

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

MIPO: Mutual Integration of Patient Journey and Medical Ontology for Healthcare Representation Learning

Xueping Peng*, Guodong Long*(✉), Tao Shen*, Sen Wang†, Chengqi Zhang‡, Allison Clarke§, Clement Schlegel§

*Australian AI Institute, Faculty of Engineering and IT, University of Technology Sydney, Australia

†School of Information Technology and Electrical Engineering, The University of Queensland, Australia

‡Department of Data Science and AI, The Hong Kong Polytechnic University, Hong Kong, China

§Health Economics and Research Division, Department of Health and Aged Care, Australia

(✉)Corresponding author email: guodong.long@uts.edu.au

Abstract—Representation learning on electronic health records (EHRs) plays a vital role in downstream medical prediction tasks. Although natural language processing techniques, such as recurrent neural networks, and self-attention, have been adapted for learning medical representations from hierarchical, time-stamped EHR data, they often struggle when either general or task-specific data are limited. Recent efforts have attempted to mitigate this challenge by incorporating medical ontologies (i.e., knowledge graphs) into self-supervised tasks like diagnosis prediction. However, two main issues remain: (1) small and uniform ontologies that lack diversity for robust learning, and (2) insufficient attention to the critical contexts or dependencies underlying patient journeys, which could further enhance ontology-based learning. To address these gaps, we propose MIPO (Mutual Integration of Patient Journey and Medical Ontology), a robust end-to-end framework that employs a Transformer-based architecture for representation learning. MIPO emphasizes task-specific representation learning through a sequential diagnosis prediction task, while also incorporating an ontology-based disease-typing task. A graph-embedding module is introduced to integrate information from patient visit records, thus alleviating data insufficiency. This setup creates a mutually reinforcing loop, where both patient-journey embedding and ontology embedding benefit from each other. We validate MIPO on two real-world benchmark datasets, showing that it consistently outperforms baseline methods under both sufficient and limited data conditions. Furthermore, the resulting diagnosis embeddings offer improved interpretability, underscoring the promise of MIPO for real-world healthcare applications.

Index Terms—Healthcare Informatics, Electronic Health Record, Knowledge Graph, Representation Learning

I. INTRODUCTION

Over the past few decades, healthcare information systems have accumulated a considerable amount of electronic health records (EHRs). The patient EHRs data typically consists of a sequence of visit records, where each visit includes clinical events such as diagnoses, procedures, medications, and laboratory tests [1], [2]. Exploiting knowledge from these voluminous EHRs has attracted attention for its potential to benefit patients and caregivers, driving advances in both academia [3]–[8] and industry [9], [10].

Recent studies [3]–[5], [11]–[13] have explored deep learning techniques to model EHR data for predictive tasks. For instance, word embedding techniques, e.g. word2vec [14], have been adopted in [3] to learn a vector representation (namely Med2Vec) for each medical concept (e.g. a diagnosis code) from the co-occurrence information without considering the temporal sequential nature of EHR data. Furthermore, considering both long-term dependency and sequential information, recurrent neural networks [4], [5], [12], [13], including LSTM [15] and GRU [16], are used to learn the contextualized representation of EHR data. These attempts to learn medical representations are still underperforming for prediction tasks and cannot be practically used for individual patients.

The reason can be two-fold. The first one is that the squeezing representation capability of EHRs will be bottle-necked for some specific prediction tasks. For example, bringing no more attention to the model when encountering the appearance of terminal disease in a patient’s visiting sequences can be detrimental to the performance accuracy in the mortality prediction task. Domain knowledge related to the specific task, with no doubt, can be borrowed to improve performance when modelling. Some NLP learning schemes [17], [18] have shed light on the extraction of task-wise information for a specific NLP task. Most commonly, they pre-train a general neural module (e.g., word embeddings [17] and contextual encoder [18]) on a large-scale unlabeled corpus with self-supervised tasks, and then leverage the pre-trained module to initialize the task-specific models for further fine-tuning. The corpus, such as Wikipedia and BookCorpus, consists of words, usually on a scale of billions, so the pre-trained module can produce enough generic representations for efficient fine-tuning convergence and superior performance. Unfortunately, compared to the sheer volume of textual data available to NLP tasks, the scale of unlabeled healthcare data is considerably smaller for pre-training to exploit sufficient task-wise medical knowledge.

The second hurdle in the applications of the above models can be blamed on the lack of interpretability, which is consid-

ered in priority when making a decision by caregivers. Graph embedding has been integrated into deep sequence models to improve the performance and provide interpretability. [4], [5], [19] train medical code embeddings upon medical ontology by using a graph-based attention mechanism, which delivers a competitive performance with interpretations aligning with medical knowledge. Note that a strict prerequisite of these works is that each medical code appears as a leaf node in the medical ontology which can be readily satisfied by healthcare data. To be clear, medical ontology here refers to a medical knowledge graph, e.g., Clinical Classifications Software (CCS)¹. Despite their success in several healthcare tasks, these methods still labour under two main limitations: (1) Unlike factoid knowledge graphs (e.g., Freebase and Wiki-Data), which store hundreds of millions of relational items, the medical ontology contains thousands of diagnosis nodes and merely “parent-child” hierarchy. Hence, it is insufficient to train expressively powerful code embeddings over the ontology; (2) Rich context or dependency information underlying each visit and the patient journey is rarely exploited during medical ontology learning which, however, contains essential information, e.g., complicated diseases.

To overcome these limitations, we propose a novel and robust healthcare representation learning model, called **Mutual Integration of Patient Journey and Medical Ontology (MIPO)**. It consists of two interactive neural modules: (1) *task-specific representation learning* module and (2) *graph-embedding* module. It aims to infuse medical knowledge into a sequential patient journey by jointly learning the task-specific and the ontology-based objectives. To clarify, a *task-specific representation learning* module is composed of two stacked Transformer encoders in a hierarchical scheme. It aims to measure local dependencies among medical codes in each patient visit and further capture long-term dependencies among multiple visits in a patient’s journey. Concurrently, the *graph-embedding* module learns code embeddings in medical ontology based on both structured knowledges in the graph and contextual information in the patient journey. Lastly, we jointly train the model to meet two objectives: one for *task-specific predictive task* based on the representation learning module, and another for *ontology-based disease typing task* based on the graph embedding module. Consequently, with such mutual integration and joint learning, MIPO can improve the prediction quality of future diagnoses, guarantee robustness regardless of sufficient or insufficient data, and make the learned patient journeys and diagnoses interpretable. Our main contributions are summarized as follows:

- We propose MIPO, an end-to-end, novel and robust model to accurately predict patients’ future visit information with mutual integration of patient journey and medical ontology.
- We design an *ontology-based disease typing task* in conjunction with the *task-specific predictive task*, to learn effective and robust healthcare representations.

