Variationally Regularized Graph-based Representation Learning for Electronic Health Records

Weicheng Zhu
jackzhu@nyu.edu
Center for Data Science
New York University
New York, NY, USA

Narges Razavian
narges.razavian@nyulangone.org
Department of Population Health
Department of Radiology
NYU School of Medicine
New York, NY, USA

ABSTRACT

Electronic Health Records (EHR) are high-dimensional data with implicit connections among thousands of medical concepts. These connections, for instance, the co-occurrence of diseases and lab-disease correlations can be informative when only a subset of these variables is documented by the clinician. A feasible approach to improving the representation learning of EHR data is to associate relevant medical concepts and utilize these connections. Existing medical ontologies can be the reference for EHR structures, but they place numerous constraints on the data source. Recent progress on graph neural networks (GNN) enables end-to-end learning of topological structures for non-grid or non-sequential data. However, there are problems to be addressed on how to learn the medical graph adaptively and how to understand the effect of medical graph on representation learning. In this paper, we propose a variationally regularized encoder-decoder graph network that achieves more robustness in graph structure learning by regularizing node representations. Our model outperforms the existing graph and non-graph based methods in various EHR predictive tasks based on both public data and real-world clinical data. Besides the improvements in empirical experiment performances, we provide an interpretation of the effect of variational regularization compared to standard graph neural network, using singular value analysis.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Computing methodologies** → **Learning latent representations**; **Neural networks**.

KEYWORDS

Graph neural networks, Electronic health records, Alzheimer's disease, Regularization, Singular value decomposition

ACM Reference Format:

Weicheng Zhu and Narges Razavian. 2021. Variationally Regularized Graph-based Representation Learning for Electronic Health Records. In *ACM*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8359-2/21/04.

<https://doi.org/10.1145/3450439.3451855>

Conference on Health, Inference, and Learning (ACM CHIL '21), April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3450439.3451855>

1 INTRODUCTION

Electronic Health Records (EHR) are rich sources of information, useful in various predictive tasks in medical application, including mortality prediction, outcomes prediction and phenotyping. The accessibility of the EHR data makes it a feasible resource for scaling screening to large populations. Especially for chronic diseases like Alzheimer's Diseases, early identification before the onset of clinical symptoms can improve the enrollment for clinical trials, and improve effectiveness of the treatments. Previous studies have explored various deep learning methodologies on the EHR [Cheng et al. 2016; Miotto et al. 2016; Tran et al. 2016]. Learning representations of medical concepts emerges as an important branch [Choi et al. 2016a, 2018, 2016c], and recent research demonstrates the significance of graph structures among medical concepts [Choi et al. 2016b, 2019]. EHR are inherently sparse and structured data with high probability of missing values. Some diseases may be recorded as diagnosis codes, while other existing conditions that are not discussed in the clinical encounter may not be documented. Graph neural network (GNN) has been considered an effective way to generalize Convolutional Neural Networks (CNN) in extracting signals from non-grid structured data [Bruna et al. 2013; Henaff et al. 2015]. As CNN can focus on the localized features of images or sequences, GNN also enable the model to highlight the significant features and infer the missing features within the topological neighbourhood. Therefore, GNN can be a strong tool for multiple machine learning tasks on EHR, including patient representation learning, medical graph learning and disease prediction.

Our work has the following main contribution on graph-based representation learning of EHR:

- We design a novel graph-based model to generalize the ability of learning implicit medical concept structures to a wide range of data sources, including short-term ICU data and long-term outpatient clinical data.
- We introduce variational regularization for node representation learning, addressing the insufficiency of self-attention in graph-based models, and difficulties of manually constructing knowledge graph from real-world noisy data sources. The novelty of our work is to enhance the learning of attention weights in GNN via regularization on node representations. With this design, our method outperforms previous

graph representation learning method in health predictive tasks based on a clinical EHR and two public EHR datasets.

- We provide interpretation on the effect of variational regularization in graph neural networks using singular value analysis, and bridge the connection between singular values and representation clustering.

2 RELATED WORK

Among the recent deep learning research on EHR, there are two dominating approaches - extracting temporal signals from time series EHR, and learning embeddings of medical concepts without directly modeling time. [Shickel et al. 2017]. In the first approach, researchers acquire temporal features from sequential biomarkers or encounter data via representation learning methods including recurrent neural networks [Che et al. 2016; Choi et al. 2015; Lipton et al. 2015], convolution blocks [Cheng et al. 2016; Razavian et al. 2016] and attention mechanism [Li et al. 2019; Song et al. 2017]. The other approach is to train neural networks that express medical concepts with high-dimensional embeddings as learned representations. Med2Vec [Choi et al. 2015] learns to represent EHR codes as embeddings, following the idea of skip-gram [Mikolov et al. 2013] in NLP tasks. These previous studies skip an important property of EHR that the diagnosis, labs and procedures are inherently associated with each other. For instance, some diseases co-occur or induce other diseases, and some lab exams indicate certain diseases.

One approach to incorporating the structure of medical concepts into representation learning is to build a medical graph that connects related codes in the EHR. Some previous research models EHR with graph structure: GRAM [Choi et al. 2016b] leverages medical ontologies to learn representations of medical concepts. These methods improve representation learning in predictive modeling by incorporating additional structural information or external knowledge. MiME [Choi et al. 2018] learns hierarchical embeddings of a subset of variables (visits, diagnosis codes, and medications) by building relationships between different levels of hierarchy. These graph-based methods mainly focus on parent-child relationship and rely on the assumption that the EHR data follows hand-designed protocols that may not accurately reflect real world data.

Compared to previous methods, the Graph Convolutional Networks (GCN) [Bruna et al. 2013; Henaff et al. 2015] are more flexible in learning graph representations. GCN generalizes translation-invariant convolution filters in standard convolutional neural networks (CNN) to a non-Euclidean localized filter [Henaff et al. 2015], that can be applied to various non-grid data. GCN can be applied to learning representations of node features and graph structure through semi-supervised learning for node classification [Kipf and Welling 2016a]. This work provides an approach for generating labels for unknown nodes given graph structures and node features. Self-attention, comparable to CNN in encoding features from spatial or sequential data [Cordonnier et al. 2020], requires less computation time and out-performs CNN in language and vision tasks [Ramachandran et al. 2019; Vaswani et al. 2017]. Similar to replacing convolutional block with self-attention, Graph Attention Network (GAT) [Veličković et al. 2018] attends each node in the graph on its neighbouring nodes and itself to learn localized features instead of using GCN spectral filters. In this architecture, GAT can assign

different importance to edges, which increases the model capacity and interpretability, and learns graph structures themselves via attention parameters.

Inspired by self-attention mechanism and GAT, we propose an encoder-decoder GNN by taking each observed EHR code as a node, initially imposing a fully connected graph on them, and implicitly learning their graph structure via self-attention mechanism. A recent work, Graph Convolution Transformer (GCT) [Choi et al. 2019], introduces a visit embedding for each patient medical encounter, and combines the visit embedding with the other medical concept embeddings through the Transformer [Vaswani et al. 2017]. Furthermore, the authors address a problem mentioned in their paper that the Transformer cannot effectively learn the attention parameters from scratch and often leads to uniformly distributed attention weights among medical concepts. GCT in [Choi et al. 2019] solves the problem by leveraging a pre-defined graph as the guidance of regularization. The graph is constructed by connecting different groups of medical concepts (e.g. diagnosis, labs and procedures) to emulate physicians' decision process.

