

Integrative Analysis of Patient Health Records and Neuroimages via Memory-based Graph Convolutional Network

Xi Zhang, Jingyuan Chou and Fei Wang

Department of Healthcare Policy and Research, Weill Cornell Medical College, Cornell University
sheryl.zhangxi@gmail.com, jic2015@med.cornell.edu, few03@med.cornell.edu

ABSTRACT

With the arrival of big data era, more and more data are becoming readily available in various real world applications and those data are usually highly heterogeneous. Taking computational medicine as an example, we have both Electronic Health Records (EHR) and medical images for each patient. For complicated diseases such as Parkinson's and Alzheimer's, both EHR and neuroimaging information are very important for disease understanding because they contain complementary aspects of the disease. However, EHR and neuroimage are completely different. So far the existing research has been mainly focusing on one of them. In this paper, we proposed a framework, Memory-Based Graph Convolution Network (MemGCN), to perform integrative analysis with such multi-modal data. Specifically, GCN is used to extract useful information from the patients' neuroimages. The information contained in the patient EHRs before the acquisition of each brain image is captured by a memory network because of its sequential nature. The information contained in each brain image is combined with the information read out from the memory network to infer the disease state at the image acquisition timestamp. To further enhance the analytical power of MemGCN, we also designed a multi-hop strategy that allows multiple reading and updating on the memory can be performed at each iteration. We conduct experiments using the patient data from the Parkinson's Progression Markers Initiative (PPMI) with the task of classification of Parkinson's Disease (PD) cases versus controls. We demonstrate that superior classification performance can be achieved with our proposed framework, comparing with existing approaches involving a single type of data.

KEYWORDS

Multimodal Deep Learning, Graph Convolutional Network, Memory Network, Patient Health Records, Neuroimages

1 INTRODUCTION

With the arrival of big data era, more and more data are becoming readily available in various real world applications. Those data are like gold mines and data mining technologies are like tools that can dig the gold out from those mines. Taking medicine as an example, we have a large amount of medical data of different types nowadays, from molecular to cellular to clinical and even environmental. As

has been envisioned in [15], one key aspect of precision medicine, which aims at recommending the right treatment to the right patient at the right time, is to integrate those multi-scale data from different sources to obtain a comprehensive understanding of a health condition.

Many data mining approaches have been proposed for analyzing medical data in recent years. For example, Ghassemi *et al.* [19] modeled the mortality risk in intensive care unit with latent variable models. Caruana *et al.* [11] utilized generalized additive model to predict the risk of pneumonia and hospital readmission. Zhou *et al.* [58] developed a matrix factorization approach for predictive modeling of the disease onset risk based on patients' Electronic Health Records (EHR) data. Tensor modeling techniques have also been leveraged in electronic phenotyping [25, 26, 52] and clinical natural language processing [35]. More recently, deep learning has emerged as a powerful data mining approach that can disentangle the complex interactions among data features and achieve superior performance. Because of the complex nature of medical problems, researchers have also been exploring the applicability of deep learning models in helping with medical problems using medical images [17, 22], EHRs [45, 58], physiological signals [46, 48], etc., and obtained promising results.

Despite the initial success, so far most of the existing works on data mining for medicine have been focusing on one single type of data (e.g., images or EHRs). However, typically different data sources contain complementary information about the patients from different aspects. For example, concerning neurological diseases, we can get general clinical information of patients, such as diagnosis, medication, lab, etc., from EHRs; while we can obtain specific biomarkers regarding white matter, gray matter, and the change of different Regions-of-Interest (ROI), from brain images. Integrative analysis of both EHR and neuroimages can help us understand the disease in a better and more comprehensive way. In reality, such integrative analysis is challenging because of the following reasons.

- *Heterogeneity.* The nature of patient EHR and neuroimages are completely different: the EHR for each patient can be regarded as a temporal event sequence, where at each timestamp multiple medical events (e.g., diagnosis, medications, lab tests, etc.) can appear; while each neuroimage is essentially a collection of pixels. Therefore the ways to process these two types of data could be very different.
- *Sequentiality.* EHR data are sequential and a specific brain image is static. The brain status reflected in a certain brain image can be related to the EHR of the corresponding patient before the acquisition of the image. Effective integration of such heterogeneous information into a unified analytics pipeline is a challenging task.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

KDD'18 Deep Learning Day, August 2018, London, UK

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

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

With the above considerations, we proposed a novel Memory-based Graph Convolutional Network (MemGCN) to perform integrative analysis with both patient EHRs and neuroimages. As its name suggests, there are two major components in MemGCN.

