

Doctor2Vec: Dynamic Doctor Representation Learning for Clinical Trial Recruitment

Siddharth Biswal,^{1,2} Cao Xiao,¹ Lucas M. Glass,¹ Elizabeth Milkovits,¹ Jimeng Sun²

¹Analytic Center of Excellence, IQVIA, Cambridge, USA

²Computational Science and Engineering, Georgia Institute of Technology, Atlanta, USA

Abstract

Massive electronic health records (EHRs) enable the success of learning accurate patient representations to support various predictive health applications. In contrast, doctor representation was not well studied despite that doctors play pivotal roles in healthcare. How to construct the right doctor representations? How to use doctor representation to solve important health analytic problems? In this work, we study the problem on *clinical trial recruitment*, which is about identifying the right doctors to help conduct the trials based on the trial description and patient EHR data of those doctors. We propose Doctor2Vec which simultaneously learns 1) doctor representations from EHR data and 2) trial representations from the description and categorical information about the trials. In particular, Doctor2Vec utilizes a dynamic memory network where the doctor's experience with patients are stored in the memory bank and the network will dynamically assign weights based on the trial representation via an attention mechanism. Validated on large real-world trials and EHR data including 2,609 trials, 25K doctors and 430K patients, Doctor2Vec demonstrated improved performance over the best baseline by up to 8.7% in PR-AUC. We also demonstrated that the Doctor2Vec embedding can be transferred to benefit data insufficiency settings including trial recruitment in less populated/newly explored country with 13.7% improvement or for rare diseases with 8.1% improvement in PR-AUC.

Introduction

The rapid growth of electronic health record (EHR) data and other health data enables the training of complex deep learning models to learn patient representations for disease diagnosis (Choi et al. 2018; 2016; Xiao, Choi, and Sun 2018), risk prediction (Xiao et al. 2018), patient subtyping (Baytas et al. 2017; Che et al. 2017), and medication recommendation (Shang et al. 2019b; 2019a). However, almost all existing works focus on modeling patients. Deep neural networks for doctor representation learning are lacking.

Doctors play pivotal roles in connecting patients and treatments, including recruiting patients into clinical trials for drug development and treating and caring for their patients. Thus an effective doctor representation will better support

a wider range of health analytic tasks. For example, identifying the right doctors to conduct the trials *site selection* so as to improve the chance of completion of the trials (Hurtado-Chong et al. 2017) and doctor recommendation for patients (Xu et al. 2018).

In this work, we focus on studying the *clinical trial recruitment* problem using doctor representation learning. Current standard practice calculates the median enrollment rate¹ for a therapeutic area as the predicted enrollment success rate for all participating doctors, which is often inaccurate. In addition, some develop a multi-step manual matching process for site selection which is labor-intensive (Hurtado-Chong et al. 2017; Potter et al. 2011). Recently, deep neural networks were applied on site selection tasks via static medical concept embedding using only frequent medical codes and simple term matching to trials (Gligorijevic et al. 2019). Despite the success, two challenges remain open.

1. Existing works do not capture the time-evolving patterns of doctors experience and expertise encoded in EHR data of patients that the doctor have seen;
2. Existing works learn a static doctor representation. However, in practice given a trial for a particular disease, the doctor's experience of relevant diseases are more important. Hence the doctor representation should change based on the corresponding trial representation.

To fill the gap, we propose Doctor2Vec which simultaneously learns 1) doctor representations from longitudinal patient EHR data and 2) trial embedding from the multimodal trial description. In particular, Doctor2Vec leverages a dynamic memory network where the representations of patients seen by the doctor are stored as memory while trial embedding serves as queries for retrieving from the memory. Doctor2Vec has the following contributions.

1. **Patient embedding as a memory for dynamic doctor representation learning.** We represent doctors' evolving experience based on the representations from the doctors' patients. The patient representations are stored as a memory for dynamic doctor representation extraction.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Enrollment rate of a doctor is the number of patients enrolled by a doctor to the trial.

2. **Trial embedding as a query for improved doctor selection.** We learn hierarchical clinical trial embedding where the unstructured trial descriptions were embedded using BERT (Devlin et al. 2018). The trial embedding serves as queries of the memory network and will attend over patient representation and dynamically assign weights based on the relevance of doctor experience and trial representation to obtain the final context vector for an optimized doctor representation for a specific trial.

We evaluated Doctor2Vec using large scale real-world EHR and trial data for predicting trial enrollment rates of doctors. Doctor2Vec demonstrated improved performance in site selection task over the best baselines by up to 8.7% in PR-AUC. We also demonstrated that the Doctor2Vec embedding can be transferred to benefit data insufficiency settings including trial recruitment in less populated/newly explored countries or for rare diseases. Experimental results show for the country transfer, Doctor2Vec achieved 13.7% relative improvement in PR-AUC over the best baseline. While for embedding transfer to rare disease trials, Doctor2Vec achieved 8.1% relative improvement in PR-AUC over the best baseline.

Related Works

Deep Patient Representation Learning The collection of massive EHR data has motivated the use of deep learning for accurate patient representation learning and disease or risk prediction (Xiao, Choi, and Sun 2018; Fu et al. 2019; Baytas et al. 2017; Choi et al. 2018). In this work, we learn hierarchical patient representation in a similar way as (Choi et al. 2018). But our focus is to construct doctor representation based on the embedding of their patients.