- We qualitatively demonstrate the interpretability of the learned representations of medical codes and quantitatively validate the effectiveness of the proposed MIPO.

II. RELATED WORK

In this section, we review the work related to mining EHRs data with deep learning techniques, especially for diagnosis prediction. We then introduce some of the latest work on the transformer-based model with knowledge graph.

A. Deep Learning for EHRs Data

In recent years, researchers have proposed various deep learning models to garner knowledge from massive EHRs and shown their superior ability in medical event predictions [6], [7], [9], [12], [20]–[25]. Previous studies recommend using recurrent neural networks (RNNs) for patient subtyping [22], [26], modelling disease progression [27], and time-series healthcare-data analysis [28]. Convolutional neural networks (CNNs) are exploited to predict unplanned readmission [29] and risk [30] with EHRs. Stacked autoencoders are employed to generate sequential EHRs data [31]. The emerging transformer-based BERT model is used for acquiring knowledge from clinical notes [21] and future visit prediction [20].

Diagnosis prediction is an important application in healthcare analytics [4], [5], [32]–[34], which leverages a patient’s sequential visit records to predict future visit information. RETAIN [13] and Dipole [12] are two representative RNN-based diagnoses predictive models. RETAIN employs RNNs to model reverse, time-ordered-EHRs sequential visits with an attention mechanism for the binary prediction tasks. Dipole applies Bi-LSTM and attention mechanisms to predict patient visit information, which enhances the temporal data modelling ability of predictive models. However, those approaches could suffer from a data insufficiency [4]. To alleviate this problem; GRAM [4] and KAME [5] exploit the information from external medical knowledge graph to learn robust representations and an RNN to model patient visits. Although this achieves state-of-the-art performance, both models lack effective aggregation of multiple medical codes in a visit and, heterogeneous information integration of patient journey and knowledge graph, which should be taken as an advantage for improving performance.

B. Transformer-based Model with Knowledge Graph

Devlin et al. [18] propose a deep bi-directional model with multiple-layer Transformers (BERT), which achieves the state-of-the-art results for various NLP tasks (eg, question answering, named entity recognition, and relation extraction). The latest, transformer-based BEHRT [20] makes direct use of the original BERT [18] to model the patient’s sequential EHRs data by taking each visit as a sentence and each medical concept as a word to predict future visit information. BioBERT [35], ClinicalBERT [36] and Med-BERT [37] achieve new state-of-the-art results on various biomedical NLP tasks through simple fine-tuning techniques with medical corpus. Furthermore, ERNIE [38], K-BERT [39] and KGET [19]

¹<https://www.hcup-us.ahrq.gov/toolssoftware/ccs/ccs.jsp>

infuse a knowledge graph into pre-trained BERT to further enhance language representation. As EHRs data have different characteristics to natural language, (e.g., medical code can provide a one-to-one map node of knowledge graph), medical codes are time-ordered in a visit. Thus, enhanced language models, such as ERNIE and K-BERT, cannot apply directly to tackling healthcare problems. However, the ideas from pairing the NLP language models with knowledge graphs motivate us to propose MIPO to mutually integrate patient journey and medical knowledge for healthcare representation learning.




Fig. 1: The model architecture of MIPO. The graph is formatted as a hierarchical tree, in which, the root node is virtual. To construct the tree, the leaf nodes (solid circles) denote fine-grained diagnoses, and the non-leaf nodes (dotted circles) denote coarse-grained disease concepts.

III. METHODOLOGY

A. Notations

We denote the set of medical codes from the EHRs data as $c_1, c_2, \dots, c_{|\mathbb{C}|} \in \mathbb{C}$ and $|\mathbb{C}|$ is the number of unique medical codes. Patients' clinical records can be represented by a sequence of visits $\mathbf{P} = \langle V_1, \dots, V_t, \dots, V_T \rangle$, which is referred to as the patient journey in the paper, where T is the visit number in the patient journey. And each visit V_t consists of a subset of medical codes ($V_t \subseteq \mathbb{C}$). For clear demonstration, all algorithms will be presented with a single patient's journey. On the other hand, a medical ontology \mathcal{G} contains the hierarchy of various medical concepts with the *parent-child* semantic relationship, which is a well-organized ontology in healthcare. In particular, the medical ontology \mathcal{G} is a directed acyclic graph (DAG) and the nodes of \mathcal{G} consist of

leaves and their ancestors. Each leaf node refers to a medical code in \mathbb{C} , which is associated with a sequence of ancestors from the leaf to the root of \mathcal{G} . And each ancestor node belongs to the set $\mathbb{N} = n_{|\mathbb{C}|+1}, n_{|\mathbb{C}|+2}, \dots, n_{|\mathbb{C}|+|\mathbb{N}|}$, where $|\mathbb{N}|$ is the number of ancestor codes in \mathcal{G} . A parent in the knowledge graph \mathcal{G} represents a related but more general concept over its children.

B. Model Architecture

As shown in Figure 1, the whole model architecture of MIPO consists of an embedded knowledge graph and two stacked modules and an attention pooling layer. Using given knowledge graph \mathcal{G} , we can obtain the embedding matrix \mathbf{G} of medical codes with graph-based attention mechanism [4]. Given the t -th visit information of a patient V_t , each medical code corresponding to a leaf node in \mathcal{G} in V_t is embedded into a vector representation with the learned \mathbf{G} , which imposes medical knowledge on the network architecture. The two stacked modules are: (1) the underlying knowledgeable encoder (V-Encoder) responsible for integrating extra medical knowledge information into visit information from basic embedding, so that we can represent the heterogeneous information of medical codes and graph nodes into a united feature space, and (2) the upper patient encoder (P-Encoder) is responsible of capturing contextual and sequential information from the underlying layer. Attention Pooling [40], [41] explores the importance of each code within an entire visit. It works by compressing a set of medical code embeddings from a visit into a single context-aware vector representation for the upper P-Encoder. We also denote the number of V-Encoder layers as N , and the number P-Encoder layers as M . The output of P-Encoder is used to predict the information of next visit.

