BiteNet: Bidirectional Temporal Encoder Network to Predict Medical Outcomes

Xueping Peng*, Guodong Long*, Tao Shen*, Sen Wang[†], Jing Jiang*, Chengqi Zhang*

* Australian Artificial Intelligence Institute, FEIT, University of Technology Sydney, Australia

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

Email: {xueping.peng, guodong.long}@uts.edu.au, Tao.Shen@student.uts.edu.au, sen.wang@uq.edu.au, {jing.jiang, chengqi.zhang}@uts.edu.au

Abstract—Electronic health records (EHRs) are longitudinal records of a patient’s interactions with healthcare systems. A patient’s EHR data is organized as a three-level hierarchy from top to bottom: patient journey - all the experiences of diagnoses and treatments over a period of time; individual visit - a set of medical codes in a particular visit; and medical code - a specific record in the form of medical codes. As EHRs begin to amass in millions, the potential benefits, which these data might hold for medical research and medical outcome prediction, are staggering - including, for example, predicting future admissions to hospitals, diagnosing illnesses or determining the efficacy of medical treatments. Each of these analytics tasks requires a domain knowledge extraction method to transform the hierarchical patient journey into a vector representation for further prediction procedure. The representations should embed a sequence of visits and a set of medical codes with a specific timestamp, which are crucial to any downstream prediction tasks. Hence, expressively powerful representations are appealing to boost learning performance. To this end, we propose a novel self-attention mechanism that captures the contextual dependency and temporal relationships within a patient’s healthcare journey. An end-to-end bidirectional temporal encoder network (BiteNet) then learns representations of the patient’s journeys, based solely on the proposed attention mechanism. We have evaluated the effectiveness of our methods on two supervised prediction and two unsupervised clustering tasks with a real-world EHR dataset. The empirical results demonstrate the proposed BiteNet model produces higher-quality representations than state-of-the-art baseline methods.

Index Terms—Electronic health records, Bidirectional encoder, Transformer, Embedding, Patient journey

I. INTRODUCTION

Healthcare information systems store huge volumes of electronic health records (EHRs) that contain detailed visit information about patients over a period of time [1]. The data is structured in three levels from top to bottom: the patient journey, the individual visit and the medical code. Fig. 1 provides a typical example of this structure. An anonymous patient visits his/her doctor, a pathology lab and is admitted to the hospital on different days. The procedures and diagnoses performed at each of these visits are recorded as industry-standard medical codes. Each medical code, i.e. International Classification of Diseases (ICD) and Current Procedure Terminology (CPT), at the lowest level, records an independent observation while the set of codes at a higher level can depict the medical conditions of a patient at a given time point. At the top level, all occurrences of medical events at different time-stamps

Fig. 1. An example of a patient’s journey that is hierarchically structured in three levels: patient journey, individual visit and medical code.

are chained together as a patient journey, which offers more informative details. Predicting sequential medical outcomes based on a patient’s journey, such as hospital re-admissions and diagnoses, is a core research task that significantly benefits for healthcare management by hospitals and governments. For example, re-admission statistics could be used to measure the quality of care; Diagnoses can be used to understand more fully a patient’s problems and relevant medical research [2]. However, researchers have encountered many challenges in their attempts to represent patient journeys and predict medical outcomes from EHR data with the characteristics of temporality, high-dimensionality and irregularity [3].

Recurrent neural networks (RNNs) have been widely used to analyze sequential data, unsurprisingly including medical events modelling for clinical prediction [3]–[7]. For example, Choi et al. [4], [6] proposed a multi-level representation learning, which integrates visits and medical concepts based on visit sequences and the co-occurrence of medical concepts. They indirectly exploited an RNN to embed the visit sequences into a patient representation for downstream prediction tasks. Some other research works directly employed RNNs to model time-ordered patient visits for predicting diagnoses [3], [5], [7], [8]. However, when the length of the patient visit sequence grows, such RNN-based models are restricted by the less expressive power of RNNs, such as vanishing gradient and forgetfulness. However, such RNN-based models are constrained by forgetfulness, i.e., their predictive power drops

significantly when the sequence of patient visits grows too long. To memorize historical records, LSTM [9] and GRU [10] have been developed to utilize memory and gate mechanism for mitigating these issues. To go further, Song et al. [11] proposed to utilize attention mechanism in a deep framework to model sequential medical events. It is worth noting that sequences of medical events are often found to be lengthy [12], especially when a patient suffers from chronic disease. Hence, due to the restricted ability of RNNs for long-term dependency modeling [13], the traditional RNNs, even with memory cells and gates, usually underperform in the cases of a long sequence of medical events. In light of this, a neural model that can overcome the performance bottleneck of RNN-based models is particularly desirable for medical predictions based on longitudinal EHR data.

Most recently, attention mechanisms [14] have sprung to the fore as effective integrations with RNNs for modeling sequential EHR data [2], [7], [8], [15]. So far, these approaches have shown satisfactory prediction accuracy, but some argue that the power of attention in an RNN is limited by weaknesses in the RNN itself [13], [16]. In particular, Vaswani et al. [13] used a sole attention mechanism, i.e., multi-head attention and self-attention, to construct a sequence-to-sequence model for neural machine translation tasks and achieved a state-of-the-art quality score. And according to Shen et al. [16], self-attention mechanism allows for more flexibility in sequence lengths than RNNs and is more task/data-driven when modeling contextual dependencies. Unlike recurrent models, attention procedure is easy to compute and the computation can also be significantly accelerated with distributed/parallel computing schemes. For example, Song et al. [11] proposed to employ 1D CNN [17] to model local context and use attention mechanism [13] to capture long-term dependency for sequential medical events. However, when applied to EHR data instead of regular sequential data (e.g., natural language text), the current attention models cannot appropriately deal with some aspects of EHR data, such as arbitrary time-stamps and hierarchical data format. Hence, to the best of our knowledge, a neural network-based entirely on attention has never been designed for analytics with EHR data.