Machine Learning Based Clinical Trial Recruitment Previously clinical trial enrollment either relies on simple statistics (e.g., medium enrollment) or manual matching (Hurtado-Chong et al. 2017). With the collection of clinical trial data, there has been some effort on developing machine learning-based models for trial site selection. For example, (van den Bor et al. 2017) applied LR with L1 regularization to determine a subset that is optimal for predicting site enrollment success. More recently, (Gligorijevic et al. 2019) learns static medical concept embedding and matches them to features derived from trial terms for site selection. However, no existing works learn trial embedding from multi-modal trial data and automatically match them to most relevant doctors.

Memory Augmented Neural Networks (MANN) have shown initial success in NLP research areas such as question answering (Weston, Chopra, and Bordes 2015; Sukhbaatar et al. 2015; Miller et al. 2016; Kumar et al. 2016). Memory Networks (Weston, Chopra, and Bordes 2015) and Differentiable Neural Computers (DNC) (Graves et al. 2016) proposed to use external memory components to assist the deep neural networks in remembering and storing things. After that, various MANN based models have been proposed such as (Sukhbaatar et al. 2015; Kumar et al. 2016; Miller et al. 2016). In healthcare, memory networks can

be valuable due to their capacities in memorizing medical knowledge and patient history. DMNC (Le, Tran, and Venkatesh 2018) proposed a MANN model for medication combination recommendation task using EHR data alone. In (Shang et al. 2019b), the authors use a memory component to fuse multi-model graphs as a memory bank to facilitate medication recommendation.

Method

Problem Formulation

Definition 1 (Doctor Records) For each doctor, the clinical experience can be represented as a sequence of patients that the doctor has seen, denoted as $\mathbf{C}(\mathbf{m}) = \{\mathbf{P}_1^{(m)}, \mathbf{P}_2^{(m)}, \dots, \mathbf{P}_k^{(m)}\}$ where m indicates the m -th doctor. Here each patient can also be represented as a sequence of multivariate observations $\mathbf{P}(\mathbf{k}) = \{\mathbf{v}_1^{(k)}, \mathbf{v}_2^{(k)}, \dots, \mathbf{v}_T^{(k)}\}$ where k indicates the k -th patient and subscript $1, 2, \dots, T$ indicates different visits for the k -th patient. Each visit $\mathbf{v}_t^{(k)}$ is the combination of diagnosis codes \mathbf{c}_d , medication codes \mathbf{c}_m , and procedure codes \mathbf{c}_p . Medical codes $\mathbf{c}_d, \mathbf{c}_p, \mathbf{c}_m \in \{0, 1\}^{|C|}$ are represented as multi-hot vectors and $|C|$ represent the total size of the code sets. We also use the demographic information available about doctors and patients available. These static information are denoted as \mathbf{Doc}_{static} which are used with other features.

Definition 2 (Clinical Trial Data) Clinical trial data comprises of two data modalities: the trial descriptions in unstructured text and the categorical features such as trial phase, primary indication, primary outcome, secondary outcome, and study type. We denote each clinical trial as a combination of text data and categorical data.

$$\begin{aligned} \mathbf{Q}(\mathbf{cat})(\mathbf{l}) &= \{\mathbf{f}_1^{(l)}, \mathbf{f}_2^{(l)}, \dots, \mathbf{f}_v^{(l)}\} \\ \mathbf{Q}(\mathbf{l}) &= [\mathbf{Q}(\mathbf{cat})(\mathbf{l}); \mathbf{Q}(\mathbf{text})(\mathbf{l})] \end{aligned}$$

where $\mathbf{l} \in \{1, 2, \dots, L\}$ is the index of clinical trials, \mathbf{f}_i is the representation for the categorical trial features, and v is the number of categorical features in a trial.

Table 1: Notations used in Doctor2Vec.

Symbol	Definition and description
\mathcal{D}	Set of unique diagnosis codes
\mathcal{P}	Set of unique procedure codes
\mathcal{M}	Set of unique medication codes
$\mathbf{C}(\mathbf{m})$	Notation for a doctor
$\mathbf{P}_k^{(m)}$	Notation for a patient for m -th doctor
$\mathbf{v}_t^{(k)}$	the t -th visit of Patient k
$\mathbf{Q}(\mathbf{l})$	The \mathbf{l} -th clinical trial
$\mathbf{Q}(\mathbf{cat})^{(\mathbf{l})}$	Categorical features of the \mathbf{l} -th clinical trial
$\mathbf{Q}(\mathbf{text})^{(\mathbf{l})}$	Text features of the \mathbf{l} -th clinical trial
$\mathbf{I}_f(\mathbf{k})$	Final patient representation
\mathbf{Doc}_{emb}	Vector representation of a Doctor
\mathbf{Doc}_{static}	Static features of a Doctor