- Graph Convolutional Network (GCN) [29]. GCN is a deep learning model that generalizes the Convolutional Neural Nets (CNN) [33] on regular lattices to irregular graphs. GCN has been proved to be very effective on extracting useful features from graphs [9, 29, 31].
- Memory Network [53]. Memory network is a new type of model that connects a regular learning process with a memory module, which is usually represented as a matrix that memorizes the historical status of the system. At each iteration some useful information is extracted from the memory to help the current inference while the same time the memory unit will be updated.

In our framework, the GCN module extracts features from the human brain networks constructed from the brain images. The longitudinal patient EHRs are stored in the memory network to encode the historical clinical information about the patient before the acquisition of the image. The information extracted from the memory network will be combined with the feature from GCN to discriminate PD cases and controls. We conduct experiments on real world data from the patients in the Parkinson's Progression Markers Initiative (PPMI) [37] and obtained superior performance comparing with conventional methodologies.

The rest of this paper is organized as follows. Section II presents the technical details of our framework. The experimental results are introduced in Section III, followed by the related work in Section IV and conclusions in Section V.

2 METHOD

2.1 Overview

As illustrated in Fig. 1, the proposed method MemGCN is a matching network that is designed for metric learning on not only brain images but also clinical records. The preprocessed brain connectivity graphs are transformed by graph convolutional networks into representations, while memory mechanism is in charge of iteratively (multiple hops) reading clinical sequences and choosing what to retrieve from memories in order to augment the representations learned by graph convolution. For the purpose of metric learning, inner product similarity and bilinear similarity are separately introduced in the matching layer. The output component is composed of a fully connected layer and a softmax for relationship classification of acquisition pairs. Accordingly, we present MemGCN, a matching network embeds multi-hop memory-augmented graph convolutions and can be trained in an end-to-end fashion with stochastic optimization.

2.2 Graph Convolution

The brain connectivity graph is characterized by defining its ROI nodes and the interactions among them. Since the graph-structured data are non-Euclidean, it is not straightforward to use a standard convolution that has impressive performances on grid. Hence, we

resort to geometric deep learning approaches [9, 41] to deal with the problem of learning features on brain connectivity network.

In general, let $\mathcal{G} = (\{1, \dots, n\}, \mathcal{E}, \mathbf{W})$ be an undirected weighted graph, where $\mathbf{W} = (w_{ij})$ is a symmetric adjacency matrix satisfying $w_{ij} > 0$ if $(i, j) \in \mathcal{E}$ and $w_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. According to spectral graph theory [14], the graph Laplacian matrix can be computed as $\Delta = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$, where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is the diagonal degree matrix with $d_{ii} = \sum_{j \neq i} w_{ij}$, and $\mathbf{I} \in \mathbb{R}^{n \times n}$ is the identity matrix. Note that Δ is a positive-semidefinite matrix and its eigendecomposition can be written as $\Delta = \Phi \Lambda \Phi^T$, where $\Phi = (\phi_0, \dots, \phi_n)$ are the orthonormal eigenvectors and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda(n))$ is the diagonal matrix of non-negative eigenvalues $0 = \lambda_0 \leq \dots \leq \lambda_n$.

In our scenario, the vertices of graph \mathcal{G} are corresponding to ROIs. Define a brain connectivity acquisition as an input signal $\mathbf{x} = (x_1, \dots, x_n)$, where $x_i \in \mathbb{R}^n$ is a feature vector associated with vertex i . The convolution operation is conducted on Fourier domain instead of the vertex domain. Consider two signals \mathbf{x} and \mathbf{g} , it can be proved that the following equation exists,

$$\begin{aligned} \mathbf{x} \star \mathbf{g} &= \Phi (\Phi^T \mathbf{x}) \odot (\Phi^T \mathbf{g}) = \Phi g_\theta(\Lambda) \Phi^T \mathbf{x} \\ &= \Phi \text{diag}(\hat{g}_1, \dots, \hat{g}_n) \hat{\mathbf{x}} \end{aligned} \quad (1)$$

where \odot is the element-wise Hadamard product and $\hat{\mathbf{x}} = \Phi^T \mathbf{x}$ defines the Graph Fourier Transform. The function $\hat{g}_\theta(\cdot)$ can be regarded as learnable spectral filters. Previous studies [10, 16, 23, 29] on geometric deep learning have proposed a variety of filter functions to achieve promising properties such as spatial localization and computational complexity. Chebyshev spectral convolution network (ChebNet) [16] is utilized in our model. Before introducing representation learning by ChebNet, we first give the details about how to construct a graph \mathcal{G} and build its edges \mathcal{E} with ROI vertices of the collection of brain image acquisitions.