To bridge the gap in this literature and address some of the open issues listed above, we propose a novel attention mechanism called **Masked Encoder** (MasEnc) for temporal context fusion. It uses self-attention to capture contextual information and temporal dependencies between a patient's visits. Then, we propose an end-to-end neural network, called **Bidirectional temporal encoder Network** (BiteNet), to predict medical outcomes by leveraging a learned representation of the patient journey, where the representation is generated solely by the proposed attention mechanism, MasEnc. BiteNet constructs a multi-level self-attention network to represent visits and patient journeys simultaneously, using attention pooling and stacked MasEnc layers. It is worth noting that, compared to the existed RNN-based methods, BiteNet can yield better prediction performance for long sequences of medical records.

Experiments conducted on two supervised prediction and

two unsupervised clustering tasks with real-world EHR datasets demonstrate that the proposed BiteNet model is superior to prior state-of-the-art baseline methods.

To summarize, our main contributions are:

- a novel attention mechanism, called MasEnc, that uses self-attention to capture the contextual information and long-term dependencies between patients' visits;
- an end-to-end neural network, called BiteNet, that predicts medical outcomes using a learned representation of a patient journey based solely on the proposed attention mechanism;
- evaluations of the proposed model on a real-world dataset with supervised and unsupervised tasks, demonstrating that the BiteNet is superior to all the comparative methods.

II. RELATED WORK

Applying deep learning to healthcare analytical tasks has recently attracted enormous interest in healthcare informatics. This section reviews two types of related studies, which are Medical Concept Embedding and patient journey embedding in EHRs.

A. Medical Concept Embedding

Word representation models are first introduced from [18]–[20]. Borrowing ideas from those models, researchers in the healthcare domain have recently explored the possibility of creating representations of medical concepts [21]. Much of this research has focused on the Skip-gram model. For example, Minarro-Gimnez et al. [22] directly applied Skip-gram to learn representations of medical text, and Vine et al. [23] did the same for UMLS medical concepts. Choi et al. [24] went a step further and used the Skip-gram model to learn medical concept embeddings from different data sources, including medical journals, medical claims, and clinical narratives. In other work [4], Choi et al. developed the Med2Vec model based on Skip-gram to learn concept-level and visit-level representations simultaneously. The shortcoming of all these models is that they view EHRs as documents in the NLP sense, which means that temporal information is ignored.

B. Patient Journey Embedding

Applying deep learning to healthcare analytical tasks has recently attracted enormous interest in healthcare informatics. RNN has been widely used to model medical events in sequential EHR data for clinical prediction tasks [3]–[7], [25]–[27]. Choi et al. [4], [6] indirectly exploit an RNN to embed the visit sequences into a patient representation by multi-level representation learning to integrate visits and medical codes based on visit sequences and the co-occurrence of medical codes. Other research works have, however, used RNNs directly to model time-ordered patient visits for predicting diagnoses [3], [5], [7], [8], [25], [28], [29]. CNN has been exploited to represent a patient journey in other way. For example, Nguyen et al. [30] transform a record into a sequence of discrete elements separated by coded time gaps, and then employ CNN

to detect and combine predictive local clinical motifs to stratify the risk. Song et al. [11] use CNN in code level to learn visit embedding. These RNN- and CNN-based models follow ideas of processing sentences [17] in documents from NLP to treat a patient journey as a document and a visit as a sentence, which only has a sequential relationship, while two arbitrary visits in one patient journey may be separated by different time intervals, an important factor in longitudinal studies. Attention-based neural networks have been exploited successfully in healthcare tasks to model sequential EHR data [2], [7], [8], [11] and have been shown to improve predictive performance.

III. BACKGROUND

This section starts with definitions of several important concepts and tasks in the paper. The remainder mainly focuses on how we adapt embeddings and attention mechanism from natural language processing (NLP) to EHR data and analytics with patient journeys.

A. Definitions and Notations

Definition 1 (Medical Code): A medical code is a term or entry to describe a diagnosis, procedure, medication, or laboratory test administered to a patient. A set of medical codes is formally denoted as $X = \{x_1, x_2, \dots, x_{|X|}\}$, where $|X|$ is the total number of unique medical codes in the EHR data.

Definition 2 (Visit): A visit is a hospital stay from admission to discharge with an admission time stamp. A visit is denoted as $V_i = \langle x_i, t_i \rangle$, where $x_i = [x_{i1}, x_{i2}, \dots, x_{ik}]$, i is the i -th visit, t_i is the admission time of the i -th visit, k is the total number of medical codes in a visit.

Definition 3 (Patient Journey): A patient journey consists of a sequence of visits over time, denoted as $J = [V_1, V_2, \dots, V_n]$, where n is the total number of visits for the patient.

Definition 4 (Temporal Interval): Temporal interval refers to a difference in days between admission time t_i of the i -th visit and admission time t_1 of the first visit in a patient journey, denoted as $p_i = |t_i - t_1|$, where $i = 1, \dots, n$.

Definition 5 (Task): Given a set of patient journeys $\{J_l\}_{l=1, \dots, L}$, the task is to predict medical outcomes.

We choose to predict diagnoses and future hospital re-admissions as examples. However, many other outcomes could be predicted.

In addition, a patient's medical data means a stored and chronological sequence of n visits in a patient journey J_l . To reduce clutter, we omit the subscript (l) indicating l -th patient, when discussing a single patient journey.

Table I summarizes notations we will use throughout the paper.

B. Medical Code Embedding

The concept of word embedding was introduced to medical analytics by Mikolov et al. [19], [20] as a way to learn low-dimensional real-valued distributed vector representations of medical codes for downstream tasks instead of using discrete medical codes. This makes medical code embedding

TABLE I
NOTATIONS FOR BITENET.