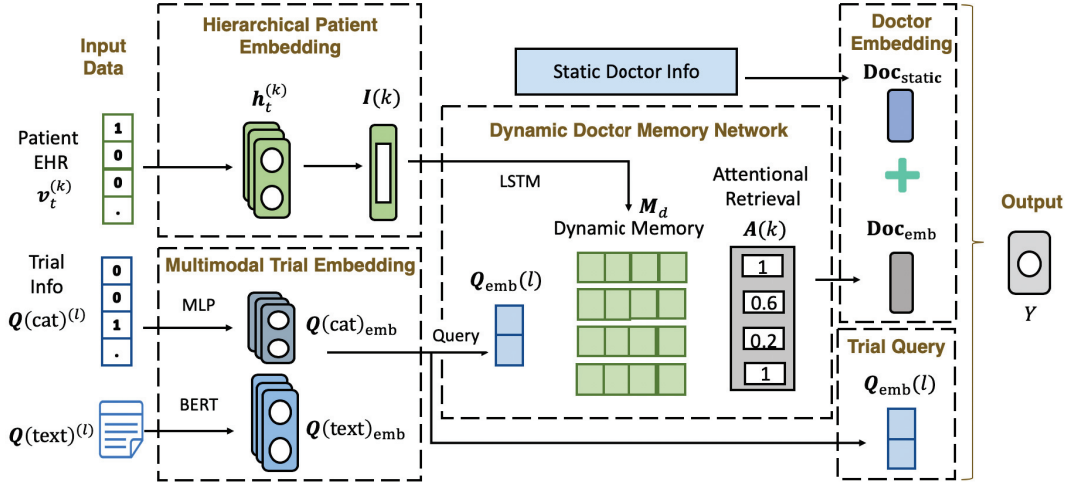


Figure 1: The Doctor2Vec Framework. (1) Hierarchical patient embedding: We obtain patient embeddings $\mathbf{h}_t^{(k)}$ from patient visits using a Bi-LSTM with attention module. Unstructured text and categorical data are modeled using MLP and BERT respectively. (2) Multimodal clinical trial information embedding: The obtained clinical trial embeddings $\mathbf{Q}_{emb}(l)$ are fused together to form the query vector for the memory network. (3) Memory network module: This query vector to used to attend over the memory bank is obtained from combination of patient embeddings. The attentional vector is used to obtain the final doctor embedding. We combine the doctor embedding \mathbf{Doc}_{emb} with clinical trial embedding $\mathbf{Q}_{emb}(l)$ and static information about the doctors \mathbf{Doc}_{static} to predict the enrollment rate of the clinical trial (output).

Doctor2Vec Framework

As illustrated in Fig. 1, Doctor2Vec includes the following components: *hierarchical patient embedding*, *multimodal clinical trial information embedding*, a *memory network module* for dynamical doctor information retrieval, and a memory output module for generating the current time doctor representation. This doctor representation is used with clinical trial representation to predict the *enrollment rate* of the clinical trial. Next, we will first introduce these modules and then provide details of training and inference.

Hierarchical Patient Embedding As discussed before, a doctor has seen a set of patients during his/her medical practice. We model that the doctor’s experience and expertise by the function over the embeddings of their patients’ EHR data. If a patient has seen two doctors, the corresponding portion of the EHR data will be modeled as two separate patients for both doctors, respectively.

To learn patient representation, motivated by (Choi et al. 2018), we leverage the inherent multilevel structure of EHR data to learn patient embedding hierarchically. The hierarchical structure includes the patients on the top, followed by visits the patient experiences over time, then at the leaf level the set of diagnosis, procedure, and medication codes recorded for each visit. Here given clinical visits of a patient, denoted as $\mathbf{v}_t^{(k)}$, we firstly pass these visits through a multi-layer perception (MLP) to get visit embedding $\mathbf{h}_t^{(k)}$ as follows.

$$\mathbf{h}_t^{(k)} = \mathbf{W}_{emb} \mathbf{v}_t^{(k)} \quad (1)$$

Without ambiguity, we will ignore the patient index (k) for brevity. Next to learn a patient embedding based on a se-

quence of visit embeddings, we pass visit embedding \mathbf{h}_k to a bi-directional long short-term memory (bi-LSTM) networks and then add an attention layer on top of the bi-LSTM to attend on important visits.

$$\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_t = \text{bi-LSTM}(\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_t) \quad (2)$$

$$\mathbf{e}_t = \mathbf{w}_\alpha^T \mathbf{g}_t + b_\alpha \quad (3)$$

$$\alpha_1, \alpha_2, \dots, \alpha_t = \text{softmax}(\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_t) \quad (4)$$

Then we obtain final patient context vector $\mathbf{I}(k)$ by summing over attended visit representations as given by Eq. 5.

$$\mathbf{I}(k) = \sum \alpha_t \cdot \mathbf{h}_t \quad (5)$$

Here $\mathbf{I}(1) \dots \mathbf{I}(k)$ are patient representations that will be fed into dynamic doctor memory bank as memory vectors.

Multimodal Trial Information Embedding Clinical Trials are conducted to evaluate a specific drug or procedure. We use public and private information about clinical trials. Here we obtained clinical trial descriptions from publicly available clinical trial database clinicaltrials.gov. The collected trial description comprises multiple data modalities, including unstructured text and categorical features. In this module, we employ a multi-modal method to embed them in a shared space.

