10.50 \pm 0.28

and DrugBank [35]. It is observed that these textual medication descriptions offer a comprehensive perspective on the functional aspects of molecules and explain their therapeutic applications. In this paper, we aim to incorporate textual knowledge into molecular representation learning.

5.2 Medication Recommendation

Existing approaches can be broadly classified into two main categories: instance-based methods and longitudinal methods. Instance-based methods [5, 40] only take information of the current visit as input. For example, *LEAP* [40] formulates the medication recommendation task into a sequential decision-making process and adopts a multi-instance multi-label learning framework to generate the medication recommendations based on the patient’s current diagnosis information.

In contrast, longitudinal methods [15, 29, 33, 34, 37, 38] use the historical information of patients and explore sequential dependency within clinical visits. Most of them model longitudinal patient information using Recurrent Neural Networks. *GAMENet* [29] and *SafeDrug* [38] further incorporated the BioKG and DDI, respectively, to improve model performance. Despite significant progress made in using structured domain knowledge in the CMR task, these methods suffer from the inherent biases and information loss in the structured data caused by the preprocessing stage. This paper aims to encode the raw texts in patient and drug modalities by using powerful PLMs as the foundation block to generate their representations, which is an effective way of integrating multimodal information in multimodal recommendation systems [22, 41].

6 CONCLUSION

In this paper, we consider the patient and medication as two distinct modalities and design a multi-modal framework *NLA-MMR* to learn their representations according to the alignment. Besides, we investigate the potential of PLMs to extract domain knowledge in the textual descriptions from the patient and medication modalities. In the medication modality, we incorporate both the chemical structures and medication descriptions to enhance molecular representation learning. Extensive experiments on three public datasets show the effectiveness and efficiency of our proposed *NLA-MMR*.

ACKNOWLEDGEMENT

Yu Rong and Kangfei Zhao are the corresponding authors. This research was partially supported by the CUHK Stanley Ho Big Data Decision Analytics Research Centre. Kangfei Zhao is supported by National Key Research and Development Plan, No. 2023YFF0725101.