Notation	Description
X	Set of unique medical codes
$ X $	The size of unique medical concept
V_i	The i -th visit of the patient
v_i	The representation of t -th visit of the patient
$\mathbf{v}$	Sequence of n visit embeddings of the patient
x_i	Set of medical codes in V_i
$\mathbf{x}$	Sequence of medical codes, $[x_1, x_2, \dots, x_n]$
e_i	Set of medical code embeddings in $\mathbf{x}_i$
$\mathbf{e}$	A sequence of medical code embeddings, $[e_1, e_2, \dots, e_n]$
J	A patient journey consisting of a sequence of visits over time
d	The dimension of medical code embedding

a fundamental processing unit for learning EHRs in a deep neural network. Formally, given a sequence or set of medical concepts $\mathbf{x} = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{|X| \times n}$, where x_i is a one-hot vector, and n is the sequence length. In the NLP literature, a word embedding method like word2vec [19], [20] converts a sequence of one-hot vectors into a sequence of low-dimensional vectors $\mathbf{e} = [e_1, e_2, \dots, e_n] \in \mathbb{R}^{d \times n}$, where d is the embedding dimension of e_i . This process can be formally written as $\mathbf{e} = W^{(e)}\mathbf{x}$, where $W^{(e)} \in \mathbb{R}^{d \times |X|}$ is the embedding weight matrix, which can be fine-tuned during the training phase.

C. Visit Embedding

As stated in Definition 2, a visit consists of a set of medical codes. Therefore, a visit can be represented by the set of medical codes embedded with real-valued dense vectors. A straightforward approach to learning this representation v_i is to sum the embeddings of medical codes in the visit, which is written as

$$v_i = \sum_{e_{ik} \in e_i} e_{ik}, \quad (1)$$

where e_i is the set of medical code embeddings in the i -th visit, and e_{ik} is the k -th code embedding in e_i . A visit can also be represented as a real-valued dense vector with a more advanced method, such as self-attention and attention pooling discussed below.

D. Attention Mechanism

1) *Vanilla Attention:* Given a patient journey consisting of a sequence of visits $\mathbf{v} = [v_1, v_2, \dots, v_m]$ ($\mathbf{v} \in \mathbb{R}^{m \times d}$) and a vector representation of a query $q \in \mathbb{R}^d$, vanilla attention [14] computes the alignment score between q and each visit v_i using a compatibility function $f(v_i, q)$. A softmax function then transforms the alignment scores $\alpha \in \mathbb{R}^n$ to a probability distribution $p(z|\mathbf{v}, q)$, where z is an indicator of which visit is important to q . A large $p(z = i|\mathbf{v}, q)$ means that v_i contributes

important information to q . This attention process can be formalized as

$$\alpha = [f(v_i, q)]_{i=1}^n, \quad (2)$$

$$p(z|v, q) = \text{softmax}(\alpha). \quad (3)$$

The output s denotes the weighted visits according to their importance, i.e.,

$$s = p(z|v, q) \odot v. \quad (4)$$

Additive attention is a commonly-used attention mechanism [14], [31], where a compatibility function $f(\cdot)$ is parameterized by a multi-layer perceptron (MLP), i.e.,

$$f(v_i, q) = w^T \sigma(W^{(1)}v_i + W^{(2)}q + b^{(1)}) + b, \quad (5)$$

where $W^{(1)} \in \mathbb{R}^{d \times d}$, $W^{(2)} \in \mathbb{R}^{d \times d}$, $w \in \mathbb{R}^d$ are learnable parameters, and $\sigma(\cdot)$ is an activation function.

Additive attention tends to be memory-intensive with long runtimes. However, it usually produces better representations for downstream tasks than multiplicative attention [32], [33], with cosine similarity as the compatibility function.

2) *Attention Pooling*: Attention Pooling [34], [35] explores the importance of each individual visit within an entire patient journey, given a specific task. It works by compressing a sequence of visit embeddings from a patient journey into a single context-aware vector representation for downstream classification or regression. Formally, q is removed from the common compatibility function, which is written as,

$$f(v_i) = w^T \sigma(W^{(1)}v_i + b^{(1)}) + b. \quad (6)$$

The probability distribution is formalized as

$$\alpha = [f(v_i)]_{i=1}^n, \quad (7)$$

$$p(z|v) = \text{softmax}(\alpha). \quad (8)$$

The final output s of self-attention is similar to the vanilla attention mechanism above, i.e.,

$$s = p(z|v) \odot v \quad (9)$$

If the output is used as sequence encoding, s is the weighted average of sampling a visit according to its importance, i.e.,

$$s = \sum_{i=1}^n p(z=i|v) \odot v_i. \quad (10)$$

3) *Multi-Head Attention*: The multi-head attention mechanism [13] relies on self-attention, which operates on a query Q , a key K and a value V :

$$\text{Attention}(Q, K, V) = p(z|Q, K)V = p(z|v)v \quad (11)$$

where Q, K, V, v is a $m \times d$ matrix, m denotes the number of visits in a patient journey, d denotes the dimension of embedding. v is a sequence of visit embeddings coming from the output of the previous layer.

Fig. 2. Masked encoder (MasEnc) block.

For scaled dot-product attention,

$$\alpha = \frac{QK^T}{\sqrt{d}}, \quad (12)$$

$$p(z|Q, K) = p(z|v) = \text{softmax}(\alpha). \quad (13)$$

The multi-head attention mechanism obtains h (i.e. one per head) different representations of (Q, K, V) , computes self-attention for each representation, concatenates the results. This can be expressed in the same notation as Eq.(9):

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \quad (14)$$

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \quad (15)$$

where the projections are parameter matrices $W_i^Q \in \mathbb{R}^{d \times d_k}$, $W_i^K \in \mathbb{R}^{d \times d_k}$, $W_i^V \in \mathbb{R}^{d \times d_v}$ and $W^O \in \mathbb{R}^{hd_v \times d}$, $d_k = d_v = d/h$.