However, the real-world EHR has significant missing data, and the hierarchy among different features cannot be strictly defined. Lab values that correspond to some diseases may not be included in patient's EHR data and vice versa. To design a generalizable model that can work with any groups of variables (diagnosis, labs, procedures, demographics), we do not prune internal connections among diseases that do not match a pre-defined graph, but rather, introduce variational regularization in our encoder-decoder GNN to address the challenges of structure learning without pre-defined graphs. Variational inference has been previously used to improve the representation learning of graph autoencoder [Hasanzadeh et al. 2019; Kipf and Welling 2016b] and link prediction with Bernoulli link inference [Hajiramezanali et al. 2019; Tiao et al. 2019]. Previous studies have reported that the regularization effect of KL-term in variational autoencoder [Prokhorov et al. 2019] and the strength of regularization is adjustable by scaling on KL-term [Higgins et al. 2017; Razavi et al. 2019]. Unlike the autoencoder framework which combines reconstruction loss with the KL-term, here we only use the KL-term to regularize node representations so that the attention weight learning can be enhanced in supervised tasks. We will discuss the benefits of this regularization in section 5.

3 METHODS

3.1 Encoder-decoder Graph Neural Networks

We embed EHR codes $\mathcal{V} = \{1, 2, \dots, N\}$ with high-dimensional vectors $\{h_i\}_{i \in \mathcal{V}}, h_i \in \mathbb{R}^d$. Compared to GAT [Veličković et al. 2018] which takes advantage of external features of nodes, we learn the representations of each medical concept. For each patient X , we denote their observed codes $\mathcal{V}_{\text{obs}} = \{x_1, x_2, \dots, x_n\}$ and initially fully connect them, and consequently learn the structure and updated representations over L additional graph layers. This part of our architecture is termed the encoder graph. We then denote additional nodes $\mathcal{V}_{\text{out}} = \{y_1, y_2, \dots, y_m\}$ for prediction task and fully connect them to output nodes of encoder graph. This sub-network is termed the decoder graph. As outlined in the top architecture of Figure 1, the medical embeddings are processed by the encoder to represent the graph, and then the decoder provides inferences based on the

(a) The full architecture of our encoder-decoder model (*Enc-dec*)

(b) The illustration of variational regularization (VGNN) on the graph architecture (blue highlighted block in (a))

Figure 1: The architecture of our encoder-decoder model (top architecture), as well as variational regularization (bottom architecture). For each patient, the observed features $i \in \mathcal{V}_{\text{obs}} \{x_1, x_2, x_3, x_4\}$ in the illustration) correspond to the EHR variables that were documented/observed in the patient records. These features are encoded as embeddings $\{h_i\}_{i \in \mathcal{V}_{\text{obs}}}$. Note that each patient has a different $\mathcal{V}_{\text{obs}}$. These observed nodes are fully connected in the first layer. In subsequent layers, the weights on the edges $\{A_{ij}\}_{i,j \in \mathcal{V}_{\text{obs}}}$, and the graph representation $\{h_i^{(l)}\}_{i \in \mathcal{V}_{\text{obs}}}$, are computed by multi-head attentions. The "encoder graph" denotes the observed nodes and the parametrized attention-based connections up to layer L . In the decoder graph, we add new nodes corresponding to the target outcomes to the encoder graph's output. Inferences on the outcomes $\{\hat{y}_c\}_{y_c \in \mathcal{V}_{\text{out}}}$ are derived from multi-head attentions on node representation and a linear feed-forward layer. In variationally regularized model, a latent layer $\{z_i\}_{i \in \mathcal{V}_{\text{obs}}}$ is placed between the encoder and decoder graph. Latent variables z_i 's are sampled from $\mathcal{N}(\mu_i, \exp(\sigma_i))$, where μ_i and σ_i are computed from $h_i^{(L)}$ by two separate feed-forward networks. The latent variables are regularized by $\mathcal{L}_{\text{div}}$, approximating the distributions of z_i 's to some priors $\mathcal{N}(\mu, \Sigma)$ (we use standard Gaussian here).

graph representations. In each graph layer, the representations are updated by graph propagation.

$$H^{(l+1)} = \text{FFN} \left[A^{(l)} \left(H^{(l)} W^{(l)} + b^{(l)} \right) \right] \quad (1)$$

where $W^{(l)} \in \mathbb{R}^{d \times d}$ and $b^{(l)} \in \mathbb{R}^d$ form a linear layer; $H^{(l)}$ is the matrix taking all the representations $h_i^{(l)}$ of observed nodes as row vectors; $A^{(l)}$ is the adjacency matrix at the l^{th} layer. The size of $H^{(l)}$ and $A^{(l)}$ varies with the sample and location of the layer. The nodes in the graph are $\mathcal{V}_{\text{obs}}$ in the encoder and $\mathcal{V}_{\text{obs}} \cup \mathcal{V}_{\text{out}}$

in the decoder. Hence, $A^{(l)} \in \mathbb{R}^{n \times n}$, $H^{(l)} \in \mathbb{R}^{d \times n}$ in the encoder, and $A^{(l)} \in \mathbb{R}^{(n+m) \times (n+m)}$, $H^{(l)} \in \mathbb{R}^{d \times (n+m)}$ in the decoder. FFNs, referring to feed-forward networks are the multilayer perceptron composed of linear layers, ReLU activations, dropout [Srivastava et al. 2014] and layer normalization [Ba et al. 2016]. In any graph layer l , the elements of adjacency matrices A_{ij} on the edge connecting node i to node j are computed using attention mechanism.

$$A_{ij} = \text{softmax}(e_{ij}) = \frac{\exp(e_{ij})}{\sum_{p \in N_i} \exp(e_{ip})} \quad (2)$$

where $e_{ip} \in \mathbb{R}$ are the attention coefficients for node i over its neighbourhood N_i . There are multiple ways of computing attention coefficients with interactions between two input vectors [Bahdanau et al. 2014; Luong et al. 2015; Vaswani et al. 2017]. The selection of attention style is discussed in Section 5. In this study, we use multi-head attention, as follows: We first concatenate two input vectors and apply a linear layer $a \in \mathbb{R}^{2d \times 1}$. The attention coefficients are computed by the following equation:

$$e_{ij} = \text{LeakyReLU} \left(a^T [Wh_i \parallel Wh_j] \right) / \sqrt{d_h} \quad (3)$$

W is a linear layer, and d_h is the dimensionality of h_i 's.

Multi-head attention [Vaswani et al. 2017] allows the model to jointly attend to information from different representation subspaces at different positions like multiple convolution filters in one CNN layer. To build our K -head attention, we alter the output size of linear layer $W^{(l)}$ and $b^{(l)}$ in equation (3) to dK , and then attention heads can be computed in parallel. Also, the outputs of multihead attention are aggregated by concatenation, so the input size of the feedforward networks is adjusted to dK as well to fit the concatenated representations.