1. **Unstructured Text.** Each trial has a text description for inclusion and exclusion criteria which describe the requirements and restrictions for recruiting to the trial. For unstructured text, we applied the BERT (Devlin et al. 2018) model. BERT builds on the Transformer (Vaswani et al. 2017) architecture and improves the pre-training

using a masked language model for bidirectional representation. Essentially BERT model is a stack of Transformer blocks. Each transformer block is a combination of self-attention block and feedforward layer. **Pretraining:** We use the same pre-training techniques as in (Devlin et al. 2018) (1) Masked Language Model: This modeling task consists of masking 15% of the input tokens and using the model to predict the masked tokens. This is trained with cross entropy loss. (2) Next sentence prediction task: In this task two sentences are fed to BERT. The model outputs a binary prediction of whether these two sentences are in consecutive order. The final pre-training objective function based on the two tasks is the sum of the log-likelihood of the masked tokens and the log-likelihood of the binary variable indicating whether two sentences are consecutive. We pretrained a BERT model using MIMIC text data (Johnson et al. 2017) to extract embeddings for each word and combine the word embeddings to obtain the final text embeddings. In order to pre-train BERT on text corpus, we first obtain preprocessed data as required by BERT which includes tokenization, indexing and masking. The procedure is formulated as below. Denote unstructured text associated with each trial as $\mathbf{Q}(\text{text})^{(l)}$, the embedding $\mathbf{Q}(\text{text})_{\text{emb}}$ is given by Eq. 6. We average over embeddings obtained from each word to compute the final embedding for the entire document.

$$\mathbf{Q}(\text{text})_{\text{emb}} = \text{BERT}(\mathbf{Q}(\text{text})^{(l)}) \quad (6)$$

- Categorical Features.** The categorical features of each clinical trial include the geographic location of the trial, hospital system, primary and secondary therapeutic area, pharmaceutical company information, phase of the trial, condition or disease, objectives of the trial, intervention model, etc. More details of the categorical features in the Appendix. The dimensions of these categorical variables ranged from 18 to 1456. For these features, we first encode them using one-hot vectors and then pass the one-hot vectors through multi-layer perception (MLP) layer. This can be expressed as below. Denote the categorical features as $\mathbf{Q}(\text{cat})^{(l)}$, the categorical feature embedding is obtained as in Eq. 7.

$$\mathbf{Q}(\text{cat})_{\text{emb}} = \mathbf{W}_c \mathbf{Q}(\text{cat})^{(l)} + \mathbf{b}_c \quad (7)$$

After obtaining embeddings from both types of clinical trial data, we fuse the embeddings from categorical data $\mathbf{Q}(\text{cat})_{\text{emb}}$ and text data $\mathbf{Q}(\text{text})_{\text{emb}}$ to obtain the final embedding $\mathbf{Q}_{\text{emb}}(l)$. We fuse these two embeddings weighted multiplicative fashion as in Eq. 8.

$$\mathbf{Q}_{\text{emb}}(l) = (\mathbf{W}_{ci} \mathbf{Q}(\text{cat})_{\text{emb}} + \mathbf{b}_{ci}) \odot (\mathbf{W}_{ti} \mathbf{Q}(\text{text})_{\text{emb}} + \mathbf{b}_{ti}) \quad (8)$$

where \odot is element-wise multiplication. The clinical trial embeddings $\mathbf{Q}_{\text{emb}}(l)$ will be fed into dynamic doctor memory network as the query to extract related patient representation memory.

Dynamic Doctor Memory Networks Since each doctor sees a diverse set of patients, doctor representation should be dynamically constructed for a given trial as opposed to a static embedding vector staying the same for all trials. The way we achieved that is by a dynamic memory network where patients are stored as memory vectors of the doctor. Then using a trial embedding as a query, we fetch the relevant patient vectors from the memory bank and dynamically assemble a doctor representation for this trial.

Inspired by (Weston, Chopra, and Bordes 2015), four memory components $\mathbf{I}, \mathbf{G}, \mathbf{O}, \mathbf{R}$ are proposed which mimics the architecture of modern computer in storing and processing information.

- Input Memory Representation.** This layer converts the patient representations to the input representation. We pass all the patient representations through a dense layer to obtain the input representations.

$$\mathbf{I}_f(\mathbf{k}) = \mathbf{W}_i \mathbf{I}(\mathbf{k}) + \mathbf{b}_i \quad (9)$$

- Generalization.** Typically generalization can be referred to as the process of updating memory representation for the memory bank. In our case, we use the patient representations to initialize the memory representation \mathbf{M}_d which is the combination of all the patient representations. We then apply an LSTM layer to update the memory via multiple iterations.

$$\mathbf{M}_d = \text{LSTM}(\mathbf{I}_f(\mathbf{1}), \dots, \mathbf{I}_f(\mathbf{k})) \quad (10)$$

- Output.** In this step, the final output memory representation is generated. We calculate the relevance between trial embedding $\mathbf{Q}_{\text{emb}}(\mathbf{l})$ and doctor embedding \mathbf{M}_d to obtain $A(k)$ as the attention vector over patient representations.

$$A(k) = \text{softmax}[\mathbf{Q}_{\text{emb}}(\mathbf{l})^T \mathbf{M}_d] \quad (11)$$

- Response.** In this step, we obtain the final Doc_{emb} using the patient embeddings and attention weights over the patients.