Spatial Graph Construction

The brain connectivity graph can be represented as a square matrix $\mathbf{x} \in \mathbb{R}^{n \times n}$ with the numerical values indicating the connectivity strength of ROI pairs. However, the region coordinates of anatomical space can provide the crucial spatial relations between ROIs which have not been taken into account in conventional works of the domain [27]. Motivated by the work [31], which applied graph convolution on a functional Magnetic Resonance Imaging (fMRI) task, a spatial graph based on 3-dimensional coordinates is constructed for our model.

Suppose we have a population of M acquisitions, which are associated with a predefined number of ROIs and share a common coordinate system. The xyz-coordinates $\{(v_{i,m}^x, v_{i,m}^y, v_{i,m}^z)\}_{m=1}^M$ of region center are able to present the spatial location of the corresponding ROI i . Specifically, the global ROI coordinates are computed by the average aggregation $\bar{v}_i = \frac{1}{M} (\sum_m^M v_{i,m}^x, \sum_m^M v_{i,m}^y, \sum_m^M v_{i,m}^z)$, $\forall i \in (1, \dots, n)$. Thus, the edges \mathcal{E} can be constructed by a Gaussian function based on k -Nearest Neighbor similarity, which is

$$w_{ij} = \begin{cases} \exp(-\frac{\|\bar{v}_i - \bar{v}_j\|^2}{2\sigma^2}), & \text{if } i \in \mathcal{N}_j \text{ or } j \in \mathcal{N}_i \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

where w_{ij} denotes the edge weights between vertex i and vertex j , \mathcal{N}_i and \mathcal{N}_j denote the neighbors for i and j respectively. In practice, we set \mathcal{G} as a 10-Nearest Neighbor graph. Therefore, the spatial

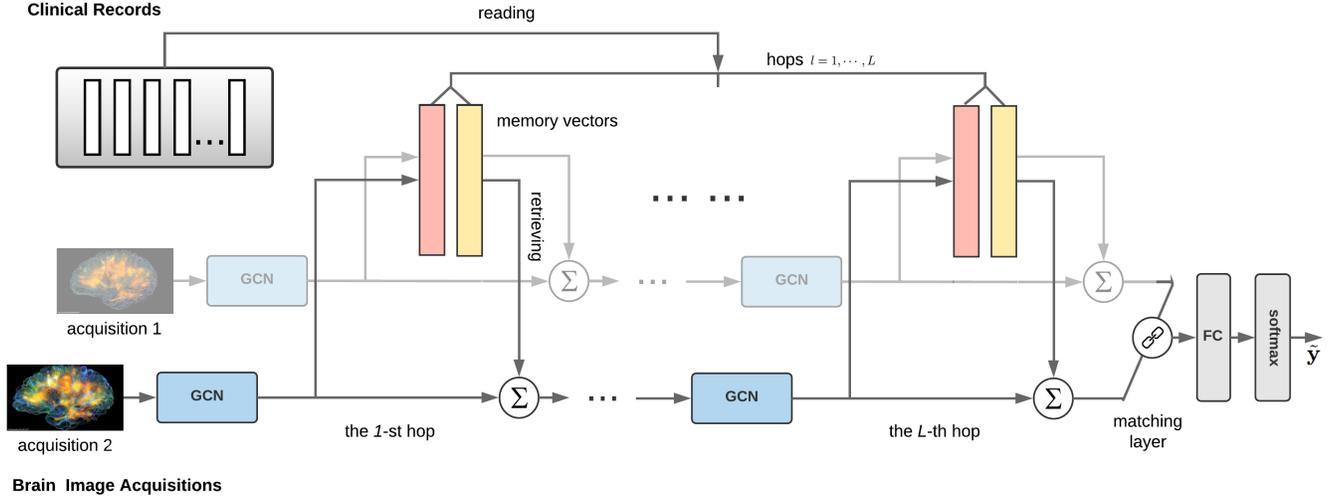


Figure 1: Memory-based Graph Convolutional Network for brain connectivity graphs with clinical records.

information of ROI is formulated into our model in terms of the graph structure.