To be specific, given a patient's visit $V_t = \{c_1, c_2, \dots, c_n\}$, where n is the number of medical codes in the visit, we can first obtain its corresponding basic code embedding $\{w_1, w_2, \dots, w_n\}$ via code embedding layer and its node embedding $\{g_1, g_2, \dots, g_n\}$ via learned knowledge graph embedding matrix \mathbf{G} . More details of the knowledge graph embedding are introduced in Section III-C. Then, MIPO adopts a knowledgeable encoder V-Encoder to inject the medical knowledge information into healthcare representation, where both $\{w_1, w_2, \dots, w_n\}$ and $\{g_1, g_2, \dots, g_n\}$ are fed into V-Encoder for fusing heterogeneous information and computing final output embeddings,

$$\{w_1^o, w_2^o, \dots, w_n^o\}, \{g_1^o, g_2^o, \dots, g_n^o\} = \text{V-Encoder}(\{w_1, w_2, \dots, w_n\}, \{g_1, g_2, \dots, g_n\}), \quad (1)$$

where $\{g_1^o, \dots, g_n^o\}$ will be used as features for ontology-based disease typing task and $\{w_1^o, w_2^o, \dots, w_n^o\}$ will be fed into upper attention pooling layer (detailed in Section III-E) to compress then to a vector v_t representing the visit. More details of the knowledgeable encoder V-Encoder will be introduced in Section III-D.

After computing v_t , MIPO employs a P-Encoder to capture the contextual and sequential information, $v_t^o =$

P-Encoder(\mathbf{v}_t), where P-Encoder(\cdot) is a multi-layer bidirectional Transformer encoder. Since P-Encoder is identical to its implementation in BERT, we exclude a comprehensive description of this module and refer readers to [18] and [42].

For simplicity, we take only one patient's visit V_t as an example. However, since most patients have multiple visits to hospital, inputs and outputs of V-Encoder and P-Encoder are multiple visits, e.g.,

$$\{\mathbf{v}_1^o, \mathbf{v}_2^o, \dots, \mathbf{v}_{T-1}^o\} = \text{P-Encoder}(\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{T-1}\}), \quad (2)$$

where $\{\mathbf{v}_1^o, \mathbf{v}_2^o, \dots, \mathbf{v}_{T-1}^o\}$ will be used as features for the task of sequential diseases prediction.

C. Knowledge Graph Embedding

To mitigate the problem of data insufficiency in healthcare and to learn knowledgeable and generalized representations of medical codes, we employ the attention-based graph embedding approach GRAM [4]. In the medical ontology \mathcal{G} , each leaf node c_i has a basic learnable embedding vector $\mathbf{E}_{i,:} \in \mathbb{R}^d$, where $1 \leq i \leq |\mathcal{C}|$, and d represent the dimensionality. And each ancestor code n_i also has an embedding vector $\mathbf{E}_{i,:} \in \mathbb{R}^d$, where $|\mathcal{C}| + 1 \leq i \leq |\mathcal{C}| + |\mathcal{N}|$. The attention-based graph embedding uses an attention mechanism to learn the d -dimensional final embedding \mathbf{G} of each leaf node i (medical code) via:

$$\mathbf{G}_{i,:} = \sum_{j \in Pa_{\mathcal{G}}(i)} \alpha_{ij} \mathbf{E}_{j,:} \quad (3)$$

where $Pa_{\mathcal{G}}(i)$ denotes the set comprised of leaf node i and all its ancestors, $\mathbf{E}_{j,:}$ is the d -dimensional basic embedding of the node j and α_{ij} is the attention weight on the embedding $\mathbf{E}_{j,:}$ when calculating $\mathbf{G}_{i,:}$, which is formulated by following the Softmax function,

$$\alpha_{ij} = \frac{\exp(f(\mathbf{E}_{i,:}, \mathbf{E}_{j,:}))}{\sum_{k \in Pa_{\mathcal{G}}(i)} \exp(f(\mathbf{E}_{i,:}, \mathbf{E}_{k,:}))}. \quad (4)$$

$$f(\mathbf{E}_{i,:}, \mathbf{E}_{j,:}) = \mathbf{w}_{\alpha}^T \tanh(\mathbf{W}_{\alpha} [\mathbf{E}_{i,:}; \mathbf{E}_{j,:}] + \mathbf{b}_{\alpha}), \quad (5)$$

where $[\mathbf{E}_{i,:}; \mathbf{E}_{j,:}]$ is to concatenate $\mathbf{E}_{i,:}$ and $\mathbf{E}_{j,:}$ in the child-ancestor order, \mathbf{w}_{α} , \mathbf{W}_{α} and \mathbf{b}_{α} are learnable parameters.

D. Knowledgeable Encoder

The design of the integrator is inspired by NLP language modelling ERNIE [38]. However, our proposed model is distinct from ERNIE in three aspects: 1) node embeddings of a knowledge graph is part of our end-to-end MIPO model, while ERNIE uses pre-trained entity embedding from a knowledge graph by TransE [43]; 2) MIPO has a hierarchical structure, where the underlying knowledgeable encoder is for medical code level and the upper patient encoder is for visit level, while ERNIE is a derivative of BERT, which has not such structure; 3) MIPO aims to improve the predictive performance with the given knowledge graph as supplementary information.

In the i -th integrator, the input code embeddings $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n\}$ and node embedding $\{\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_n\}$ are fed into two different multi-head self-attentions (MSA) [42].

$$\begin{aligned} \{\tilde{\mathbf{w}}_1^{(i)}, \tilde{\mathbf{w}}_2^{(i)}, \dots, \tilde{\mathbf{w}}_n^{(i)}\} &= \text{MSA}(\{\mathbf{w}_1^{(i-1)}, \mathbf{w}_2^{(i-1)}, \dots, \mathbf{w}_n^{(i-1)}\}), \\ \{\tilde{\mathbf{g}}_1^{(i)}, \tilde{\mathbf{g}}_2^{(i)}, \dots, \tilde{\mathbf{g}}_n^{(i)}\} &= \text{MSA}(\{\mathbf{g}_1^{(i-1)}, \mathbf{g}_2^{(i-1)}, \dots, \mathbf{g}_n^{(i-1)}\}). \end{aligned} \quad (6)$$

Then, the i -th integrator adopts an information integration layer for the mutual integration of the code and node embedding in a visit, and computes the output embedding for each code and node. For a code \mathbf{w}_j and its corresponding node \mathbf{g}_j , the information integration process is as follows,

$$\begin{aligned} \mathbf{h}_j &= \sigma(\tilde{\mathbf{W}}_c^{(i)} \tilde{\mathbf{w}}_j^{(i)} + \tilde{\mathbf{W}}_g^{(i)} \tilde{\mathbf{g}}_j^{(i)} + \tilde{\mathbf{b}}^{(i)}), \\ \mathbf{w}_j^{(i)} &= \sigma(\mathbf{W}_c^{(i)} \mathbf{h}_j + \mathbf{b}_t^{(i)}), \\ \mathbf{g}_j^{(i)} &= \sigma(\mathbf{W}_g^{(i)} \mathbf{h}_j + \mathbf{b}_e^{(i)}). \end{aligned} \quad (7)$$

where \mathbf{h}_j is the inner hidden state integrating the information of both the code and the node. $\sigma(\cdot)$ is the non-linear activation function, which is usually the ReLU function.