$$\text{Doc}_{\text{emb}} = \sum A(k) \mathbf{I}_f(\mathbf{k}) \quad (12)$$

We use the doctor representation which is composed of patients and the clinical trial representation to obtain a final context vector. Besides dynamic doctor embedding Doc_{emb} , we also include static information about doctors in the final embedding such as their educational history, length of practice, length of practice into the feature vector. The resulting final embedding vectors are then fed into a fully connected layer and passed through a softmax to obtain class labels.

$$\mathbf{Y} = \text{Softmax}([\text{Doc}_{\text{emb}}; \mathbf{Q}_{\text{emb}}(\mathbf{l}); \text{Doc}_{\text{static}}]) \quad (13)$$

where the input to Softmax are concatenation of dynamic doctor embedding Doc_{emb} , trial query embedding $\mathbf{Q}_{\text{emb}}(\mathbf{l})$ and static doctor embedding $\text{Doc}_{\text{static}}$.

The enrollment rate category is obtained by binning the continuous enrollment rate. We divide the continuous enrollment scores into five discrete classes ranging at $0 \sim 0.2$, $0.2 \sim 0.4$, $0.4 \sim 0.6$, $0.6 \sim 0.8$, $0.8 \sim 1.0$. The 5 enrollment categories are used labels for classification.

Training and Inference

During training, We train our models by minimizing the cross entropy loss to optimize \mathbf{W}_{emb} , \mathbf{W}_c , \mathbf{W}_i , weight matrices of Bi-LSTM. We denote the network parameters by θ_c which is updated by optimizing for the loss function.

$$\mathcal{L} = -\frac{1}{N} \sum (y_i \log(\hat{y}_i) + (1 - y_i)^T \log(1 - \hat{y}_i)) \quad (14)$$

In the inference phase, we use calibrated threshold value for obtaining predicted labels from the predicted probability values where we obtain the probability values from the final layer of the network as mentioned in Eq. 13. Our Doctor2Vec model is summarized in Algorithm 1.

Algorithm 1: Model Training for Doctor2Vec

Input: Training dataset, input (\mathbf{C} , \mathbf{Q}) and target \mathbf{Y} ;
epochs N_{epoch}

Output: Trained model for enrollment rate prediction
with parameter θ_c

```
1 Initialization;
2 for  $i = 1, \dots, N_{epoch}$  do
3   foreach mini-batch in the training set do
4     Obtain  $\mathbf{h}_k(\mathbf{m})$  using MLP in Eq. 1 ;
5     Compute  $\mathbf{I}(k)$  using BiLSTM and Attn. Eq. 5;
6     Compute  $\mathbf{Q}_{emb}$  using combination of MLP and
       BERT by trial embedding;
7      $\mathbf{I}_f(k)$  is obtained from  $\mathbf{I}(k)$  ;
8      $\mathbf{A}(k)$  is generated from inner product of  $\mathbf{Q}_{emb}$ 
       and  $\mathbf{I}_f(k)$ 
9     Compute Doctor representation  $\mathbf{Doc}_{emb}$  in Eq.
       12;
10    Combine  $\mathbf{Doc}_{emb}$ ,  $\mathbf{Q}_{emb}(l)\mathbf{Doc}_{static}$  for the
       final prediction Eq. 13;
11    Calculate prediction loss  $\mathcal{L}$  using Eq. 14;;
12    Update parameters according to the gradient of
        $\mathcal{L}$ ;
13  end
14 end
```

Experiment

We designed experiments to answer the following questions.

Q1: Does Doctor2Vec have better performance in predicting clinical trial enrollment to support site selection?

Q2: Can Doctor2Vec embedding perform in transfer learning setting for trials across countries or across diseases?

Implementation We implemented Doctor2Vec² with PyTorch 1.0 (Paszke et al. 2017). For training the model, we used Adam (Kingma and Ba 2014a) with the mini-batch of 128 samples. The training was performed on a machine equipped with an Ubuntu 16.04 with 128GB memory and Nvidia Tesla P100 GPU.

²Code: <https://github.com/sidsearch/Doctor2vec>

Data Source We obtained patient and trial information from the following three data sources.

1. We extracted trial data from IQVIA’s real-world patient and clinical trial database, which can be accessed by request³. It contains 2609 clinical trials formed during 2014 and 2019. This dataset includes 25894 doctors across 28 countries. It includes both unstructured eligibility criteria and categorical features including the geographic location of the trial, hospital system, primary, secondary therapeutic areas, drug names, etc. The data also includes outcome measures such as the trial enrollment rate. In ground truth, the distribution of the enrollment categories are 12%, 33%, 37%, 12%, 6% respectively for $0 \sim 0.2$, $0.2 \sim 0.4$, $0.4 \sim 0.6$, $0.6 \sim 0.8$, $0.8 \sim 1.0$ bins of enrollment score.
2. We also obtained real world patient claims dataset from Database 1. This dataset contains a longitudinal treatment history from 430,239 patients over 7 years. In addition to medical codes about diagnosis, procedure, medication, it also includes information about doctors such as specialty, education, hospital location, geographical location.
3. We also extract clinical trial descriptions from publicly available clinical trial database clinicaltrials.gov. We match the trial information with our Database 1 on NCT ID which is a universal clinical trial ID.