As the majority of predictive tasks in EHR have imbalance labels (more negative than positive), we use the weighted binary cross-entropy loss to train this model:

$$\mathcal{L}_{\text{bce}} = - \sum_{y_c \in \mathcal{V}_{\text{out}}} w_c \cdot Y_c \cdot \log[\sigma(\hat{y}_c)] + (1 - Y_c) \cdot \log[1 - \sigma(\hat{y}_c)] \quad (4)$$

where σ is the sigmoid function; Y_c is the ground truth of y_c ; $\hat{y}_c \in \mathbb{R}$ is the output of the last feed-forward layer in the decoder; w_c is the negative-to-positive ratio, putting more weight on the minority. This loss function describes the loss of all the outcomes for one sample, and the loss of mini-batch is the mean of losses over the batch.

3.2 Variationally regularized Encoder-Decoder Model

In our experiments of the encoder-decoder graph networks, we observe that node representations after the encoder layer often collapse to a tight clustered and lack implicit structures, which leads to uniformly-distributed attention weights and prevents graph layers from learning meaningful edges among medical concepts. The uniformity of the attention weights is also observed by Choi et al. [Choi et al. 2019], and solved by regularizing the links with a pre-defined knowledge graph. In this study, we trace the problem to the distribution of embeddings and introduce variational regularization to encourage the node representations to be centered around the

origin with moderate distances, which as we will show, leads to learning more expressive connections.

Inspired by VGAE [Kipf and Welling 2016b] which improves link inference by assuming a Gaussian prior on the node representations, we add a latent layer between the encoder and decoder to regularize the graph representation. Let $Z = \{z_i\}_{i \in \mathcal{V}_{\text{obs}}}$, $z_i \in \mathbb{R}^d$ be the latent variables corresponding to each observed node representations after encoder layers $h_i^{(L)}$. We assume a standard normal prior distributions $p(z_i) \sim \mathcal{N}(0, I)$ and the generative encoder distribution of $q(z_i|X) \sim \mathcal{N}(\mu_i, \exp(\sigma_i))$ where μ_i and σ_i are learned from encoder outputs with a linear layer (i.e. $\mu_i = W_\mu h_i^{(L)} + b_\mu$ and $\sigma_i = W_\sigma h_i^{(L)} + b_\sigma$). The variance is parameterized as an exponential to assure non-negativity. Then the sampled latent variables z_i 's, replacing $h_i^{(L)}$'s, become the inputs to the decoder layer. Let $p(Z) = \prod_{i \in Z} p(z_i)$ and $q(Z|X) = \prod_{i \in Z} q(z_i|X)$. The variational formulation on auto-encoders solves the maximization problem on the posterior $p(X|Z)$ by maximizing the Evidence lower bound (ELBO) [Kingma and Welling 2013]:

$$\text{ELBO}_{\text{VAE}} = \mathbb{E}_q [\log p(\hat{X}|Z)] - \text{KL} [q(Z|X) \| p(Z)] \quad (5)$$

where the first term is the loss for reconstructing the input and $\text{KL}(\cdot \| \cdot)$ is Kullback–Leibler divergence $\mathcal{L}_{\text{div}}$ between prior distribution and likelihood of the latent space Z . From an empirical perspective, the KL-term $\mathcal{L}_{\text{div}}$ regularizes z_i 's to center around the origin, while the reconstruction term ensures sufficient distance between the z_i 's to prevent mode collapse and retain expressiveness. Here, we use the KL-term $\mathcal{L}_{\text{div}}$ to regularize representations of medical concepts in our supervised model, and combine that with cross-entropy loss $\mathcal{L}_{\text{bce}}$ in Equation (4) as the loss function to minimize:

$$\mathcal{L}(y, \hat{y}) = \mathcal{L}_{\text{bce}}(y, \hat{y}) + \text{KL} [q(Z|X) \| p(Z)] \quad (6)$$

For the rest of the paper, we denote our encoder-decoder graph neural network as *Enc-dec* and the variationally regularized formation as *VGNN*. Our implementation is open source and available at: https://github.com/NYUMedML/GNN_for_EHR.

4 EXPERIMENTS

In this study, we test the proposed methods in the context of three clinical tasks: readmission prediction at discharge, based on eICU cohort [Pollard et al. 2018], mortality prediction at 24 hour after admission, based on MIMIC-III cohort [Johnson et al. 2016] and Alzheimer's Disease prediction within 12 to 24 months based on inpatient and outpatient EHR data from NYU Langone Health, (AD-EHR for short). The first two datasets consist of short-term records from ICUs (inpatient); the third dataset is a long-term inpatient and outpatient clinical EHR spanning over 4 years. With the experiments on various type of EHR data, we demonstrate the capacity of our method on EHR representation learning.

All the EHR dataset are partitioned into training, validation and test set by patient unique IDs at ratio of 8-1-1, respectively. We train models based on training sets, tune hyperparameters with validation sets and report the performance of models with test sets.

Table 1: Dataset Statistics Summary (number of average observed features / number of total features)

Dataset	AD-EHR	MIMIC-III	eICU
Diagnosis	10.1/6028	11.5/6778	6.5/3250
Procedures	— / —	4.5/2006	5.0/2212
Lab Values	6.1/3073	54.2/1817	— / —
Demographic	3.3/38	— / —	— / —
# of positives	8174	5377	7051
# of total patients	1613088	50391	41026

4.1 Alzheimer’s Disease Prediction

Alzheimer’s Disease (AD) leads to the majority of dementia, but the cause of AD is poorly understood. This gap in disease mechanism prevents researchers from constructing a medical knowledge graph that associates AD-related diseases and variables. This motivates us to attempt to learn the graph connections from scratch. In this experiment, we use the EHR from NYU Langone Health corresponding to 1.64M distinct patients with unique Medical Record Numbers (MRN), spanning from 2016 to 2019. AD-EHR includes diagnosis, recorded as ICD-10 codes, and lab values recorded as LOINC codes. More details on cohort selection and data preprocessing are described in Appendix A. After the data preprocessing, the whole dataset includes 1.61M patient records and the encounter-based records for each patient are transformed into a 9139-dimensional one-hot vector. The detailed statistics on feature distributions are presented at Table 1. To assess whether aggregation across time lead to any major loss of information, we compare our models built on aggregated EHR data with two encounter-based time series baseline models. As presented in the results section, we find that for the AD-EHR task, aggregated EHR is more effective than time series data.