IV. PROPOSED MODEL

This section begins with introducing the vanilla and masked encoder, followed by the masked temporal encoder and bidirectional encoder representation network for patient journeys to predict medical outcomes.

A. Masked Encoder Block

Because self-attention was originally designed for NLP tasks, it does not consider temporal relationships within inputs. Obviously, this is very important when modeling sequential medical events. Inspired by the work of Vaswani et al. [13] with a transformer and Shen et al. [16], [36] with masked multi-dimensional self-attention, we developed the masked encoder to capture the contextual and temporal relationships between visits. The structure of our masked encoder is shown in Fig. 2, which has two sub-layers.

The first is a multi-head self-attention mechanism, which is rewritten into a temporal-dependent format,

$$\text{Attention}(Q, K, V, M) = \text{softmax}(\alpha + M)V, \quad (16)$$

where α outlined in Eq.(7) and (12) work to capture the contextual relationship between the visits, and M to capture the temporal dependency between those visits. There are three types of temporal order information forward, backward and diagonal-disabled mask. There is no temporal order information between medical codes in a visit. In this context, self-attention usually requires disabling the attention of each code to itself [37], which is effectively the same as applying a diagonal-disabled mask such that,

$$M_{ij}^{diag} = \begin{cases} 0, & i \neq j \\ -\infty, & i = j \end{cases} \quad (17)$$

Similarly, masks can be used to encode temporal order information between visits as an attention output. Our approach also incorporates two masks, i.e., a forward mask M^{fw} and backward mask M^{bw} , i.e.,

$$M_{ij}^{fw} = \begin{cases} 0, & i < j \\ -\infty, & \text{otherwise} \end{cases} \quad (18)$$

$$M_{ij}^{bw} = \begin{cases} 0, & i > j \\ -\infty, & \text{otherwise} \end{cases} \quad (19)$$

The forward mask M^{fw} , only relates later visits j to earlier visits i , and vice versa with the backward one.

The second sub-layer is a simple, position-wise, fully connected feed-forward network. We employ a residual connection [38] around each of the two sub-layers, followed by layer normalization [39]. That is, the output of each layer is $\text{LayerNorm}(v + \text{sublayer}(v))$, where $\text{sublayer}(v)$ is the function implemented by the layer itself [13].

B. Bidirectional Temporal Encoder Network

We propose a patient journey embedding model, called **Bidirectional temporal encoder Network** (BiteNet), with MasEnc as its major components. The architecture of BiteNet is shown in Fig. 3. In BiteNet, the embedding layer is applied to the input medical codes of visits, and its output is processed by a stack of N MasEnc blocks with diagonal mask M^{diag} in Eq.(17) and the code-level attention layer to generate visit embeddings $[v_1, v_2, \dots, v_m]$ given in Eq.(10). The visit interval encodings are then added to the visit embeddings, followed by two parameter-untied stacks of N MasEnc blocks with forward mask M^{fw} in Eq.(18) and M^{bw} in Eq.(19), respectively. Their outputs are denoted by $u^{fw}, u^{bw} \in \mathbb{R}^{d \times m}$. The visit-level attention layers in Eq.(10) are applied to the outputs followed by the concatenation layer to generate output $u^{bi} \in \mathbb{R}^{2d}$. To complete the process, a feed-forward layer consisting of dense and softmax layers is employed to generate a categorical distribution over targeted categories.

a) *Interval Encoding*: Although MasEnc incorporates information on the order of visits in a patient journey, the relative time intervals between visits is an important factor in longitudinal studies. We include information on the visit intervals $p = [p_1, p_2, \dots, p_n]$ in the sequence. In particular, we add interval encodings to the visit embeddings of the sequence. There are two steps to perform interval encoding: 1) building a lookup table with dimensions of $m \times d$, where m is the days of the dataset span (e.g., $m = 11 \times 365$ over 11 years), and d is the dimension of code embedding; 2) updating the lookup table during training. The d -dimensional interval encoding is then added to the visit embedding with the same dimension.

Fig. 3. Bidirectional temporal encoder network (BiteNet) for patient journey embedding.

V. EXPERIMENTS

We conducted experiments based on the MIMIC-III dataset to compare the performance of our proposed method, BiteNet, with several state-of-the-art methods. We also evaluate the influence of the model components via ablation studies. Building on this, we visualize the importance of visits for patients and diagnoses in visits derived from the visit-level and code-level attention weightings, respectively. At last, we compare the performance between BiteNet and baseline models on unsupervised tasks. The source code is available in <https://github.com/Xueping/BiteNet>.

A. Data Description

1) *Dataset*: MIMIC-III [12] is an open-source, large-scale, de-identified dataset of ICU patients and their EHRs. The dataset consists of 46k+ ICU patients with 58k+ visits collected over 11 years. In this paper, we consider two sub-datasets: Dx and Dx&Tx, where Dx is a dataset which only includes diagnosis codes for each visit, and Dx&Tx is another dataset which includes diagnosis and procedure codes for each visit.

2) *Data Pre-processing*: We chose patients who made at least two visits. All infrequent diagnoses codes were removed and the threshold was empirically set to 5. In summary, we extract 7,499 patients, each with an average of 2.66 visits; the

average number of diagnoses and procedures in each visit are 13 and 4, respectively.

B. Experiment Setup

1) *Supervised Prediction Tasks*: The two tasks we selected to evaluate the performance of our proposed model are to predict re-admission [40] and future diagnosis [2].

- **Re-admission (Readm)** is a standard measure of the quality of care. We predicted unplanned admissions within 30 days following discharge from an indexed visit. A visit is considered a re-admission if its admission date is within 30 days of discharge from an eligible indexed visit [2].
- **Diagnoses** reflect the model’s understanding of a patient’s problems. In the experiments, we aim to predict diagnosis categories instead of the real diagnosis codes, which are the nodes in the second hierarchy of the ICD9 codes¹.