Enrollment Rate Enrollment rate for each investigator is defined as

$$\text{Enrollment Rate} = \frac{\# \text{ sub. randomized} - \# \text{ sub. discontinued}}{\text{enrollment window}}$$

After obtaining the enrollment numbers, we perform a min-max normalization step to obtain normalized enrollment rate which is between 0 and 1. In this normalization step, we only consider the investigators associated with each clinical trial.

Table 2: Data Statistics

# of clinical trials	2,609
# of doctors	25,894
# of doctor-trial pair(samples)	102,487
# of patients	430,239
Avg # of Dx codes per visit	4.23
Max # of Dx codes per visit	56
Avg # of Procedure codes per visit	1.23
Max # of Procedure codes per visit	18
Avg # of Med codes per visit	9.36

Baselines We consider the following baselines.

1. Median Enrollment (Median). Current industry standard that considers the median enrollment rate for each therapeutic area as estimated rate for all trials in that area.
2. Logistic Regression (LR). We combine the medication, diagnosis and procedure codes along with the clinical trial information to create feature vectors, and then apply LR to predict the enrollment rate category.

³<https://www.iqvia.com/insights/the-iqvia-institute>

3. Random Forest (RF) (Breiman 2001). We combine the medication, diagnosis and procedure codes along with the clinical trial information to create feature vectors and then pass it to RF to predict the enrollment rate category.
4. AdaBoost (Schapire 1999). We combine the medication, diagnosis and procedure codes along with the clinical trial information to create feature vectors and then apply AdaBoost classifier to predict the enrollment rate categories.
5. Multi-layer Perceptron (MLP). We use MLP to process doctor features. In this case, we obtain the doctor features by converting all the visit vectors associated with a doctor to a count vector of different diagnosis, medication, procedure codes. We convert categorical information of clinical trials to multi-hot vectors and obtain TF-IDF features from text information of clinical trials.
6. Long Short-Term Memory Networks (LSTM) (Hochreiter and Schmidhuber 1997). We process all the temporally ordered visit vectors associated with a doctor using an LSTM. The embedding obtained from LSTM is concatenated with embedding obtained from categorical and text information of clinical trials to predict enrollment rate.
7. DeepMatch (Gligorijevic et al. 2019) In this model, the features for the doctors are obtained by collecting the top 50 most frequent medical codes and passed through an MLP layer to obtain an embedding vector. This embedding is concatenated with embedding obtained from categorical and text information of clinical trials via MLP and TF-IDF to finally predict enrollment rate.

Evaluation Metrics To evaluate the performance of enrollment prediction, We used PR-AUC as the metric for the classification task, and the coefficient of determination (R^2) score for the regression task. Details of the metrics are provided in appendix.

Experiment Setup and Evaluation Strategies We split our data into train, test, validation split with 70:20:10 ratio. We also ensured that the clinical trials are unique and no overlap in train, test, validation split. We used Adam (Kingma and Ba 2014b) optimizer at learning rate 0.001 with learning rate decay. We fix the best model on evaluation set within 200 epochs and report the performance in test set. Details about reproducibility including hyperparameters are provided in Appendix.

Q1: Doctor2Vec achieved the best predictive performance in clinical trial enrollment prediction

We conducted experiments for both classification (e.g., predict enrollment rate category) and regression (e.g., predict actual rate) tasks. Results are provided in Table 3. From the results, we observe that Doctor2Vec achieved the best performance in both settings.

For category classification, Doctor2Vec has 8.7% relative improvement in PR-AUC over the best baseline LSTM. Among the baselines, the Median method performs the

Table 3: Doctor2Vec achieves the best performance on both metrics in predicting actual enrollment rate (regression task) and rate categories (classification task) compared to state-of-the-art baselines. Results of ten independent runs.

	PR-AUC	R^2 Score
Median	0.571 \pm 0.014	0.54 \pm 0.072
LR	0.672 \pm 0.041	0.314 \pm 0.082
RF	0.731 \pm 0.034	0.618 \pm 0.034
AdaBoost	0.747 \pm 0.002	0.684 \pm 0.146
MLP	0.761 \pm 0.019	0.762 \pm 0.049
LSTM	0.792 \pm 0.034	0.780 \pm 0.621
DeepMatch	0.735 \pm 0.068	0.821 \pm 0.073
Doctor2Vec	0.861 \pm 0.021	0.841 \pm 0.072

worst, indicating the population level information is not accurate enough for each individual trial. Tree based models such as RF and Adaboost performs better than Median enrollment and LR, which can be attributed to the large number of features they leverage as well as their ability of distilling complex features. MLP performs better than tree based models due to having adequate number of layers for better capturing information and ability control overfitting. The LSTM network further improves over MLP since it is able to extract the temporal information present from the visits of patients. Compared with these approaches, the DeepMatch models achieved much lower PR-AUC since the model leverages the 50 most frequent codes for medical concept embedding, thus missing many important information of the doctors.

In actual rate prediction task, Doctor2Vec gains 2.4% relative improvement in R^2 over best baseline DeepMatch. As for the baselines, the LR model performs the worst, indicating linear models cannot capture the complex and temporal information in the data. Median enrollment is worse than most baselines but better than LR since for some more common diseases median enrollment can be a good predictor. Again, MLP and LSTM work better than tree-based models due to they can better capture complex features. DeepMatch in the regression settings tends to perform better than MLP and LSTM, which can be attributed to the majority of actual scores being in the range of [0.4-0.65] leading to improved performance.