For simplicity, the i -th integrator operation is denoted as follows,

$$\{\mathbf{w}_1^{(i)}, \dots, \mathbf{w}_n^{(i)}\}, \{\mathbf{g}_1^{(i)}, \dots, \mathbf{g}_n^{(i)}\} = \text{Integrator}(\{\mathbf{w}_1^{(i-1)}, \dots, \mathbf{w}_n^{(i-1)}\}, \{\mathbf{g}_1^{(i-1)}, \dots, \mathbf{g}_n^{(i-1)}\}). \quad (8)$$

The output embeddings of codes will be used by following attention pooling to compress a set of codes in a visit to a vector, with the output embeddings of nodes used to guarantee the proposed model can learn the reasonable knowledge from given medical ontology.

Note that we exclude position embedding in V-Encoder, as medical codes in a visit are not time-ordered.

E. Attention Pooling

Attention Pooling [40], [41] explores the importance of each individual code within a patient visit. It works by compressing a set of medical code embeddings from a patient visit into a single context-aware vector representation. Formally, it is written as,

$$f(\mathbf{w}_i^o) = \mathbf{w}^T \sigma(\mathbf{W}^{(1)} \mathbf{w}_i^o + \mathbf{b}^{(1)}) + \mathbf{b}, \quad (9)$$

where \mathbf{w}_i^o ($1 \leq i \leq n$) is one output of V-Encoder. The probability distribution is formalized as

$$\boldsymbol{\alpha} = \text{softmax}([f(\mathbf{w}_i^o)]_{i=1}^n). \quad (10)$$

The final output \mathbf{v} of the attention pooling is the weighted average of sampling a code according to its importance, i.e.,

$$\mathbf{v} = \sum_{i=1}^n \boldsymbol{\alpha} \odot [\mathbf{w}_i^o]_{i=1}^n. \quad (11)$$

F. Learning Healthcare Representation with Predictive Tasks

We jointly train the MIPO model with a *task-specific predictive task* and an *ontology-based disease typing task*, such that the mutual integration of knowledge graph and patient journey improves the performance of the healthcare representation learning.

1) *Task-specific Predictive Task*: Given a patient's visit records $\mathbf{P} = \{V_1, V_2, \dots, V_{T-1}\}$, to capture the EHRs sequential visit behaviour information, we perform the sequential diagnoses predictive task with the objective of predicting the disease codes of the next visit V_t , which can be expressed as follows,

$$\hat{\mathbf{y}}_{t-1}^P = \hat{\mathbf{v}}_t = \text{Softmax}(\mathbf{W}_P \mathbf{v}_{t-1}^o + \mathbf{b}_P), \quad (12)$$

$$\mathcal{L}_P(V_1, \dots, V_T) = \frac{1}{T-1} \sum_{t=1}^{T-1} \left(\mathbf{y}_t^{P^\top} \log \hat{\mathbf{y}}_t^P + (1 - \mathbf{y}_t^P)^\top \log (1 - \hat{\mathbf{y}}_t^P) \right). \quad (13)$$

where $\mathbf{v}_{t-1}^o \in \mathbb{R}^d$ is the output of P-Encoder to denote the representation of the $(t-1)$ -th visit, $\mathbf{W}_P \in \mathbb{R}^{|\mathcal{C}| \times d}$ and $\mathbf{b}_P \in \mathbb{R}^{|\mathcal{C}|}$ are the learnable parameters.

2) *Ontology-based Disease Typing Task*: To Enable MIPO to inject knowledge into healthcare representation by informative graph, we design the task using the output node embeddings of the knowledgeable encoder V-Encoder. This task is a multi-label prediction task. In particular, the non-leaf nodes located from the second layer in medical ontology \mathcal{G} are also known as the disease categories (or types), and each fine-grained diagnosis corresponds to the only disease category by finding its ancestor in the second layer.

As mentioned in Section III-A, knowledge graph \mathcal{G} contains the hierarchy of various medical concepts with the *parent-child* semantic relationship, and the medical codes \mathbb{C} come from its leaf nodes. Ideally, the disease categories in \mathcal{G} will acquire knowledge from the leaf nodes and represent more general medical concepts. Thus, we use the disease categories as targets of the task and the output embeddings of nodes of V-Encoder as input. To be specific, given the codes $V_t = \{c_1, c_2, \dots, c_n\}$ in a visit V_t , and its corresponding disease categories $\{n_1, n_2, \dots, n_m\}$ (shown in Figure 1) in multi-level hierarchy \mathcal{G} , where $m = 18$ for CCS Multi-level ontology, we define the disease categories distribution for the medical code c_i in V_t as follows,

$$\hat{\mathbf{y}}_{t,i}^V = \text{Softmax}(\mathbf{W}_V \mathbf{g}_{t,i}^o + \mathbf{b}_V), \quad (14)$$

where $\mathbf{g}_{t,i}^o \in \mathbb{R}^d$ is the output of V-Encoder, and t the t -th visit, i is the i -th code in t -th visit, $\mathbf{W}_V \in \mathbb{R}^{m \times d}$ and $\mathbf{b}_V \in \mathbb{R}^m$ are the learnable parameters.

Based on Equation 14, we use the cross-entropy between the ground truth visit $\mathbf{y}_{t,i}^V$ and the predicted visit $\hat{\mathbf{y}}_{t,i}^V$ to calculate the loss for each medical code from all the timestamps as follows:

$$\mathcal{L}_V(V_1, \dots, V_{T-1}) = \frac{1}{n(T-1)} \sum_{t=1}^{T-1} \sum_{i=1}^n \left(\mathbf{y}_{t,i}^{V^\top} \log \hat{\mathbf{y}}_{t,i}^V + (1 - \mathbf{y}_{t,i}^V)^\top \log (1 - \hat{\mathbf{y}}_{t,i}^V) \right). \quad (15)$$

where $T-1$ is the number of the patient's visits, and n is the number of medical codes in a visit.