REFERENCES

- [1] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. In *In Proc. EMNLP-IJCNLP. ACL*, 3613–3618.
- [2] Tian Bian, Yuli Jiang, Jia Li, Tingyang Xu, Yu Rong, Yi Su, Timothy C. Y. Kwok, Helen Meng, and Hong Cheng. 2023. Decision Support System for Chronic Diseases Based on Drug-Drug Interactions. In *International Conference on Data Engineering. IEEE*, 3467–3480.
- [3] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter F. Stewart. 2016. RETAIN: An Interpretable Predictive Model for Healthcare using Reverse Time Attention Mechanism. In *Proc. NeurIPS*, 3504–3512.
- [4] David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. 2015. Convolutional Networks on Graphs for Learning Molecular Fingerprints. In *Proc. NeurIPS*, 2224–2232.
- [5] Fan Gong, Meng Wang, Haofen Wang, Sen Wang, and Mengyue Liu. 2021. SMR: Medical Knowledge Graph Embedding for Safe Medicine Recommendation. *Big Data Res.* 23 (2021), 100174.
- [6] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2022. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *ACM Trans. Comput. Heal.* 3, 1 (2022), 2:1–2:23.
- [7] Jiaqi Han, Wenbing Huang, Tingyang Xu, and Yu Rong. 2022. Equivariant Graph Hierarchy-Based Neural Networks. In *Proc. NeurIPS*, Vol. 35. Curran Associates, Inc., 9176–9187.
- [8] Kexin Huang, Cao Xiao, Trong Hoang, Lucas Glass, and Jimeng Sun. 2020. Caster: Predicting drug interactions with chemical substructure representation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 702–709.
- [9] Wenbing Huang, Jiaqi Han, Yu Rong, Tingyang Xu, Fuchun Sun, and Junzhou Huang. 2022. Equivariant Graph Mechanics Networks with Constraints. In *Proc. ICLR*.
- [10] Rui Jiao, Jiaqi Han, Wenbing Huang, Yu Rong, and Yang Liu. 2023. Energy-motivated equivariant pretraining for 3d molecular graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 8096–8104.
- [11] 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.
- [12] Alistair EW Johnson, David J Stone, Leo A Celi, and Tom J Pollard. 2018. The MIMIC Code Repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association* 25, 1 (2018), 32–39.
- [13] Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A. Shoemaker, Paul A. Thiessen, Bo Yu, Leonid Zaslavsky, Jian Zhang, and Evan Bolton. 2021. PubChem in 2021: new data content and improved web interfaces. *Nucleic Acids Res.* 49, Database-Issue (2021), D1388–D1395.
- [14] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alán Aspuru-Guzik. 2020. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Mach. Learn. Sci. Technol.* 1, 4 (2020), 45024.
- [15] Hung Le, Truyen Tran, and Svetla Venkatesh. 2018. Dual Memory Neural Computer for Asynchronous Two-view Sequential Learning. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining. ACM*, 1637–1645.
- [16] Cecilia S Lee and Aaron Y Lee. 2020. Clinical applications of continual learning machine learning. *The Lancet Digital Health* 2, 6 (2020), e279–e281.
- [17] Jinhuyk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinform.* 36, 4 (2020), 1234–1240.
- [18] Xuan Li, Zhanke Zhou, Jiangchao Yao, Yu Rong, Lu Zhang, and Bo Han. 2024. Neural Atoms: Propagating Long-range Interaction in Molecular Graphs through Efficient Communication Channel. In *Proc. ICLR*.
- [19] Shengchao Liu, Mehmet Furkan Demirel, and Yingyu Liang. 2019. N-Gram Graph: Simple Unsupervised Representation for Graphs, with Applications to Molecules. In *Proc. NeurIPS*, 8464–8476.
- [20] Yunan Luo, Guangde Jiang, Tianhao Yu, Yang Liu, Lam Vo, Hantian Ding, Yufeng Su, Wesley Wei Qian, Huimin Zhao, and Jian Peng. 2021. ECNet is an evolutionary context-integrated deep learning framework for protein engineering. *Nature Communications* 12, 1 (Sept. 2021).
- [21] Hehuan Ma, Yatao Bian, Yu Rong, Wenbing Huang, Tingyang Xu, Weiyang Xie, Geyan Ye, and Junzhou Huang. 2022. Cross-dependent graph neural networks for molecular property prediction. *Bioinform.* 38, 7 (01 2022), 2003–2009.
- [22] Zongshen Mu, Yueting Zhuang, Jie Tan, Jun Xiao, and Siliang Tang. 2022. Learning hybrid behavior patterns for multimedia recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*. 376–384.
- [23] Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. Transfer Learning in Biomedical Natural Language Processing: An Evaluation of BERT and ELMo on Ten Benchmarking Datasets. In *Proceedings of the 18th BioNLP Workshop and Shared Task, BioNLP@ACL*, Dina Demner-Fushman, Kevin Bretonnel Cohen, Sophia Ananiadou, and Junichi Tsujii (Eds.). ACL, 58–65.
- [24] Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi. 2018. The eICU Collaborative Research Database, a freely available multi-center database for critical care research. *Scientific data* 5, 1 (2018), 1–13.
- [25] Jesse Read, Bernhard Pfahringer, Geoffrey Holmes, and Eibe Frank. 2009. Classifier Chains for Multi-label Classification. In *ECML PKDD. Springer*, 254–269.
- [26] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying WEI, Wenbing Huang, and Junzhou Huang. 2020. Self-Supervised Graph Transformer on Large-Scale Molecular Data. In *Proc. NeurIPS*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 12559–12571.
- [27] Suebarn Ruksakulpiwat, Wendie Zhou, Atsadaporn Niyomyart, Tongyao Wang, and Aaron Kudlowitz. 2023. How does the COVID-19 pandemic impact medication adherence of patients with chronic disease?: A systematic review. *Chronic illness* 19, 3 (2023), 495–513.
- [28] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. 2021. E(n) Equivariant Graph Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 9323–9332.
- [29] Junyuan Shang, Cao Xiao, Tengfei Ma, Hongyan Li, and Jimeng Sun. 2019. GAMeNet: Graph Augmented MEemory Networks for Recommending Medication Combination. In *Proceedings of the AAAI Conference on Artificial Intelligence. AAAI Press*, 1126–1133.
- [30] Kimberly A Shoenbill, Suranga N Kasturi, and Eneida A Mendonca. 2023. Artificial Intelligence, Machine Learning, and Natural Language Processing. In *Chronic Illness Care: Principles and Practice*. Springer, 469–479.
- [31] Hongda Sun, Shufang Xie, Shuqi Li, Yuhuan Chen, Ji-Rong Wen, and Rui Yan. 2022. Debiased, Longitudinal and Coordinated Drug Recommendation through Multi-Visit Clinic Records. In *Proc. NeurIPS*.
- [32] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [33] Lu Wang, Wei Zhang, Xiaofeng He, and Hongyuan Zha. 2018. Supervised Reinforcement Learning with Recurrent Neural Network for Dynamic Treatment Recommendation. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining. ACM*, 2447–2456.
- [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 ACM International Conference on Information and Knowledge Management. ACM*, 1623–1632.
- [35] David S. Wishart, Yannick D. Feunang, An Chi Guo, Elvis J. Lo, Ana Marcu, Jason R. Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, Nazanin Assempour, Ithayavani Iynkkaran, Yifeng Liu, Adam Maciejewski, Nicola Gale, Alex Wilson, Lucy Chin, Ryan Cummings, Diana Le, Allison Pon, Craig Knox, and Michael Wilson. 2018. DrugBank 5.0: a major update to the DrugBank database for 2018. *Nucleic Acids Res.* 46, Database-Issue (2018), D1074–D1082.
- [36] 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. ACM*, 935–945.
- [37] Chaoqi Yang, Cao Xiao, Lucas Glass, and Jimeng Sun. 2021. Change Matters: Medication Change Prediction with Recurrent Residual Networks. In *Proceedings of the International Joint Conference on Artificial Intelligence. ijcai.org*, 3728–3734.
- [38] Chaoqi Yang, Cao Xiao, Fenglong Ma, Lucas Glass, and Jimeng Sun. 2021. Safe-Drug: Dual Molecular Graph Encoders for Recommending Effective and Safe Drug Combinations. In *Proceedings of the International Joint Conference on Artificial Intelligence, Zhi-Hua Zhou (Ed.)*, 3735–3741.
- [39] 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. ACM*, 4075–4085.
- [40] Yutao Zhang, Robert Chen, Jie Tang, Walter F Stewart, and Jimeng Sun. 2017. LEAP: learning to prescribe effective and safe treatment combinations for multimorbidity. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1315–1324.
- [41] Hongyu Zhou, Xin Zhou, Zhiwei Zeng, Lingzi Zhang, and Zhiqi Shen. 2023. A Comprehensive Survey on Multimodal Recommender Systems: Taxonomy, Evaluation, and Future Directions. *CoRR abs/2302.04473* (2023).