Q2: Doctor2Vec can perform well in trial recruitment prediction even across countries and across diseases

One major challenge for clinical trial recruitment is when conducting trials in a less populated country or a country that is newly explored, or building a trial for a rare disease, the recruitment rate is often hard to estimate since there is not enough historical data to refer to. In this section we design two experiments to explore whether the embedding learned by Doctor2Vec will be useful in order to benefit the aforementioned data insufficiency settings.

1. Trained on United States data and transfer to a less populated/newly explored country;

2. Trained on common diseases and transfer to rare/low prevalence diseases.

For the first experiment, we trained Doctor2Vec on 1443 clinical trials in the United states during the time 2014-2019 and test on 47 clinical trials in South Africa during the time 2014-2019. We perform the same model transfer for all baselines. Results are provided in Table 4.

For the second experiment, we test the model on 38 clinical trials for drugs about idiopathic pulmonary fibrosis (IPF, a rare lung disease) and inflammatory bowel disease (IBD, a low prevalence chronic inflammatory bowel disease). The model was trained on 2569 clinical trials from the rest of the available diseases. We perform the same model transfer for all baselines. Results are provided in Table 5.

Table 4: Doctor2Vec achieves the best performance when we transfer the model trained on US data to predict trial enrollment in South Africa.

	PR-AUC	R^2 Score
Median	0.524 ± 0.032	0.420 ± 0.039
LR	0.601 ± 0.023	0.279 ± 0.014
RF	0.661 ± 0.038	0.552 ± 0.048
AdaBoost	0.672 ± 0.01	0.581 ± 0.039
LSTM	0.758 ± 0.013	0.721 ± 0.025
DeepMatch	0.703 ± 0.087	0.756 ± 0.031
Doctor2Vec	0.862 ± 0.003	0.817 ± 0.025

Table 5: Doctor2Vec achieves the best performance when we transfer the model trained on common disease trials to rare and low prevalence disease trials.

	PR-AUC	R^2 Score
Median	0.413 ± 0.013	0.387 ± 0.001
LR	0.521 ± 0.021	0.225 ± 0.028
RF	0.610 ± 0.019	0.517 ± 0.032
AdaBoost	0.623 ± 0.002	0.548 ± 0.046
LSTM	0.725 ± 0.002	0.623 ± 0.038
DeepMatch	0.638 ± 0.021	0.678 ± 0.049
Doctor2Vec	0.784 ± 0.032	0.716 ± 0.014

For both settings, Doctor2Vec performs much better than state-of-the-art baselines. For the country transfer, Doctor2Vec achieved 13.7% relative improvement in PR-AUC over best baseline LSTM and 8.1% relative improvement in R^2 over best baseline DeepMatch. While for embedding transfer to rare disease trials, Doctor2Vec achieved 8.1% relative improvement in PR-AUC over the best baseline LSTM and 5.6% relative improvement in R^2 over best baseline DeepMatch.

For country transfer, we also examine the R^2 scores. Based on the R^2 values, the DeepMatch model and LSTM model accounts for 69.2% and 67.3% of the variance in the data, respectively. While our Doctor2Vec accounts for 83.6% of the variance. This shows our model and the prediction fit more to the real observation.

Case Study

We present case studies to demonstrate the effectiveness of the proposed Doctor2Vec model.

Phase I trial for Gemcitabine plus Cisplatin This phase I trial is a combination cancer therapy. A doctor in the US who has worked in internal medicine during the past 3 years has run the trial. The actual enrollment rate is 0.72. The rate estimation provided by the best baseline LSTM is 0.57. While the estimated rate from Doctor2Vec is 0.69, which is much closer to the ground truth. The reason for Doctor2Vec to perform more accurately is the internal medicine doctor has a broader coverage of diseases. Baseline models consider all these diseases that the doctor treated when measuring the match between the doctor and the trial. While Doctor2Vec was able to focus more on the patients who had cancer diagnosis instead of all patients which leads to improved prediction.

Phase II trial for Alzheimer’s Disase This phase II trial is about an amyloid drug for treating Alzheimer’s patients. A doctor in the US who has treated cancer during the past 4 years runs the trial and has a trial enrollment rate at 0.62. The estimated rate from the best baseline LSTM model is 0.45. Doctor2Vec predicts the enrollment rate will be 0.58, which is much closer to the ground truth. For this case, Doctor2Vec is more accurate because Doctor2Vec is able to learn better doctor representations for this doctor by focusing on the neurological disease patients compared to other disease type patients seen by the doctor.

Conclusion

In this work, we proposed Doctor2Vec, a doctor representation learning based on both patient representations from longitudinal patient EHR data and trial embedding from the multimodal trial description. Doctor2Vec leverages a dynamic memory network where the representations of patients seen by the doctor are stored as memory while trial embedding serves as queries for retrieving the memory. Evaluated on real world patient and trial data, we demonstrated via trial enrollment prediction tasks that Doctor2Vec can learn accurate doctor embeddings and greatly outperform state-of-the-art baselines. We also show by additional experiments that the Doctor2Vec embedding can also be transferred to benefit the data insufficient setting (e.g., model transfer to less populated/newly explored country or from common disease to rare disease) that is highly valuable yet extremely challenging for clinical trials.