3) *Objective Function*: In order to take advantage of the mutual integration of informative knowledge graph and sequential patient journey, we train the two tasks together to improve the performance of the healthcare representation learning, which can be formulated as follows,

$$\mathcal{L}(V_1, \dots, V_T) = \mathcal{L}_P(V_1, \dots, V_T) + \mathcal{L}_V(V_1, \dots, V_{T-1}). \quad (16)$$

Note that in our implementation, we take the average of the individual cross entropy error for multiple patients. Algorithm 1 describes the overall training procedure of the proposed MIPO with one individual patient journey.

Algorithm 1: The MIPO model

Input: Medical knowledge graph \mathcal{G} , the set of medical codes \mathcal{C} and Patient records $\mathbf{P} = \{V_1, V_2, \dots, V_{T-1}\}$

- 1 Initialize medical code embedding matrix \mathbf{W} ;
- 2 Knowledge graph embedding matrix \mathbf{G} via Eq. 3;
- 3 Initialize v-list to None and $\hat{\mathbf{y}}^K$ -list to None ;
- 4 **for** $t \leftarrow 1$ **to** $(T-1)$ **do**
- 5 $\mathbf{W}_t = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n\}$ # medical code embedding;
- 6 $\mathbf{G}_t = \{\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_n\}$ # graph node embedding;
- 7 $\{\mathbf{w}_1^o, \mathbf{w}_2^o, \dots, \mathbf{w}_n^o\}, \{\mathbf{g}_1^o, \mathbf{g}_2^o, \dots, \mathbf{g}_n^o\} =$
 V-Encoder($\mathbf{W}_t, \mathbf{G}_t$) via Eq. 1;
- 8 $\mathbf{v}_t = \text{Att-Pool}(\{\mathbf{w}_1^o, \mathbf{w}_2^o, \dots, \mathbf{w}_n^o\})$ via Eq. 11 #
 t -th visit representation;
- 9 Add \mathbf{v}_t to v-list;
- 10 Compute predicted first-level category $\hat{\mathbf{y}}_{t,i}^K$ via
 Eq. 14; Add $\hat{\mathbf{y}}_{t,i}^K$ to $\hat{\mathbf{y}}^K$ -list;
- 11 $\{\mathbf{v}_1^o, \mathbf{v}_2^o, \dots, \mathbf{v}_{T-1}^o\} = \text{P-Encoder}(\text{v-list})$ via
 Eq. 2;
- 12 Compute predicted sequential diagnoses $\hat{\mathbf{y}}^P$ via Eq. 12;
- 13 Update the model's parameters by optimizing the loss
 via Eq. 16 using $\hat{\mathbf{y}}^P$ and $\hat{\mathbf{y}}^K$ -list.

IV. EXPERIMENTS

A. Data Description

1) *MIMIC-III Dataset*: The MIMIC-III dataset [10] is an open-source, large-scale, de-identified dataset of ICU patients and their EHRs. The diagnosis codes in the dataset follow the ICD9 standard. The dataset consists of medical records of 7,499 intensive care unit (ICU) patients over 11 years, where we chose patients who had made at least two visits. We use MIMIC to represent MIMIC-III in the experiment.

2) *eICU Dataset*: The eICU dataset [44] is another publicly available EHRs dataset, which is a multi-center database comprising de-identified health data associated with over 200,000 admissions to ICUs across the United States between 2014-2015. The dataset consists of medical records of 16,180 ICU patients, where we follow MIMIC to choose patients who had made at least two visits.

Table I shows the statistical details about the two datasets. As the table shown, those two representative datasets can be

TABLE I: Statistics of the datasets.

Dataset	MIMIC	eICU
# of patients	7,499	16,180
# of visits	19,911	39,912
Avg. # of visits per patient	2.66	2.47
# of unique ICD9 codes	4,880	758
Avg. # of ICD9 codes per visit	13.06	5.21
Max # of ICD9 codes per visit	39	57
# of category codes	272	167
Avg. # of category codes per visit	11.23	4.72
Max # of category codes per visit	34	33
# of disease typing code	18	18
Avg. # of disease typing codes per visit	6.57	3.42
Max # of disease typing codes per visit	15	14

used to extensively evaluate different aspects of the models. The number of patients and visits in the eICU dataset is big enough to validate the performance of the proposed MIPO with long visit records. The MIMIC dataset consists of very short visits, and the number of patients is smaller. With these two different types of datasets, we can fully and correctly validate the performance of all the diagnosis prediction approaches.

B. Predictive Tasks

The proposed model consists of two predictive tasks to simultaneously learn the integration between knowledge graph and sequential patient journey.

A task-specific predictive task is to predict the diagnosis information of the next visit. In the experiments, true labels y_t^P are prepared by grouping the ICD9 codes into 283 groups using CCS single-level diagnosis grouper². This aims to improve the training speed and predictive performance, while preserving sufficient granularity for all the diagnoses. The second hierarchy of the ICD9 codes can also be used as category labels [5]. These two grouping methods obtain similar predictive performances.

An ontology-based disease typing task is to predict the disease category given the medical code (leaf node). The disease categories come from “CCS_LVL_1” of CCS multi-level diagnosis grouper³, which groups the ICD9 codes into 18 categories. In the experiments, we prepared the 18 categories as true labels $y_{t,i}^V$. This is to guarantee parent nodes learn general knowledge from their children, following the *parent-child* semantic relationship. The results will be shown in Section IV-F.

C. Experimental Setup

1) *Baseline Approaches*: We compare the performance of our proposed model against the following baseline models:

- **Graph-based Models** include GRAM [4] and KAME [5]. They incorporate the medical ontology with an attention mechanism and recurrent neural networks for representation learning with the application to diagnosis prediction.

- **Attention-based Models** include Dipole [12] which uses the bidirectional GRU as the backbone and assigns an attention weight for each visit and RETAIN [13] which learns the medical concept embeddings and performs heart failure prediction via the reversed RNN with the attention mechanism.

- **Plain RNNs** include BRNN which is the basic framework of most risk prediction models with the bidirectional GRU as the backbone.

2) *Evaluation Measures*: We measure the predictive performance by $Prec@k$ and $Acc@k$, which are defined as:

$$Prec@k = \frac{\# \text{ of true positives in the top } k \text{ predictions}}{\min(k, \# \text{ of positives})}$$

$$Acc@k = \frac{\# \text{ of true positives in the top } k \text{ predictions}}{\# \text{ of positives}}$$

We report the average values of $Prec@k$ and $Acc@k$ and vary k from 5 to 30 in the experiments, where $Prec@k$ aims to evaluate the coarse-grained performance, and $Acc@k$ is proposed to evaluate the fine-grained performance [5]. For all the measures, greater values reflect better performance.