4.2 MIMIC-III and eICU Predictive Tasks

MIMIC-III and eICU data are two publicly available EHR datasets collected from ICU patients. There has been several clinically meaningful predictive tasks studied for these population, including mortality prediction, readmission prediction and phenotype classification. Choi et al. [Choi et al. 2019] leverages graph structures of EHR in predicting readmission and mortality of ICU patients, using eICU data. We use the same preprocessing steps on eICU to compare our models with the published pre-defined guidance knowledge graphs. Besides eICU, we also empirically evaluate our methods on a more common public dataset, MIMIC-III. Mortality prediction at early days after ICU admission (i.e. 24 hours or 48 hours) is among widely studied and clinically useful predictive tasks based on MIMIC-III dataset [Che et al. 2016; Harutyunyan et al. 2017; Purushotham et al. 2018], although the selection on feature sets varies a lot in different studies. We follow the schemas used in AD-EHR and eICU and make the features more cohesive. We not only extract ICD-9 codes and CPT procedure codes referring to the schema of eICU, but also categorize lab values into buckets according to the schema of AD-EHR. To avoid potential data leakage between mortality and the preventative events immediately preceding it, we only include

the chart events within the first 24 hours after ICU admission as the input for the mortality prediction task. This clinical task setting is based on the benchmark study by [Purushotham et al. 2018].

4.3 Baseline Models

We introduce several baseline models in various machine learning domains to demonstrate the necessity of our design on the model architecture and the statistically significant improved performance of our methodology.

- **Random Forest** Random forest [Breiman 2001] is an ensemble model of decision trees that takes advantage of bagging mechanism to reduce overfitting. It exams whether deep learning methods are over-complex.
- **Multilayer Perceptron** MLP is the multiple feed-forward network, previously used in predictive tasks in EHR [Tangri et al. 2008]. We use MLP as a non-embedding baseline that takes one-hot vectors of disease codes as inputs.
- **Temporal Methods*** We use two temporal deep learning methods on AD-EHR to investigate the impact of time: an RNN with 3-layer LSTM [Lipton et al. 2015], and a CNN with convolutional block [Razavian et al. 2016] in Table 4 (Appendix) on both feature axis and time axis.
- **Neural Bag of Words** Using the average or the sum of embeddings [Iyyer et al. 2015; Mikolov et al. 2013] is a method for representing patients. We embed the medical concepts with embeddings h_i , and represent a patient by averaging embeddings of all the positive features [Iyyer et al. 2015; Mikolov et al. 2013]. Then a feed-forward network outputs the inference on target classes.
- **Transformer** Choi et al. [Choi et al. 2019] adapt Transformer [Vaswani et al. 2017] to learn the graph structures of EHR via interactions among the medical concept embeddings h_i ’s and the visit embeddings v ’s. This work also shows that Transformer has superior performances in EHR predictive tasks than graph convolution networks (GCN) with pre-defined graphs and random graphs.
- **Graph Convolutional Transformer** GCT [Choi et al. 2019] takes advantage of pre-defined medical ontologies as the prior guidance to regularize the attention weights in Transformer. The prior graphs takes the medical concepts as vertices and the connections between diagnosis and procedures, procedures and lab values as edges. The edges are weighted by empirical conditional probabilities derived from the co-occurrence relationship among medical concepts. We use the publicly available codes and hyperparameters to train this baseline.

4.4 Results

In Table 2, we report the performance of different models on three tasks. Since all of three tasks have imbalances class labels, the precision-recall curve is a more informative evaluation metric on the prediction performance than ROC curve [Saito and Rehmsmeier 2015]. To quantify PR-curve, we compute the area under PR-curve (AUPRC) to summarize the curve. We compute the mean and confidence interval by bootstrapping the test data 100 times. The results in Table 2 show that the 99% confidence intervals of our VGNN

Table 2: Model evaluation on the test set using precision-recall curves (99% confidence interval)

Method	AD-EHR		MIMIC-III Mortality	eICU Readmission
	AUPRC	PPV@0.4Recall	AUPRC	AUPRC
Random Forest [Breiman 2001]	0.2316 $\pm$ 0.0043	0.0890 $\pm$ 0.0029	0.5976 $\pm$ 0.0056	0.3614 $\pm$ 0.0049
MLP[Tangri et al. 2008]	0.3775 $\pm$ 0.0050	0.5623 $\pm$ 0.0182	0.6646 $\pm$ 0.0045	0.3639 $\pm$ 0.0045
RNN* [Lipton et al. 2015]	0.2590 $\pm$ 0.0045	0.3038 $\pm$ 0.0041	—	—
CNN* [Razavian et al. 2016]	0.3566 $\pm$ 0.0053	0.4267 $\pm$ 0.0056	—	—
NBOW [Iyyer et al. 2015]	0.3386 $\pm$ 0.0049	0.5265 $\pm$ 0.0138	0.6787 $\pm$ 0.0054	0.3730 $\pm$ 0.0049
Transformer [Choi et al. 2019]	0.3957 $\pm$ 0.0044	0.6844 $\pm$ 0.0165	0.6777 $\pm$ 0.0051	0.3792 $\pm$ 0.0042
GCT [Choi et al. 2019]	0.3409 $\pm$ 0.0040	0.5174 $\pm$ 0.0095	0.6810 $\pm$ 0.0046	0.3794 $\pm$ 0.0045
Enc-dec (Ours)	0.4216 $\pm$ 0.0047	0.6756 $\pm$ 0.0109	0.6962 $\pm$ 0.0051	0.3881 $\pm$ 0.0047
VGNN (Ours)	0.4580 $\pm$ 0.0048	0.7489 $\pm$ 0.0075	0.7102 $\pm$ 0.0046	0.3986 $\pm$ 0.0050

method have no overlap with the baseline models, indicating VGNN outperforms the other baseline models statistically significantly. The optimal hyperparameter settings are in Table 3 (Appendix), and the precision recall curves for three tasks are shown in Figure 7 (Appendix). We observe that the precision for AD-EHR task have a sharp drop around 0.4 recall, so we introduce PPV@0.4Recall to depict the precision at a relatively high classifier threshold for AD-EHR. Table 2 shows the graph-based models are superior to the simple embedding model like NBOW. This comparison demonstrates the importance of connections among different medical concepts. We notice that for the AD-EHR prediction, the performance of GCT are worse than other graph-based methods and close to NBOW. This problem is caused by the nature of dataset: unlike the data from ICU where the variables are measured frequently, AD-EHR is primarily outpatient clinical data. It has more lab values missing and most patients only have diagnosis codes. However, the design of GCT prunes the connections among diagnosis codes, so when lab measures are missing, GCT will reduce to NBOW on diagnosis codes. The insufficiency of GCT on AD prediction demonstrates the importance of learning connections among diagnosis codes. These

connections can be overlooked in short-term ICU records, but they are in fact crucial in learning the overall graph representation of patients. The interaction among the code representation helps reduce the impact of missing codes. The lower performance of the temporal models on AD-EHR shows that for long-term EHR tasks, the temporal information is not a dominating signal.