Acknowledgement

This work was in part supported by the National Science Foundation award IIS-1418511, CCF-1533768 and IIS-1838042, the National Institute of Health award NIH R01 1R01NS107291-01 and R56HL138415.

References

- Baytas, I. M.; Xiao, C.; Zhang, X.; Wang, F.; Jain, A. K.; and Zhou, J. 2017. Patient subtyping via time-aware lstm networks. In *KDD*.
- Breiman, L. 2001. Random forests. *Mach. Learn.* 45(1):5–32.
- Che, C.; Xiao, C.; Liang, J.; Jin, B.; Zho, J.; and Wang, F. 2017. An rnn architecture with dynamic temporal matching for personalized predictions of parkinson’s disease. In *SDM*.
- Choi, E.; Bahadori, M. T.; Sun, J.; Kulas, J.; Schuetz, A.; and Stewart, W. 2016. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems*.
- Choi, E.; Xiao, C.; Stewart, W.; and Sun, J. 2018. Mime: Multilevel medical embedding of electronic health records for predictive healthcare. In *Advances in Neural Information Processing Systems*, 4547–4557.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fu, T.; Gao, T.; Xiao, C.; Ma, T.; and Sun, J. 2019. Pearl: Prototype learning via rule learning. In *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics, BCB ’19*, 223–232. New York, NY, USA: ACM.
- Gligorijevic, J.; Gligorijevic, D.; Pavlovski, M.; Milkovits, E.; Glass, L.; Grier, K.; Vankireddy, P.; and Obradovic, Z. 2019. Optimizing clinical trials recruitment via deep learning. *Journal of the American Medical Informatics Association*.
- Graves, A.; Wayne, G.; Reynolds, M.; Harley, T.; Danihelka, I.; Grabska-Barwińska, A.; Colmenarejo, S. G.; Grefenstette, E.; Ramalho, T.; Agapiou, J.; et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature* 538(7626):471.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9.
- Hurtado-Chong, A.; Joeris, A.; Hess, D.; and Blauth, M. 2017. Improving site selection in clinical studies: a standardised, objective, multistep method and first experience results. *BMJ open* 7(7):e014796.
- Johnson, A. E.; Stone, D. J.; Celi, L. A.; and Pollard, T. J. 2017. The mimic code repository: enabling reproducibility in critical care research. *JAMIA*.
- Kingma, D. P., and Ba, J. 2014a. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980.
- Kingma, D. P., and Ba, J. 2014b. Adam: A method for stochastic optimization. *ICLR*.
- Kumar, A.; Irsoy, O.; Ondruska, P.; Iyyer, M.; Bradbury, J.; Gulrajani, I.; Zhong, V.; Paulus, R.; and Socher, R. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning*, 1378–1387.
- Le, H.; Tran, T.; and Venkatesh, S. 2018. Dual memory neural computer for asynchronous two-view sequential learning. In *Proceedings of the 24rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1637–1645. ACM.
- Miller, A.; Fisch, A.; Dodge, J.; Karimi, A.-H.; Bordes, A.; and Weston, J. 2016. Key-value memory networks for directly reading documents. In *Empirical Methods in Natural Language Processing*, 1400–1409.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch.
- Potter, J. S.; Donovan, D. M.; Weiss, R. D.; Gardin, J.; Lindblad, R.; Wakim, P.; and Dodd, D. 2011. Site selection in community-based clinical trials for substance use disorders: Strategies for effective site selection. *The American journal of drug and alcohol abuse* 37(5):400–407.
- Schapire, R. E. 1999. A brief introduction to boosting. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI’99*, 1401–1406. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Shang, J.; Ma, T.; Xiao, C.; and Sun, J. 2019a. Pre-training of graph augmented transformers for medication recommendation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI’19*, 5953–5959. AAAI Press.
- Shang, J.; Xiao, C.; Ma, T.; Li, H.; and Sun, J. 2019b. Gamenet: Graph augmented memory networks for recommending medication combination. *AAAI*.
- Sukhbaatar, S.; Weston, J.; Fergus, R.; et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*, 2440–2448.
- van den Bor, R. M.; Grobbee, D. E.; Oosterman, B. J.; Vaessen, P. W.; and Roes, K. C. 2017. Predicting enrollment performance of investigational centers in phase iii multicenter clinical trials. *Contemporary clinical trials communications* 7:208–216.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Weston, J.; Chopra, S.; and Bordes, A. 2015. Memory networks. In *ICLR*.
- Xiao, C.; Ma, T.; Dieng, A.; Blei, D. M.; and Wang, F. 2018. Readmission prediction via deep contextual embedding of clinical concepts. *PLOS ONE* 13(4):1–15.
- Xiao, C.; Choi, E.; and Sun, J. 2018. Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review. *JAMIA*.
- Xu, X.; Fu, Y.; Xiong, H.; Jin, B.; Li, X.; Hu, S.; and Yin, M. 2018. Dr. right!: Embedding-based adaptively-weighted mixture multi-classification model for finding right doctors with healthcare experience data. In *ICDM*.