D. Results of Diagnosis Prediction

Table II shows both the precision and accuracy of the proposed MIPO and baselines with different k on two real-world datasets for task-specific predictive task. From Table II, we can observe that the performance of the proposed MIPO, in both precision and accuracy, is better than that of all the baselines on the two datasets.

On the MIMIC dataset, compared with KAME and GRAM, the precision of MIPO improves 5.94% and 5.76% with accuracy improving 3.11% and 2.78% when $k = 5$, respectively. These results suggest that it is effective to integrate medical knowledge and sequential patient journey when predicting diagnoses. Comparably, Dipole, RETAIN and BRNN do not use external knowledge in the diagnosis prediction task. Dipole and RETAIN directly learn the medical code embeddings from the input data with location-based attention mechanisms, and BRNN learns the code embeddings from the input data with bi-directional RNN. Compared with KAME and GRAM, the performances of Dipole, RETAIN and BRNN are lower, indicating that employing knowledge graph is effective with data insufficiency. However, instead of adding attention mechanisms on the past visits like Dipole and RETAIN, and simply integrating medical knowledge into visits like KAME and GRAM, the proposed MIPO aims to integrate the given knowledge graph and sequential patient journeys to improve predictive performance.

Though the number of visits and patients on the eICU is larger than that on the MIMIC dataset, the number of labels observed are much less. On this significantly insufficient dataset, MIPO still outperforms all the baselines. In the five described, Dipole, RETAIN and BRNN achieves a better performance than KAME and GRAM, which suggests that with

²<https://www.hcup-us.ahrq.gov/toolssoftware/ccs/AppendixASingleDX.txt>

³<https://hcup-us.ahrq.gov/toolssoftware/ccs/AppendixCMultiDX.txt>

TABLE II: Performance comparison of sequential diagnoses prediction.

Dataset	Model	Prec@k						Acc@k					
		5	10	15	20	25	30	5	10	15	20	25	30
MIMIC	BRNN	0.5707	0.5112	0.5270	0.5718	0.6234	0.6690	0.2692	0.4028	0.4933	0.5636	0.6220	0.6690
	RETAIN	0.5769	0.5071	0.5280	0.5700	0.6214	0.6721	0.2721	0.3976	0.4936	0.5617	0.6201	0.6721
	Dipole	0.5750	0.5104	0.5334	0.5813	0.6303	0.6753	0.2753	0.4028	0.5000	0.5732	0.6290	0.6753
	GRAM	0.5870	0.5248	0.5498	0.6024	0.6523	0.6956	0.2792	0.4210	0.5211	0.5954	0.6507	0.6955
	KAME	0.5852	0.5195	0.5389	0.5873	0.6384	0.6799	0.2759	0.4164	0.5111	0.5808	0.6370	0.6799
	MIPO	0.6446	0.5661	0.5813	0.6305	0.6752	0.7189	0.3070	0.4522	0.5502	0.6229	0.6739	0.7188
eICU	BRNN	0.6221	0.7011	0.7756	0.8229	0.8620	0.8845	0.5480	0.6892	0.7733	0.8226	0.8619	0.8845
	RETAIN	0.6332	0.7124	0.7796	0.8277	0.8655	0.8907	0.5571	0.7000	0.7772	0.8274	0.8654	0.8907
	Dipole	0.6264	0.7018	0.7696	0.8255	0.8610	0.8898	0.5514	0.6895	0.7673	0.8252	0.8609	0.8898
	GRAM	0.6048	0.6846	0.7571	0.8101	0.8485	0.8791	0.5277	0.6719	0.7549	0.8098	0.8485	0.8791
	KAME	0.6004	0.6795	0.7509	0.8093	0.8466	0.8770	0.5226	0.6668	0.7487	0.8090	0.8465	0.8770
	MIPO	0.6848	0.75202	0.8127	0.8532	0.8847	0.9086	0.5986	0.73852	0.8102	0.8527	0.8845	0.9086




Fig. 2: Acc@20 of diagnoses prediction on MIMIC and eICU, size of training data is varied from 20% to 80%.

enough data, even without external knowledge, attention-based models can still learn reasonable medical code embeddings to make accurate predictions. However, compared with the proposed MIPO, the precision and accuracy of these three approaches are lower, which again argues that integration of medical knowledge and sequential patient journeys can improve prediction performance. The performance of KAME is the weakest since this approach explicitly incorporates knowledge from leaf nodes and parent nodes, which cannot adequately balance the knowledge and sequential visits. However, the proposed model learns the healthcare representations by taking advantage of task-specific predictive task and ontology-based disease typing task to harmoniously fuse medical knowledge and sequential patient journey.

E. Data Sufficiency Analysis

In order to analyze the influence of data sufficiency on the predictions, we conduct the following experiments on the MIMIC and eICU datasets, respectively. We randomly split the data into training set, validation set and test set, and fix the size of the validation set at 10%. To validate robustness against insufficient data, we vary the size of the training set to form four groups: 20%, 40%, 60% and 80%, and use the remaining part as the test set. The training set in the 20% group is the most insufficient for training the proposed and baseline models, while the data in the 80% group are the most sufficient for training models. Finally, we calculate the accuracy of labels in each group. Figures 2 show the Acc@20 on both the MIMIC and eICU datasets. Note that similar results can be obtained when $k = 5, 10, 15, 25$ or 30 .

From Figure 2, we can observe that the accuracy of the proposed MIPO is higher than that of baselines in all groups on both MIMIC and eICU datasets. KAME and GRAM achieve better performances on MIMIC than other approaches on MIMIC, which shows that, with insufficient data, KAME and GRAM still learn reasonable medical code embeddings and improve predictions. The performance of BRNN in the groups 20%, 40%, 60% is the worst since this approach does not use any attention mechanism or external knowledge. When the training data on the eICU dataset is significantly insufficient, the proposed MIPO still significantly outperforms baselines in all groups. We observe that the performance obtained by the models using medical knowledge remains approximately the same (GRAM) or even drop (KAME). The underlying reason may be that KAME and GRAM over-fit the insufficient data using the medical knowledge. Thus, the models learn larger weighting for knowledge than with sequential visits. Furthermore, as shown in Figure 2b, the average accuracy of Dipole, RETAIN and BRNN is better than that of both KAME and GRAM, indicating that information of sequential visits plays a more important role under insufficient data. These observations can also be found in Table II. It is again to demonstrate that the proposed MIPO harmoniously balances medical knowledge and patient journeys when the EHRs data is insufficient.