In addition to the comparison across methods on the same task vertically, we also compare and analyze the results among different tasks. The size of public EHR datasets, like eICU and MIMIC-III, are limited, while in the real-world the hospital system usually has more patients in the EHR database. Table 2 shows that not all of the graph based methods dominate other non-graph based methods such as bag-of-words, while the experiments on AD-EHR demonstrates outstanding performances of models that learn appropriate graphs. The size of dataset can be a cause on this phenomenon. Hence, we also evaluate the model capacity with various data sizes. AD-EHR includes 1.6M patients records, which allows us to analyze the impact of the size of dataset on learning graph structures, by training models with different sizes of training data. We experiment training with different proportion of AD-EHR and evaluate the performance change on the test set. Figure 2 shows that at the low data size level, NBOW is only slightly inferior to VGNN. However, as the training data size grows, the performance of VGNN far exceeds NBOW. Also, the slope of curves in Figure 2 indicates that the graph-based models have more performance gain with the increment on data size than NBOW. This finding explains why graph-based methods have the most performance gain in AD prediction among all three tasks. It also indicates that learning the medical graph connections enlarges the model capacity, so the graph based model has more potential to improve by learning from more data. Therefore, even though the improvement of incorporating graph structure on MIMIC-III and eICU are less than AD-EHR, it follows the trend of performance growth with the size of data.

Figure 2: The evaluation of graph and non-graph based methods on different levels of training data sizes using AD-EHR. The shaded area denotes $\pm$ one standard deviation.

5 DISCUSSION

In this section, we discuss the main components of the development of our models that improve the model performance and provide support for our choices of various settings. We develop the analysis on the methods with both quantitative statistics and qualitative interpretation.

Figure 3: The patterns of the first layer graph with different attention functions. Method 1 is the attention used in Transformer; method 2 is the attention in our models. Results are visualized based on trained model representations of a randomly selected patient from the held-out test set of the AD-EHR task. The patient has between 10 to 20 observed features and a positive outcome label.

5.1 Behavior of Attention Functions

Attention mechanism is widely used in deep learning to express the “soft” links among representations. It is a collections of functions $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$. There are several common ways to compute attentions in previous literature. Transformers model [Vaswani et al. 2017] uses feed-forward networks to create three vectors - key K , query Q and value V . The attention weights are computed by (Method 1):

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (7)$$

Our method in Equation (3) is in another attention style [Bahdanau et al. 2014; Luong et al. 2015; Veličković et al. 2018] (Method 2) that takes the inner product of a learnable vector and the concatenation of two relevant representations. Table 2 shows the performance difference between two methods, Transformer/GCT (Method 1) vs. the Enc-dec/VGNN (Method 2). We unveil the functional properties of two attention functions by analyzing the graph structure patterns of two attention mechanisms. Figure 3 shows that attention function in Equation (7) directs the model to a sparser graph that only includes 0 and 1. This pattern is desirable for tasks such as alignment in machine translation where there are one-to-one bijections between words most of the times. However, in the depicted EHR example, the “hard” edges lead to a problem that the 14 node representations in the example are reduced to 3 node representations.

The attention mechanism in our model (Method 2) is more robust: in the sample of VGNN, we can observe the that head-related diseases D32.0 (Benign neoplasm of cerebral meninges), H65.1 (Other acute non-suppurative otitis media) and J01.1 (Acute frontal sinusitis) have more attention weights, while the representations of the other nodes are also included in the outputs towards deeper layers. In general, compared to method 1, the attention matrix computed by method 2 has positive weights on multiple medical concepts, so the graph representations are able to receive the message different medical embeddings. In Figure 3, we also notice that some row attention weights in Enc-dec are numerically close to each other. This phenomenon leads to the graph representation to be similar.

We then quantitatively analyze the impact of this phenomenon and how the problem is solved by the regularization of VGNN in Section 5.2.

5.2 Singular Value Analysis on Graphs

We now assess the performance of the attention mechanism quantitatively by decomposing the graph attention layer. The adjacency matrices $A^{(l)}$ contain the graph structural information learned by the GNN. Previous studies show that the spectral analysis of graph convolution kernels can lead to a deeper understanding of the Laplacian matrices in graph convolution [Balcilar et al. 2020; Chen et al. 2020]. Similarly, to characterize the learned $A^{(l)}$'s, we start by using singular value decomposition (SVD):

$$A^{(l)} = U \text{diag}(s_i)_{i \in \mathcal{V}_{\text{obs}}} V^T \quad (8)$$

where U, V are the collections of orthonormal vectors and s_i are singular values of $A^{(l)}$. Through SVD, the graph convolution can be transformed to a linear combination of the node representations' projections onto subspaces spanned by these orthonormal basis.

$$A^{(l)} H_{:,j}^{(l-1)} = \sum_{i=1}^{|\mathcal{V}_{\text{obs}}|} s_i u_i v_i^T H_{:,j}^{(l-1)}, (1 \leq j \leq d) \quad (9)$$

where $H_{:,j}$ are the collections of the j^{th} elements in node representations at the layer $l-1$. The magnitudes and directions of transformation given by graph attention layers can be interpreted from singular values and orthonormal basis U . Hence, from Equation (9) we learn that the singular values of $A^{(l)}$ correspond to the magnitude of the messages passed among nodes in the graph layer.

The softmax normalization on attention coefficients restricts adjacency matrices $A^{(l)}$ such that $\sum_{j \in \mathcal{V}_{\text{obs}}} A_{ij}^{(l)} = 1$. Under this constrain, we can have a lower bound on the largest singular value.

Lemma 5.3. *Let matrix $A \in \mathbb{R}^{d \times d}$. Suppose $\sum_{j=1}^d A_{ij} = 1, 1 \leq i \leq d$, and A has singular values $s_1 \geq s_2 \geq \dots \geq s_d$, then $s_1 \geq 1$.*

PROOF. Appendix C. $\square$

By Lemma 5.3, all the first singular values of graph adjacency matrices are greater or equal than 1. Hence, we analyze the distribution of the remaining singular values to interpret the adjacency

matrices, as follows: We perform SVD on the learned adjacency matrices of the heldout testset patients in AD-EHR dataset. We limit this analysis to patients with a positive label, who had more than 10 observed features. In Figure 4, the singular values of first graph layers are visualized. The plots demonstrate that singular values of Enc-dec model has a larger range, but the majority of eigenvalues are close to 0. Vanishing singular values leads to fewer effective dimensions in the graph layers, according to Equation (9). As depicted in Figure 4, more than half of the samples only have one non-zero dimension under Enc-dec. But we observe that they contain at least two non-zero singular values using VGNN. Also, in VGNN model, we can observe that most analysed patients have at least 5 singular values significantly greater than 0. It indicates that the variational regularization enables the graph kernel to avoid mode collapse, and combine more node representations during inference, which helps with generalization of the model across different samples.

Figure 5 shows different clustering behavior of the learned node representations, at varying numbers of non-zero singular values. We visualize the 2D PCA projection of learned h_i representations after graph encoder layer in the VGNN model. Blue dots represent projections of every feature (medical code) observed in the training cohort. In red, we overlay the observed features of two different patients, one with a low count of non-zero singular values (left), and one with high count of non-zero singular values (right). We observe that the increment in non-zero singular values corresponds to more clusters in the projection of the learned representation of the patient features. The clusters that form for each patient also exhibit meaningful medical meanings: i.e. The red nodes with annotations in the right plot correspond to sepsis (red ICD code annotations),

diabetes (green ICD code annotations) and urinary incontinence (black ICD code annotations). Previous medical studies indicate that these diseases are related to AD [Chou et al. 2017; Lee et al. 2018] and are correlated [Hsiao et al. 2015]. This example helps demonstrate that a graph with more significant singular values aids learning more expressive representations, and we can interpret the semantics of the nodes based on the emerged clusters.