2) *Unsupervised Tasks*: Two tasks of **clustering** and **nearest neighbor search (NNS)** [41] in unsupervised learning were conducted to evaluate the quality of the medical code embedding results. We selected the ground truth by using a well-organized ontology, Clinical Classifications Software (CCS)². CCS provides a way to classify diagnoses and procedures into a limited number of categories by aggregating individual ICD9 codes into broad diagnosis and procedure groups to facilitate statistical analysis and reporting³. CCS aggregates ICD9 diagnosis codes into 285 mutually exclusive categories. We obtained 265 categories and 61,630 near neighbor pairs for clustering and the nearest neighbor search, respectively.

3) *Baseline Methods*: We compare the performance of our proposed model against several baseline models:

- **RNN** and **BRNN**, we directly embed visit information into the vector representation v_t by summation of embedded medical codes in the visit, and then feed this embedding to the GRU. The hidden state h_t produced by the GRU is used to predict the $(t+1)$ -th visit information.
- **RETAIN** [8]: which learns the medical concept embeddings and performs heart failure prediction via the reversed RNN with the attention mechanism.
- **DeepR** [30]: which is a multilayered architecture based on convolutional neural networks (CNNs) that learn to extract features from medical records and predict future risk.
- **Dipole** [7], which uses bidirectional recurrent neural networks and three attention mechanisms (location-based, general, concatenation-based) to predict patient visit information. We chose concatenation-based Dipole as a baseline method.
- **SanD** [11]: which employs a masked, self-attention mechanism, and uses positional encoding and dense interpolation strategies to incorporate temporal order to generate a sequence-level prediction.

4) *Evaluation metrics*: The two evaluation metrics used are:

- **PR-AUC**: (Area under Precision-Recall Curve), to evaluate the likelihood of re-admission. PR-AUC is considered to be a better measure for imbalanced data like ours.
- **precision@k**: which is defined as the correct medical codes in top k divided by $\min(k, |y_t|)$, where $|y_t|$ is the number of category labels in the $(t+1)$ -th visit [25]. We report the average values of precision@k in the experiments. We vary k from 5 to 30. The greater the value, the better the performance.
- **Accuracy@k**: to evaluate the performance of the nearest neighbor search, which is defined as the number of the nearest neighbors in top k divided by k . We report the average values of Accuracy@k in the experiments. We vary k from 1 to 10. The greater the value, the better the performance.
- **NMI**: normalized mutual information for clustering performance.

5) *Implementation*: We implement all the approaches with Tensorflow 2.0. For training models, we use RMSprop with a minibatch of 32 patients and 10 epochs. The drop-out rate is 0.1 for all approaches. The data split ratio is 0.8:0.1:0.1 for training, validation and testing sets. In order to fairly compare the performance, we set the same embedding dimension $d = 128$ for all the baselines and the proposed model.

C. Results of Prediction Tasks

1) *Prediction Performance Comparison*: Table II reports the results of the two prediction tasks on the two datasets - future re-admissions and diagnoses. The results show that BiteNet outperforms all the baselines. This demonstrates that the superiority of our framework results from the explicit consideration of the inherent hierarchy of EHRs and the contextual and temporal dependencies which are incorporated into the representations. In addition, we note that the performance obtained by the models remains approximately the same (RNN), or even drops by up to 0.43% (DeepR), after adding the procedure to the training data over Precision@5. The underlying reason might be that the future diagnosis prediction is less sensitive to the procedures, with the result that the relationships among them cannot be well captured when predicting few diagnoses. Still, by using attention mechanisms, Dipole, SanD and our BiteNet model, we achieve a marginal improvement when comparing the performance of using Dx&Tx for training to that of Dx. This implies that attention could play an important role in the learned representations of the procedures. For the performance of future re-admission, we note that all models achieve an improvement when adding procedure information in the training stage, indicating that the additional procedures benefit for the future re-admission prediction.

2) *Robustness for the lengthy sequence of visits*: We conducted a set of experiments to evaluate the robustness of BiteNet by varying the length of sequential visits of a patient from 6 to 16, considered to be long patient journeys in the

¹<http://www.icd9data.com>

²<https://www.hcup-us.ahrq.gov>

³ <https://www.hcup-us.ahrq.gov/toolsoftware/ccs/CCSUsersGuide.pdf>

TABLE II
PREDICTION PERFORMANCE COMPARISON OF FUTURE RE-ADMISSION AND DIAGNOSES (DX IS FOR DIAGNOSIS, AND TX IS FOR PROCEDURE).