F. Interpretable Representation Analysis

To qualitatively demonstrate the interpretability of the learned medical code representations by all the predictive models on the MIMIC dataset, we randomly select 2000 medical codes and then plot on a 2-D space with t-SNE [45] shown in Figures 3 and 4. Each dot represents a diagnosis code, and their color represents the disease categories while the text annotations represent the detailed disease categories in CCS multi-level hierarchy.

From Figure 3, we can observe that MIPO learns interpretable disease representations that are in accord with the hierarchies of the given knowledge graph \mathcal{G} , and obtains 18 non-overlapping clusters. As shown in Figure 4, KAME and GRAM learn reasonably interpretable disease representations for partial categories, as there is large number of dots overlapping in the centers of Figures 4b and 4c. Figures 4d, 4e and 4f confirm that without a knowledge graph, simply using

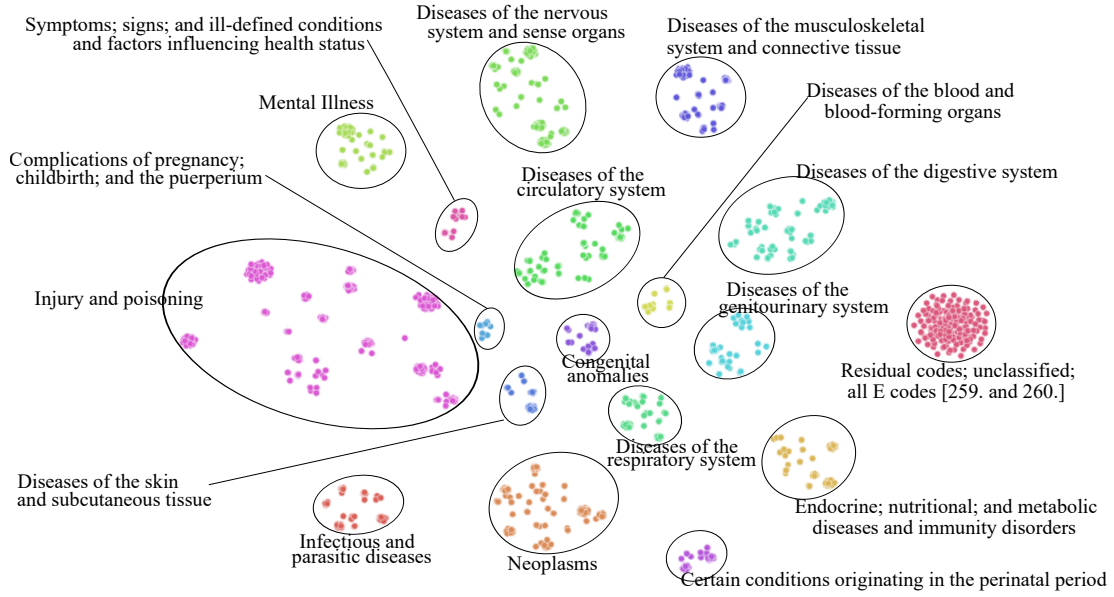


Fig. 3: Annotations of MIPO Diagnosis Embedding




Fig. 4: t -SNE Scatterplots of Medical Codes Learned by Predictive Models on the MIMIC dataset.

the co-occurrence or supervised predictions cannot easily provide for learning interpretable representations. In addition, the predictive performance of MIPO is much better than that of KAME and GRAM, as shown in Table II, which proves that the proposed model does not affect the interpretability of medical codes. Moreover, it significantly improves the prediction accuracy.

V. CONCLUSIONS

This paper presents MIPO, a novel approach that integrates medical knowledge and the patient journey to enhance health-

care representation. By introducing a knowledgeable encoder and employing two predictive tasks—sequential diagnosis prediction and disease category classification—MIPO effectively captures heterogeneous information from both patient records and a knowledge graph. Empirical evaluations on two real-world medical datasets reveal that MIPO not only outperforms baseline methods but also remains robust in scenarios with limited data. Moreover, the visualizations of medical code representations offer insights into MIPO’s interpretability, underscoring its potential for real-world healthcare applications.