5.4 Interpretation of Variationally Regularized Graph Representations

According to variational regularization, and as seen in Equation (6), the learnable parametrized distribution of representations of each observed code is forced closer to a standard Gaussian prior. This strategy prevents the representations from converging to only one direction in high dimensional space. Looking at the patients to which we applied SVD in Figure 4, we measure the compactness of clusters by the mean of ℓ_2 distances each points to the mass center after the encoder graph layer.

$$\text{compactness} = \frac{1}{M} \sum_{i=1}^M \sum_{i \in \mathcal{V}_{\text{obs}}} \frac{\|h_i^{(l)} - c(h_i^{(l)})\|_2}{|\mathcal{V}_{\text{obs}}|} \quad (10)$$

where M is the number of samples, and $c(\cdot)$ is the mass center of all node representations at given layer. The compactness of node representations in Enc-dec is 0.5786, while in VGNN the compactness is 3.1036. Additionally, for the same samples, we visualize the distribution of the values of the learned node representations. Figure 6 in Appendix D shows that the representations in Enc-dec only learned by self-attention are biased.

Figure 4: The magnitudes of singular values of the first graph convolution layer on AD-EHR data. The black vertical lines show the range, and the curves in the vertical direction show the smoothed distribution of magnitudes. Note that more than half of the first singular values of Enc-dec vanish.

Figure 5: 2D PCA projection of learned representations of h_i after graph encoder layer in the VGNN model. Blue dots represent projections of every feature (medical code) observed in the training cohort. In red, we overlay the observed features of two different patients, one with a low count of non-zero singular values (left), and one with high count of non-zero singular values (right). We observe that the increment in non-zero singular values corresponds to more clusters in the projection of the learned representation of the patient features. The clusters that form for each patient also exhibit meaningful medical meanings: i.e. The red nodes with annotations in the right plot correspond to sepsis (red annotated ICD codes), diabetes (green annotated ICD codes) and urinary incontinence (black annotated ICD codes).

Together, these two statistics indicate that the KL-term regularizes the representations to avoid over-clustering and biases.

6 CONCLUSION

Bridging connections among medical concepts contributes towards learning more expressive representation of the EHR, and therefore, improves the performance on predictive tasks in population health. We proposed an encoder-decoder graph neural network that adaptively learns the connections among observed medical codes in EHR. Our method also addresses the problem of learning more expressive representations via variational regularization. We showed that our model achieves superior performance on three EHR-based predictive tasks. Singular value analysis presented here helped explain some of the empirically observed benefits of our proposed regularization, compared to standard graph based methods. Our future studies include exploration of self-supervised learning to further improve generalization of graph based EHR representation learning.

REFERENCES

- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. *CoRR* abs/1607.06450 (2016). arXiv:1607.06450 <http://arxiv.org/abs/1607.06450>
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. <http://arxiv.org/abs/1409.0473> cite arxiv:1409.0473Comment: Accepted at ICLR 2015 as oral presentation.
- Muhammet Balcilar, Guillaume Renton, Pierre Heroux, Benoit Gauzere, Sebastien Adam, and Paul Honeine. 2020. Bridging the Gap Between Spectral and Spatial Domains in Graph Neural Networks. arXiv:2003.11702 [cs.LG]
- Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (2001), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2013. Spectral Networks and Locally Connected Networks on Graphs. *CoRR* abs/1312.6203 (2013).
- Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David A. Sontag, and Yan Liu. 2016. Recurrent Neural Networks for Multivariate Time Series with Missing Values. *CoRR* abs/1606.01865 (2016). arXiv:1606.01865 <http://arxiv.org/abs/1606.01865>
- Zhiqian Chen, Fanglan Chen, Lei Zhang, Taoran Ji, Kaiqun Fu, Liang Zhao, Feng Chen, and Chang-Tien Lu. 2020. Bridging the Gap between Spatial and Spectral Domains: A Survey on Graph Neural Networks. arXiv:2002.11867 [cs.LG]
- Yu Cheng, Feng Wang, Ping Zhang, and Jianying Hu. 2016. Risk Prediction with Electronic Health Records: A Deep Learning Approach. In *SDM*.
- Edward Choi, Mohammad Taha Bahadori, Elizabeth Searles, Catherine Coffey, Michael Thompson, James Bost, Javier Tejedro-Sojo, and Jimeng Sun. 2016a. Multi-Layer Representation Learning for Medical Concepts. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California, USA) (*KDD '16*). Association for Computing Machinery, New York, NY, USA, 1495–1504. <https://doi.org/10.1145/2939672.2939823>
- Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F. Stewart, and Jimeng Sun. 2016b. GRAM: Graph-based Attention Model for Healthcare Representation Learning. *CoRR* abs/1611.07012 (2016). arXiv:1611.07012 <http://arxiv.org/abs/1611.07012>
- Edward Choi, Mohammad Taha Bahadori, and Jimeng Sun. 2015. Doctor AI: Predicting Clinical Events via Recurrent Neural Networks. *CoRR* abs/1511.05942 (2015). arXiv:1511.05942 <http://arxiv.org/abs/1511.05942>
- Edward Choi, Cao Xiao, Walter F. Stewart, and Jimeng Sun. 2018. MiME: Multilevel Medical Embedding of Electronic Health Records for Predictive Healthcare. *CoRR* abs/1810.09593 (2018). arXiv:1810.09593 <http://arxiv.org/abs/1810.09593>
- Edward Choi, Zhen Xu, Yujia Li, Michael W. Dusenberry, Gerardo Flores, Yuan Xue, and Andrew M. Dai. 2019. Graph Convolutional Transformer: Learning the Graphical Structure of Electronic Health Records. *CoRR* abs/1906.04716 (2019). arXiv:1906.04716 <http://arxiv.org/abs/1906.04716>
- Youngduck Choi, Chill Chiu, and David Sontag. 2016c. Learning Low-Dimensional Representations of Medical Concepts. *AMIA Joint Summits on Translational Science*