Data	Model	Readm (PR-AUC)	Diagnosis Precision@k					
			k=5	k=10	k=15	k=20	k=25	k=30
Dx	RNN	0.3021±0.0051	0.6330±0.0055	0.5874±0.0035	0.6309±0.0047	0.6977±0.0064	0.7601±0.0053	0.8068±0.0011
	BRNN	0.3119±0.0078	0.6362±0.0051	0.5925±0.0031	0.6365±0.0026	0.7014±0.0035	0.7596±0.0049	0.8128±0.0026
	RETAIN	0.3014±0.0127	0.6498±0.0075	0.5948±0.0055	0.6368±0.0054	0.6999±0.0060	0.7627±0.0039	0.8102±0.0042
	DeepPr	0.2999±0.0190	0.6434±0.0030	0.5865±0.0025	0.6257±0.0033	0.6981±0.0018	0.7604±0.0044	0.8113±0.0047
	Dipole	0.2841±0.0128	0.6484±0.0062	0.5997±0.0042	0.6339±0.0041	0.7034±0.0026	0.7620±0.0024	0.8121±0.0012
	SAnD	0.2979±0.0263	0.6179±0.0147	0.5709±0.0136	0.6100±0.0126	0.6805±0.0123	0.7464±0.0119	0.7959±0.0109
	BiteNet	0.3266±0.0047	0.6615±0.0024	0.6019±0.0056	0.6432±0.0031	0.7104±0.0031	0.7757±0.0046	0.8245±0.0053
Dx&Tx	RNN	0.3216±0.0047	0.6317±0.0055	0.5857±0.0029	0.6310±0.0037	0.6973±0.0015	0.7616±0.0021	0.8093±0.0019
	BRNN	0.3270±0.0065	0.6402±0.0064	0.5961±0.0025	0.6386±0.0032	0.7088±0.0036	0.7662±0.0018	0.8138±0.0018
	RETAIN	0.3161±0.0107	0.6552±0.0058	0.6048±0.0047	0.6429±0.0060	0.7085±0.0042	0.7675±0.0042	0.8166±0.0039
	DeepPr	0.3142±0.0150	0.6391±0.0038	0.5947±0.0040	0.6330±0.0042	0.7004±0.0043	0.7633±0.0041	0.8125±0.0031
	Dipole	0.2899±0.0093	0.6515±0.0078	0.6097±0.0050	0.6415±0.0044	0.7121±0.0064	0.7691±0.0065	0.8149±0.0046
	SAnD	0.2996±0.0135	0.6242±0.0104	0.5774±0.0106	0.6233±0.0092	0.6878±0.0089	0.7510±0.0107	0.8004±0.0122
	BiteNet	0.3357±0.0045	0.6705±0.0045	0.6117±0.0033	0.6511±0.0035	0.7187±0.0046	0.7799±0.0067	0.8289±0.0068

Fig. 4. Robustness comparison regarding Precision @20 on two datasets. Length of the visit sequence varies from 6 to 16.

Fig. 5. Diagnoses Precision@{5,10,15,20,25,30} on two datasets. The MasEnc layer number N varies from 1 to 8.

medical domain. Fig. 4 shows the performance of future diagnoses prediction on Precision@20. Overall, the model BiteNet outperforms the other models as the length of a patient's sequential visits increases diagnoses prediction over the dataset with only diagnoses information. Our BiteNet

model achieves the best performance when the lengths of sequential visits are 6, 8 and 12 over Dx&Tx. In particular, the performances of Dipole and BRNN are comparative with our BiteNet and follow a similar trend with an increase in the length of sequential visits. From these results, a conclusion

Fig. 6. Diagnoses Precision@{5,10,15,20,25,30} on two datasets when $N = 2$. The number of heads in multi-head attention varies from 4 to 32.

can be drawn that the temporal order masks in the MasEnc module and interval encoding in our framework play a vital role in capturing the lengthy temporal dependencies between patients' sequential visits.

3) *Effectiveness of varying stack N of MasEnc layers*: In this set of experiments, we evaluated the performance of our proposed model BiteNet by varying the stack N of MasEnc layers from 1 to 8, where the number of heads in multi-head attention is set to 4. Fig. 5 shows the results. Overall, BiteNet achieved the best performance of future diagnoses prediction when the layer number of MasEnc block is 2. After this, the performance was slightly declining as the number of layers increased. However, the precision of future diagnoses prediction dropped to a minimum when $N = 7$ and then slightly increased when $N = 8$ over dataset Dx. Overall, we find the performance of Precision@30 outperformed the others, and Precision@10 was the worth over both Dx and Dx&Tx.

4) *Effectiveness of varying heads h of multi-head attention*: In this set of experiments, we varied the number of heads h in multi-head attention from 4 to 32 to evaluate the effectiveness of the proposed BiteNet, where the stack N of MasEnc layers is set to 2. From Fig. 6, we can observe that the performance of future diagnoses prediction drops when the number of heads becomes larger. Particularly, the prediction precision of BiteNet remains stable or drops slightly over both Dx and Dx&Tx. It demonstrates that our proposed model, BiteNet, is less sensitive to the number of heads in multi-head attention.

5) *Ablation Study*: We performed a detailed ablation study to examine the contributions of each of the model's components to the prediction tasks. There are three components: (Attention) the two attention pooling layers to learn the visits from the embedded medical codes and learn the patient journey from the embedded visits; (DireMask) the direction mask in MasEnc module, and (Interval) the interval encodings to be added to the learned visit embeddings.

- **Attention**: replace each of the two attention pooling layers with a simple summation layer;
- **DireMask**: remove the diagonal-disabled, forward and

backward directional mask in the MasEnc module;

- **Interval**: remove the interval encodings module;
- **BiteNet**: our model as proposed in the paper.

From Table III, we see that the full complement of BiteNet achieves superior accuracy over the ablated models. Specifically, we note that Interval and Attention contribute the highest accuracy to re-admission prediction over Dx and Dx&Tx, respectively. It gives us confidence in using the Interval encoding and Attention pooling to learn the patient journey representations without sufficient data. Significantly, it is clear that the MasEnc component provides valuable information with the learned embeddings of the patient journey for the performance of diagnoses prediction over Dx&Tx. Specifically, BiteNet predicted re-admissions with 3.45% more accuracy on Dx and with 3.40% over Dx&Tx. Similarly, BiteNet predicted diagnoses precision@5 with 1.46% more accurately than Dx and diagnoses precision@30 with 1.41% more accuracy than Dx&Tx.