REFERENCES

- [1] B. Shickel, P. J. Tighe, A. Bihorac, and P. Rashidi, "Deep ehr: a survey of recent advances in deep learning techniques for electronic health record (ehr) analysis," *IEEE J Biomed Health Inform.*, vol. 22, no. 5, pp. 1589–1604, 2018.
- [2] L. Song, C. W. Cheong, K. Yin, W. K. Cheung, and B. Cm, "Medical concept embedding with multiple ontological representations," in *IJCAI*, 2019, pp. 4613–4619.
- [3] E. Choi, M. T. Bahadori, E. Searles, C. Coffey, M. Thompson, J. Bost, J. Tejedor-Sojo, and J. Sun, "Multi-layer representation learning for medical concepts," in *SIGKDD*. ACM, 2016, pp. 1495–1504.
- [4] E. Choi, M. T. Bahadori, L. Song, W. F. Stewart, and J. Sun, "Gram: graph-based attention model for healthcare representation learning," in *SIGKDD*. ACM, 2017, pp. 787–795.
- [5] F. Ma, Q. You, H. Xiao, R. Chitta, J. Zhou, and J. Gao, "KAME: Knowledge-based attention model for diagnosis prediction in healthcare," in *CIKM*. ACM, Oct. 2018, pp. 743–752.
- [6] J. Luo, M. Ye, C. Xiao, and F. Ma, "Hitanet: Hierarchical time-aware attention networks for risk prediction on electronic health records," in *SIGKDD*. ACM, 2020, pp. 647–656.
- [7] X. Peng, G. Long, T. Shen, S. Wang, and J. Jiang, "Sequential diagnosis prediction with transformer and ontological representation," in *ICDM*. IEEE, 2021, pp. 489–498.
- [8] P. Mukherjee, M. Humbert-Droz, J. H. Chen, and O. Gevaert, "Scope: predicting future diagnoses in office visits using electronic health records," *Scientific Reports*, vol. 13, no. 1, p. 11005, 2023.
- [9] H. Ren, J. Wang, W. X. Zhao, and N. Wu, "Rapt: Pre-training of time-aware transformer for learning robust healthcare representation," in *SIGKDD*. ACM, 2021, pp. 3503–3511.
- [10] A. E. Johnson, T. J. Pollard, L. Shen, H. L. Li-wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "Mimic-iii, a freely accessible critical care database," *Scientific data*, vol. 3, p. 160035, 2016.
- [11] X. Zhang, B. Qian, X. Li, J. Wei, Y. Zheng, L. Song, and Q. Zheng, "An interpretable fast model for predicting the risk of heart failure," in *SDM*. SIAM, 2019, pp. 576–584.
- [12] F. Ma, R. Chitta, J. Zhou, Q. You, T. Sun, and J. Gao, "Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks," in *SIGKDD*. ACM, Aug. 2017, pp. 1903–1911.
- [13] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, "Retain: An interpretable predictive model for healthcare using reverse time attention mechanism," in *NeurIPS*, 2016, pp. 3504–3512.
- [14] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv:1301.3781*, 2013.
- [15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [16] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *arXiv:1406.1078*, 2014.
- [17] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *NeurIPS*, 2013, pp. 3111–3119.
- [18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [19] X.-y. Wang, T. Yang, X.-y. Gao, and K.-f. Hu, "Knowledge graph enhanced transformers for diagnosis generation of chinese medicine," *Chinese journal of integrative medicine*, vol. 30, no. 3, pp. 267–276, 2024.
- [20] Y. Li, S. Rao, J. R. A. Solares, A. Hassaine, R. Ramakrishnan, D. Canoy, Y. Zhu, K. Rahimi, and G. Salimi-Khorshidi, "Behrt: transformer for electronic health records," *Scientific Reports*, vol. 10, no. 1, pp. 1–12, 2020.
- [21] F. Li, Y. Jin, W. Liu, B. P. S. Rawat, P. Cai, and H. Yu, "Fine-tuning bidirectional encoder representations from transformers (bert)-based models on large-scale electronic health record notes: An empirical study," *JMIR medical informatics*, vol. 7, no. 3, p. e14830, 2019.
- [22] J. Gao, C. Xiao, Y. Wang, W. Tang, L. M. Glass, and J. Sun, "Stagenet: Stage-aware neural networks for health risk prediction," in *The Web Conference*, 2020, pp. 530–540.
- [23] T. Bai and S. Vucetic, "Improving medical code prediction from clinical text via incorporating online knowledge sources," in *The Web Conference*, 2019, pp. 72–82.
- [24] H. Song, D. Rajan, J. J. Thiagarajan, and A. Spanias, "Attend and diagnose: Clinical time series analysis using attention models," in *AAAI*, 2018.
- [25] X. Zeng, Y. Feng, S. Moosavinasab, D. Lin, S. Lin, and C. Liu, "Multilevel self-attention model and its use on medical risk prediction," in *Pac Symp Biocomput.* World Scientific, 2020, pp. 115–126.
- [26] I. M. Baytas, C. Xiao, X. Zhang, F. Wang, A. K. Jain, and J. Zhou, "Patient subtyping via time-aware lstm networks," in *SIGKDD*. ACM, 2017, pp. 65–74.
- [27] T. Pham, T. Tran, D. Phung, and S. Venkatesh, "Deepcare: A deep dynamic memory model for predictive medicine," in *PAKDD*. Springer, 2016, pp. 30–41.
- [28] T. Ruan, L. Lei, Y. Zhou, J. Zhai, L. Zhang, P. He, and J. Gao, "Representation learning for clinical time series prediction tasks in electronic health records," *BMC Medical Informatics and Decision Making*, vol. 19, no. 8, p. 259, 2019.
- [29] P. Nguyen, T. Tran, N. Wickramasinghe, and S. Venkatesh, "Deepr: A convolutional net for medical records," 2016.
- [30] F. Ma, J. Gao, Q. Suo, Q. You, J. Zhou, and A. Zhang, "Risk prediction on electronic health records with prior medical knowledge," in *SIGKDD*. ACM, Jul. 2018, pp. 1910–1919.
- [31] D. Lee, H. Yu, X. Jiang, D. Rogith, M. Gudala, M. Tejjani, Q. Zhang, and L. Xiong, "Generating sequential electronic health records using dual adversarial autoencoder," *Journal of the American Medical Informatics Association*, vol. 27, no. 9, pp. 1411–1419, 2020.
- [32] S. Niu, Q. Yin, Y. Song, Y. Guo, and X. Yang, "Label dependent attention model for disease risk prediction using multimodal electronic health records," in *ICDM*. IEEE, 2021, pp. 449–458.
- [33] C. Zang and F. Wang, "Scehr: Supervised contrastive learning for clinical risk prediction using electronic health records," in *ICDM*. IEEE, 2021, pp. 857–866.
- [34] X. Peng, G. Long, S. Pan, J. Jiang, and Z. Niu, "Attentive dual embedding for understanding medical concepts in electronic health records," in *IJCNN*. IEEE, 2019, pp. 1–8.
- [35] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang, "Biobert: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
- [36] K. Huang, J. Altaaar, and R. Ranganath, "Clinicalbert: Modeling clinical notes and predicting hospital readmission," *arXiv preprint arXiv:1904.05342*, 2019.
- [37] L. Rasmy, Y. Xiang, Z. Xie, C. Tao, and D. Zhi, "Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction," *NPJ digital medicine*, vol. 4, no. 1, pp. 1–13, 2021.
- [38] Z. Zhang, X. Han, Z. Liu, X. Jiang, M. Sun, and Q. Liu, "Ernie: Enhanced language representation with informative entities," *arXiv preprint arXiv:1905.07129*, 2019.
- [39] W. Liu, P. Zhou, Z. Zhao, Z. Wang, Q. Ju, H. Deng, and P. Wang, "K-bert: Enabling language representation with knowledge graph," 2019.
- [40] Z. Lin, M. Feng, C. N. d. Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio, "A structured self-attentive sentence embedding," *arXiv:1703.03130*, 2017.
- [41] Y. Liu, C. Sun, L. Lin, and X. Wang, "Learning natural language inference using bidirectional lstm model and inner-attention," *arXiv:1605.09090*, 2016.
- [42] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017, pp. 5998–6008.
- [43] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," in *NeurIPS*, 2013, pp. 2787–2795.
- [44] T. J. Pollard, A. E. Johnson, J. D. Raffa, L. A. Celi, R. G. Mark, and O. Badawi, "The eicu collaborative research database, a freely available multi-center database for critical care research," *Scientific data*, vol. 5, p. 180178, 2018.
- [45] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.