- proceedings. *AMIA Summit on Translational Science* 2016 (07 2016), 41–50.
- Hsusan Chou, Jiunn-Tay Lee, Chun-Chieh Lin, Yueh-Feng Sung, Che-Chen Lin, Chih-Hsin Muo, Fu-Chi Yang, Chi Pang Wen, I-Kuan Wang, Chia-Hung Kao, Chung Hsu, and Chun-Hung Tseng. 2017. Septicemia is associated with increased risk for dementia: A population-based longitudinal study. *Oncotarget* 8 (09 2017). <https://doi.org/10.18632/oncotarget.20899>
- Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. 2020. On the Relationship between Self-Attention and Convolutional Layers. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=HJlnC1rKPB>
- Ehsan Hajiramezani, Arman Hasanzadeh, Nick Duffield, Krishna R Narayanan, Mingyuan Zhou, and Xiaoning Qian. 2019. Variational Graph Recurrent Neural Networks. arXiv:1908.09710 [cs.LG]
- Hrayr Harutyunyan, Hrant Khachatryan, David Kale, and Aram Galstyan. 2017. Multi-task Learning and Benchmarking with Clinical Time Series Data. *Scientific Data* 6 (03 2017). <https://doi.org/10.1038/s41597-019-0103-9>
- Arman Hasanzadeh, Ehsan Hajiramezani, Nick Duffield, Krishna R. Narayanan, Mingyuan Zhou, and Xiaoning Qian. 2019. Semi-Implicit Graph Variational Auto-Encoders. arXiv:1908.07078 [cs.LG]
- Mikael Henaff, Joan Bruna, and Yann LeCun. 2015. Deep Convolutional Networks on Graph-Structured Data. *CoRR* abs/1506.05163 (2015). arXiv:1506.05163 <http://arxiv.org/abs/1506.05163>
- I. Higgins, Loïc Matthey, A. Pal, C. Burgess, Xavier Glorot, M. Botvinick, S. Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR*.
- Chih-Yen Hsiao, Huang-Yu Yang, Chih-Hsiang Chang, Hsing-Lin Lin, Chao-Yi Wu, Meng-Chang Hsiao, Peir-Haur Hung, Su-Hsun Liu, Cheng-Hao Weng, cheng-chia Lee, Tzung-Hai Yen, Yung-Chang Chen, and Tzu-Chin Wu. 2015. Risk Factors for Development of Septic Shock in Patients with Urinary Tract Infection. *BioMed Research International* 2015 (07 2015), 7 pages. <https://doi.org/10.1155/2015/717094>
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep Unordered Composition Rivals Syntactic Methods for Text Classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Beijing, China, 1681–1691. <https://doi.org/10.3115/v1/P15-1162>
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3 (2016), 160035.
- Diederik P Kingma and Max Welling. 2013. Auto-Encoding Variational Bayes. <http://arxiv.org/abs/1312.6114> cite arxiv:1312.6114.
- Thomas N. Kipf and Max Welling. 2016a. Semi-Supervised Classification with Graph Convolutional Networks. *CoRR* abs/1609.02907 (2016). arXiv:1609.02907 <http://arxiv.org/abs/1609.02907>
- Thomas N Kipf and Max Welling. 2016b. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* (2016).
- Hee Lee, Hye Seo, Hee Cha, Yun Yang, Soo Kwon, and Soo Jin Yang. 2018. Diabetes and Alzheimer's Disease: Mechanisms and Nutritional Aspects. *Clinical Nutrition Research* 7 (10 2018), 229. <https://doi.org/10.7762/cnr.2018.7.4.229>
- Yikuan Li, Shishir Rao, Jose Roberto Ayala Solares, Abdelaali Hassaine, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Gholamreza Salimi Khorshidi. 2019. BEHRT: Transformer for Electronic Health Records. *CoRR* abs/1907.09538 (2019). arXiv:1907.09538 <http://arxiv.org/abs/1907.09538>
- Zachary Lipton, David Kale, Charles Elkan, and Randall Wetzell. 2015. Learning to Diagnose with LSTM Recurrent Neural Networks. (11 2015).
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. *CoRR* abs/1508.04025 (2015). arXiv:1508.04025 <http://arxiv.org/abs/1508.04025>
- Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *CoRR* abs/1301.3781 (2013).
- Riccardo Miotto, Li Li, Brian A. Kidd, and Joel T. Dudley. 2016. Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records. *Scientific Reports* 6 (17 May 2016), 26094 EP –. <https://doi.org/10.1038/srep26094> Article.
- Tom Pollard, Alistair Johnson, Jesse Raffa, Leo Celi, Roger Mark, and Omar Badawi. 2018. The eICU Collaborative Research Database, a freely available multi-center database for critical care research. *Scientific Data* 5 (09 2018), 180178. <https://doi.org/10.1038/sdata.2018.178>
- Victor Prokhorov, Ehsan Shareghi, Yingzhen Li, Mohammad Taher Pilehvar, and Nigel Collier. 2019. On the Importance of the Kullback-Leibler Divergence Term in Variational Autoencoders for Text Generation. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*. Association for Computational Linguistics, Hong Kong, 118–127. <https://doi.org/10.18653/v1/D19-5612>
- Sanjay Purushotham, Chuizheng Meng, Zhengping Che, and Yan Liu. 2018. Benchmarking deep learning models on large healthcare datasets. *Journal of Biomedical Informatics* 83 (2018), 112 – 134. <https://doi.org/10.1016/j.jbi.2018.04.007>
- Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jonathon Shlens. 2019. Stand-Alone Self-Attention in Vision Models. *CoRR* abs/1906.05909 (2019). arXiv:1906.05909 <http://arxiv.org/abs/1906.05909>
- Ali Razavi, Aaron van den Oord, Ben Poole, and Oriol Vinyals. 2019. Preventing Posterior Collapse with delta-VAEs. *CoRR* abs/1901.03416 (2019). arXiv:1901.03416 <http://arxiv.org/abs/1901.03416>
- Narges Razavian et al. 2016. Multi-task Prediction of Disease Onsets from Longitudinal Lab Tests. *CoRR* abs/1608.00647 (2016). arXiv:1608.00647 <http://arxiv.org/abs/1608.00647>
- Takaya Saito and Marc Rehmsmeier. 2015. The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. *PLOS ONE* 10, 3 (03 2015), 1–21. <https://doi.org/10.1371/journal.pone.0118432>
- Benjamin Shickel, Patrick Tighe, Azra Bihorac, and Parisa Rashidi. 2017. Deep EHR: A Survey of Recent Advances on Deep Learning Techniques for Electronic Health Record (EHR) Analysis. *CoRR* abs/1706.03446 (2017). arXiv:1706.03446 <http://arxiv.org/abs/1706.03446>
- Huan Song, Deepta Rajan, Jayaraman J. Thiagarajan, and Andreas Spanias. 2017. Attend and Diagnose: Clinical Time Series Analysis using Attention Models. arXiv:1711.03905 [stat.ML]
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 56 (2014), 1929–1958. <http://jmlr.org/papers/v15/srivastava14a.html>
- Navdeep Tangri et al. 2008. Predicting technique survival in peritoneal dialysis patients: Comparing artificial neural networks and logistic regression. *Nephrology, dialysis, transplantation : official publication of the European Dialysis and Transplant Association - European Renal Association* 23 (05 2008), 2972–81. <https://doi.org/10.1093/ndt/gfn187>
- Louis C. Tiao, Pantelis Elinas, Harrison Nguyen, and Edwin V. Bonilla. 2019. Variational Spectral Graph Convolutional Networks. *CoRR* abs/1906.01852 (2019). arXiv:1906.01852 <http://arxiv.org/abs/1906.01852>
- Truyen Tran, Trang Pham, Dinh Phung, and Svetha Venkatesh. 2016. DeepCare: A Deep Dynamic Memory Model for Predictive Medicine.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems* 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5998–6008. <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. *International Conference on Learning Representations* (2018). <https://openreview.net/forum?id=rjXmpikCZ>

A DATA DETAILS

We preprocessed 1.6M EHR with 100K features, including diagnosis, labs and procedures. Since the ICD-10 codes are hierarchical and get granular at different levels, we merge the codes that share same number up to the first place after the decimal point to the first codes in their subdivisions. We set the target variable by aggregating all the ICD-10 codes for AD¹ and exclude all these codes from our feature set. The original data is encounter-wise. We track each patient by partitioning his/her encounters into history window (before 2016.02.19), feature window (2016.02.20 - 2017.02.19) and gap window (2017.02.20 - 2018.02.19). We use observations from encounters in the feature window as our inputs, and exclude patients who are AD positive within any three windows. These patients are dropped to avoid data leakage, as our goal is to predict new-onset AD 12-24 months in the future. We then aggregate all the encounters in the feature window temporally to allow the model to focus on learning graph representations between the features, rather than focusing on sparsity patterns in the time dimension.