6) *Visualization and Explainability*: One aspect of this method is that it hierarchically compresses medical codes into visits and visits into patient journeys. At each level of aggregation, the model decides the importance of the lower-level entries on the upper-level entry, which makes the model explainable. To showcase this feature, we visualize two patient journeys. These examples come from the re-admission prediction task with the MIMIC dataset. From the importance distribution of the patient visits, we analyze the most essential visits to future re-admission. For example, Visit 1 in Fig. 7 was the most influential factor in Patient 1's re-admission. After closely examining Visit 1, we found a vital insight into Diagnosis 3 (ICD 585.6), i.e., end-stage renal disease, which would obviously cause frequent and repeated re-admissions to the hospital. in Fig. 8, Patient 2 had 7 visits. Visit 4 contributed most to the re-admissions. Again, the diagnoses reveal the cause: long-term use of insulin (ICD V58.6.7) and pure hypercholesterolemia (ICD 272.0).

7) *Results of Unsupervised Tasks*: We use the clustering and nearest neighbor search tasks to evaluate the embedding results on two datasets: Dx and Dx&Tx. We choose K-Means

TABLE III
ABLATION PERFORMANCE COMPARISON.

Data	Ablation	Readm (PR-AUC)	Diagnosis Precision@k					
			k=5	k=10	k=15	k=20	k=25	k=30
Dx	Attention	0.3015±0.0068	0.6506±0.0050	0.6014±0.0028	0.6440±0.0032	0.7125±0.0022	0.7712±0.0020	0.8174±0.0023
	DireMask	0.3026±0.0132	0.6528±0.0051	0.6047±0.0031	0.6460±0.0029	0.7128±0.0035	0.7711±0.0030	0.8191±0.0034
	Interval	0.2921±0.0111	0.6469±0.0041	0.5947±0.0050	0.6361±0.0029	0.7054±0.0025	0.7667±0.0028	0.8139±0.0020
	BiteNet	0.3266±0.0047	0.6615±0.0124	0.6019±0.0056	0.6432±0.0031	0.7104±0.0031	0.7757±0.0046	0.8245±0.0053
Dx&Tx	Attention	0.3017±0.0070	0.6453±0.0046	0.5984±0.0030	0.6424±0.0027	0.7092±0.0036	0.7701±0.0020	0.8176±0.0009
	DireMask	0.3046±0.0124	0.6487±0.0050	0.6015±0.0021	0.6437±0.0023	0.7117±0.0009	0.7728±0.0013	0.8193±0.0008
	Interval	0.3109±0.0109	0.6420±0.0018	0.5945±0.0027	0.6352±0.0034	0.7060±0.0021	0.7664±0.0013	0.8148±0.0016
	BiteNet	0.3357±0.0045	0.6705±0.0045	0.6117±0.0033	0.6511±0.0035	0.7187±0.0046	0.7799±0.0067	0.8289±0.0068

TABLE IV
PREDICTION PERFORMANCE COMPARISON OF FUTURE RE-ADMISSION AND DIAGNOSES (DX IS FOR DIAGNOSIS, AND TX IS FOR PROCEDURE).

Data	Model	Clustering (NMI)	NNS Accuracy@k		
			k=1	k=5	k=10
Dx	RNN	0.4824±0.0054	0.1462±0.0133	0.1236±0.0043	0.1154±0.0030
	BRNN	0.4842±0.0035	0.1524±0.0078	0.1204±0.0053	0.1084±0.0062
	RETAIN	0.4844±0.0084	0.1556±0.0154	0.1254±0.0041	0.1100±0.0030
	DeepPr	0.4571±0.0059	0.1120±0.0112	0.0936±0.0040	0.0914±0.0041
	Dipole	0.4822±0.0046	0.1504±0.0117	0.1156±0.0038	0.1084±0.0030
	SAnD	0.4717±0.0044	0.1243±0.0096	0.1148±0.0071	0.0984±0.0041
	BiteNet	0.4856±0.0035	0.1584±0.0075	0.1284±0.0035	0.1154±0.0028
Dx&Tx	RNN	0.4838±0.0036	0.1505±0.0118	0.1236±0.0032	0.1148±0.0028
	BRNN	0.4851±0.0028	0.1643±0.0069	0.1240±0.0036	0.1150±0.0048
	RETAIN	0.4867±0.0078	0.1744±0.0153	0.1312±0.0032	0.1142±0.0036
	DeepPr	0.4658±0.0051	0.1283±0.0100	0.1116±0.0028	0.1010±0.0006
	Dipole	0.4858±0.0056	0.1744±0.0110	0.1343±0.0040	0.1188±0.0042
	SAnD	0.4722±0.0040	0.1264±0.0095	0.1084±0.0045	0.1014±0.0047
	BiteNet	0.4873±0.0034	0.1758±0.0061	0.1492±0.0030	0.1244±0.0024

Fig. 7. Importance of visits for patient 1, and importance of diagnoses in visit 1.

Fig. 8. Importance of visits for patient 2, and importance of diagnoses in visit 4.

as the clustering algorithm, and NMI to evaluate the learned representations for the diagnosis codes. The embedding of diagnosis codes is derived from the proposed model BiteNet on the task of future diagnoses prediction.

NMI for clustering and Accuracy@k ($k = 1, 5, 10$) for NNS

are reported in Table IV, where we highlight the best results. Overall, the model BiteNet outperforms the other baseline models on the two unsupervised tasks over two datasets. From the table, we find that the RNN-based models (RNN, BRNN, RETAIN and Dipole) obtained better performance in medical code embedding compared to the CNN-based models (DeepPr and SAnD). This demonstrates that temporal dependency on EHR data is significantly hidden information when embedding medical codes and RNN-based models can capture the temporal relationship. However, our proposed model, BiteNet, outperforms the baseline models including RNN-based, CNN-based, and attention integrated models, because BiteNet captures the contextual information and long-term dependencies between patients' visits, and considers the hierarchical structure of EHR data as well. The superior performance of BiteNet over the other models can also be explained by the introduction of the temporal self-attention model and the incorporation of the contextual information and the temporal interval from the data, which create a better learner of the medical code embeddings.