We use the following schema to aggregate the encounter data. The diagnosis features are set to observed/positive if they have positive outcomes in any encounter. The lab values are binned into ranges of -10, -3, -1, -0.5, 0.5, 1, 3, 10 standard deviations where the statistics of each lab are computed independently from training set. The lab value features are defined as observed/positive if the lab values of any encounter fall into the corresponding range.

B TRAINING DETAILS

Graph neural networks can be memory consuming. In our setting, as the graph is fully-connected, for $O(n)$ nodes, $O(n^2)$ edge weights should be allocated, where n is the number of positive/observed features. However, for each patient graph, only a few nodes have observed values, so we implement the model in sparse form with Pytorch 1.1.0 to free the memory of unobserved nodes and their edges in the graph of each patient.

To improve the robustness and ability of the model inferring missing features through graph, we randomly mask 10% of nodes during training. Since all of our predictive tasks have imbalanced labels (Table 2), we use weighted cross-entropy loss based on class weights. For AD-EHR data, the labels are extremely imbalanced, so the weighted loss cannot effectively improve the performance. Hence, we upsample positive samples of the training set by 50 times to see more positive samples in given epochs. We also randomly downsample 80% negative patients with age under 50 each epoch to accelerate training, as they may not contain significant signals related to AD. For MIMIC-III and eICU, since the label are less imbalanced, we only upsample the positive samples 2 time. Validation and test set retain their original distribution.

We tune the hyper-parameters including the number of heads (1-4) and layers of graph in graph based models (1-3), the number of layers in feed-forward networks (1-2), embedding sizes (128-1024), dropout rates [0, 1] and learning rates [10^{-5} , 10^{-3}]. The optimal hyper-parameters in our experiments are listed in Table 3. Learning

rate decay is used to avoid overfitting. We half the learning rate if the AUPRC stops growing for two epochs. For these experiments, we use Tesla V100 GPUs to train our model.

C PROOF OF LEMMA

Lemma C.1. Let matrix $A \in \mathbb{R}^{d \times d}$. Suppose $\sum_{j=1}^d A_{ij} = 1$ and A has singular values $s_1 \geq s_2 \geq \dots \geq s_d$, then $s_1 \geq 1$.

PROOF. Let $e = (1, 1, \dots, 1)^T$, we have $A^T e = e$, because $\sum_{j=1}^d A_{ij} = 1$. Therefore, 1 is an eigenvalue of A^T . Since A and A^T has the same eigenvalues with same multiplicities, 1 is an eigenvector of A .

According to the Min-max theorem, let $X \subseteq \mathbb{R}^d$,

$$s_1 = \min_{\dim(X)=d} \max_{\|x\|_2=1, x \in X} \|Ax\|_2 = \max_{\|x\|_2=1} \|Ax\|_2$$

Suppose λ is the greatest eigenvalue of A and x be an eigenvector corresponding to λ such that $\|x\|_2 = 1$, we have $\|Ax\|_2 = |\lambda|$. Therefore, $s_1 \geq |\lambda| \geq 1$. $\square$

D SUPPLEMENTARY FIGURES AND TABLES

Figure 6: Distribution of node embedding entries in test samples of AD-EHR.

¹Agency for Healthcare Research and Quality(AHRQ) at United States Department of Health and Human Services defines the family of Alzheimer's related dementia, including ICD-10 codes: F01.50, F01.51, F02.80, F02.81, F03.90, F03.91, F04, F05, F07.0, F07.81, F07.89, F07.9, F09, F48.2, G30.0, G30.1, G30.8, G30.9, G31.01, G31.09, G31.1, G31.83, R41.81, R54.

Figure 7: Precision-Recall curves of experiments corresponding to Table 2. The precision for AD-EHR task have a sharp drop around 0.4 recall, so we introduce PPV@0.4Recall to depict the precision at a relatively high classifier threshold for AD-EHR.

Task	Model	No. Heads	No. Graph Layers	No. Feed-forward Layers	Learning Rate	Dropout Rate	Embedding Size	Batch Size
AD-EHR <i>Alzheimer's Disease Prediction</i>	MLP	—	—	2	0.0001	0.3	—	64
	CNN*	—	—	2	0.0003	0.2	—	64
	RNN*	—	—	2	0.0003	0.2	—	64
	NBOW	—	—	2	0.0003	0.3	1024	64
	Transformer	1	3	2	0.0002	0.4	768	32
	GCT	1	3	2	0.0002	0.1	768	32
	Enc-dec	1	2	1	0.0001	0.4	768	32
	VGNN	1	2	1	0.00003	0.1	1024	32
MIMIC-III <i>Mortality Prediction</i>	MLP	—	—	2	0.0001	0.5	—	64
	NBOW	—	—	3	0.0003	0.4	128	64
	Transformer	1	3	2	0.0002	0.4	256	32
	GCT	1	3	2	0.0002	0.1	256	32
	Enc-dec	1	2	1	0.0001	0.4	768	32
	VGNN	1	2	1	0.0001	0.2	768	32
eICU <i>Readmission Prediction</i>	MLP	—	—	2	0.0001	0.5	—	64
	NBOW	—	—	3	0.0003	0.4	128	64
	Transformer	1	3	2	0.0002	0.45	128	32
	GCT	1	3	2	0.00022	0.08	128	32
	Enc-dec	1	2	1	0.0001	0.5	128	32
	VGNN	1	2	1	0.0001	0.4	128	32

Table 3: Hyperparameters of experiments in Table 2. The hyperparameter settings of the previous study on the same dataset remain the same.

Type	Block	Hyperparameters
Vectical Conv	Conv2d	$k-9139 \times 1; c-64; p-0; s-1$
	ReLU	
	BatchNorm2d	
Vectical Conv	Conv2d	$k-64 \times 1; c-128; p-0; s-1$
	ReLU	
	BatchNorm2d	
Temporal	Maxpool1d	$k-5; p-1; s-1$
Temporal Conv	Conv1d	$k-5; c-256; p-1; s-1$
	ReLU	
	BatchNorm1d	
	Avgpool1d	$k-T; p-1; s-1$
FC $\times 2$	Linear	256×256
	ReLU	
	BatchNorm1d	
	Dropout	
	Linear	256×1

Table 4: Architecture of CNN for temporal signals. k is kernel size; c is the output channel; s is the stride; p is the padding size. T is the maximum number of encounters observed.