VI. CONCLUSION

In this paper, we proposed a novel prediction model called **BiteNet**. The model framework comprises a MasEnc module that captures the contextual information and the temporal relationships between the visits in a patient's healthcare journey and attention pooling that construct the hierarchical structure of three-levelled EHR data. The output is a representation of a patient journey that, once learned by the model, can be used to predict medical outcomes with an end-to-end sole self-attention network. We evaluated BiteNet's performance of the model against several baseline methods with supervised and unsupervised tasks, and conducted an ablation study to examine the contributions of each component. The results show that BiteNet produces more accurate predictions than baseline methods.

VII. ACKNOWLEDGMENTS

This work was supported in part by the Australian Research Council (ARC) under Grant LP160100630, LP180100654 and DE190100626.

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] A. Rajkomar, E. Oren, K. Chen, A. M. Dai, N. Hajaj, M. Hardt, P. J. Liu, X. Liu, J. Marcus, M. Sun *et al.*, "Scalable and accurate deep learning with electronic health records," *npj Digital Medicine*, vol. 1, no. 1, p. 18, 2018.
- [3] Z. Qiao, S. Zhao, C. Xiao, X. Li, Y. Qin, and F. Wang, "Pairwise-ranking based collaborative recurrent neural networks for clinical event prediction," in *IJCAI*, 2018, pp. 3520–3526.
- [4] 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.
- [5] 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.
- [6] E. Choi, C. Xiao, W. Stewart, and J. Sun, "Mime: Multilevel medical embedding of electronic health records for predictive healthcare," in *NeurIPS*, 2018, pp. 4552–4562.
- [7] 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, 2017, pp. 1903–1911.
- [8] 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.
- [9] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [10] 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.
- [11] H. Song, D. Rajan, J. J. Thiagarajan, and A. Spanias, "Attend and diagnose: Clinical time series analysis using attention models," in *AAAI*, 2018.
- [12] 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.
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017, pp. 5998–6008.
- [14] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv:1409.0473*, 2014.
- [15] K. Niu, J. Guo, Y. Pan, X. Gao, X. Peng, N. Li, and H. Li, "Multi-channel deep attention neural networks for the classification of autism spectrum disorder using neuroimaging and personal characteristic data," *Complexity*, vol. 2020, 2020.
- [16] T. Shen, T. Zhou, G. Long, J. Jiang, S. Pan, and C. Zhang, "Disan: Directional self-attention network for rnn/cnn-free language understanding," in *AAAI*, 2018.
- [17] Y. Kim, "Convolutional neural networks for sentence classification," in *EMNLP*, Oct. 2014, pp. 1746–1751.
- [18] K. Jha, Y. Wang, G. Xun, and A. Zhang, "Interpretable word embeddings for medical domain," in *ICDM*. IEEE, 2018, pp. 1061–1066.
- [19] 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.
- [20] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv:1301.3781*, 2013.
- [21] 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.
- [22] J. A. Minarro-Giménez, O. Marin-Alonso, and M. Samwald, "Exploring the application of deep learning techniques on medical text corpora," *Stud Health Technol Inform*, vol. 205, pp. 584–588, 2014.
- [23] L. De Vine, G. Zuccon, B. Koopman, L. Sitbon, and P. Bruza, "Medical semantic similarity with a neural language model," in *CIKM*. ACM, 2014, pp. 1819–1822.
- [24] Y. Choi, C. Y.-I. Chiu, and D. Sontag, "Learning low-dimensional representations of medical concepts," *AMIA Summits on Translational Science Proceedings*, vol. 2016, p. 41, 2016.
- [25] 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.
- [26] X. Peng, G. Long, T. Shen, S. Wang, J. Jiang, and M. Blumenstein, "Temporal self-attention network for medical concept embedding," in *ICDM*. IEEE, 2019, pp. 498–507.
- [27] X. Peng, G. Long, T. Shen, S. Wang, and J. Jiang, "Self-attention enhanced patient journey understanding in healthcare system," *arXiv preprint arXiv:2006.10516*, 2020.
- [28] 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.
- [29] T. Ma, C. Xiao, and F. Wang, "Health-atm: A deep architecture for multifaceted patient health record representation and risk prediction," in *SDM*. SIAM, 2018, pp. 261–269.
- [30] P. Nguyen, T. Tran, N. Wickramasinghe, and S. Venkatesh, "Deepr: A convolutional net for medical records," 2016.
- [31] L. Shang, Z. Lu, and H. Li, "Neural responding machine for short-text conversation," *arXiv:1503.02364*, 2015.
- [32] S. Sukhbaatar, J. Weston, R. Fergus *et al.*, "End-to-end memory networks," in *NeurIPS*, 2015, pp. 2440–2448.
- [33] A. M. Rush, S. Chopra, and J. Weston, "A neural attention model for abstractive sentence summarization," *arXiv:1509.00685*, 2015.
- [34] 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.
- [35] Y. Liu, C. Sun, L. Lin, and X. Wang, "Learning natural language inference using bidirectional lstm model and inner-attention," *arXiv:1605.09090*, 2016.
- [36] T. Shen, T. Zhou, G. Long, J. Jiang, and C. Zhang, "Bi-directional block self-attention for fast and memory-efficient sequence modeling," *arXiv:1804.00857*, 2018.
- [37] M. Hu, Y. Peng, Z. Huang, X. Qiu, F. Wei, and M. Zhou, "Reinforced mnemonic reader for machine reading comprehension," *arXiv:1705.02798*, 2017.
- [38] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [39] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [40] C. Xiao, T. Ma, A. B. Dieng, D. M. Blei, and F. Wang, "Readmission prediction via deep contextual embedding of clinical concepts," *PLoS one*, vol. 13, no. 4, p. e0195024, 2018.
- [41] X. Cai, J. Gao, K. Y. Ngiam, B. C. Ooi, Y. Zhang, and X. Yuan, "Medical concept embedding with time-aware attention," in *IJCAI*, 2018, pp. 3984–3990.