ChebNet

With the constructed graph \mathcal{G} , its graph Laplacian matrix Δ can be obtained. Now our goal is to learn a high-level representation for each image acquisition by feeding its input signal \mathbf{x} as well as the shared Δ into the neural network. From the general sense, it can capture the local traits of each individual brain images and the global traits of the population of subjects.

To address the issues of localization and computational efficiency for convolution filters on graphs, ChebNet exploited a series of polynomial filters represented in the Chebyshev basis,

$$g_{\theta}(\Delta) = \sum_{p=0}^{r-1} \theta_p T_p(\tilde{\Delta}) = \sum_{p=0}^{r-1} \theta_p \Phi T_p(\tilde{\Delta}) \Phi^T \quad (3)$$

where $\tilde{\Delta} = 2\lambda_n^{-1}\Delta - \mathbf{I}$ is the rescaled Laplacian which leads to its eigenvalues $\tilde{\Lambda} = 2\lambda_n^{-1}\Lambda - \mathbf{I}$ in the interval $[-1, 1]$. θ is the r -dimensional vector Chebyshev coefficients parameterizing the filters. And $T_p(\tilde{\lambda}) = 2\lambda T_{p-1}(\tilde{\lambda}) - T_{p-2}(\tilde{\lambda})$ defines the Chebyshev polynomial in a recursive manner with $T_0(\tilde{\lambda}) = 1$ and $T_1(\tilde{\lambda}) = \tilde{\lambda}$.

To explicitly express filter learning of the graph convolution, without loss of generality, let k^l denote the index of feature map in layer l , the k^{l+1} -th feature map in its layer of sample m is given by

$$\mathbf{y}_{m,k^{l+1}} = \sum_{k^l=1}^{f_{in}} g_{\theta_{k^l,l+1}}(\mathbf{L})\mathbf{y}_{m,k^l} \in \mathbb{R}^n \quad (4)$$

yielding $f_{in} \times f_{out}$ vectors of trainable Chebyshev coefficients $\theta_{k^l,l+1} \in \mathbb{R}^r$. In detail, \mathbf{y}_{m,k^l} denotes the feature maps of the l -th layer. For the input layer, \mathbf{y}_{m,k^1} can be simply set as $\mathbf{x}_{m,i}$, $i = 1, \dots, n$, which is the i -th row vector of the $n \times n$ brain connectivity matrix. Given a sample $\mathbf{x} \in \mathbb{R}^{n \times n}$, its output of graph convolution can be collected into a feature matrix $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{f_{out}}) \in \mathbb{R}^{n \times f_{out}}$, where each row represents the learned high-level feature vector of an ROI.

2.3 Memory Augmentation

The key contribution of MemGCN is incorporating sequential records into the representation learning of brain connectivity in terms of memories. Our model is proposed based on Memory Networks [50, 53] which has a variety of successful uses in natural language processing tasks [8, 21, 39] including complex reasoning or question answering. When we define a memory, it could be viewed as an array of slots that can encode both long-term and short-term context. By pushing the clinical sequences into the memories, the continuous representations of this external information are processed with brain graphs together so that a more comprehensive diagnosis could be made. Inspired by the observation, the memory-augmented graph convolution are designed.

We start by introducing the MemGCN in the single hop operation, and then show the architecture of stacked hops in multiple steps. Concretely, the memory augmentation can be divided into two procedures: reading and retrieving (see Fig. 1).

Clinical Sequences Reading

Suppose there is a discrete input clinical sequences $s_j, j = 1, \dots, t$, where j is the index of a clinical record extracted from the certain timestamp. In memory network, it needs to be transformed as continuous vectors $\mathbf{z}_j, j = 1, \dots, t$ and stored into the memory. We use a fixed number of timestamps t to define the memory size. The dimension of the continuous space is denoted as d while the dimension of the original clinical features is denoted as D . To embed the sequential vectors s_1, \dots, s_t , a $d \times D$ embedding matrix \mathbf{A} is used. That is $\mathbf{z}_j = \mathbf{A}s_j$. The matrix $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_t)$ can be regarded as a new input memory representation.

Meanwhile, similar to the method in [50], an output memory to generate continuous vectors $\{\mathbf{e}_j\}$ is involved. The corresponding embeddings is obtained from $\mathbf{e}_j = \mathbf{B}s_j$, where \mathbf{B} also is a $d \times D$ embedding matrix. Different from other computational forms of attentive weights [3, 54], two memories in our model are maintained by the separate sequence reading procedures, which are responsible

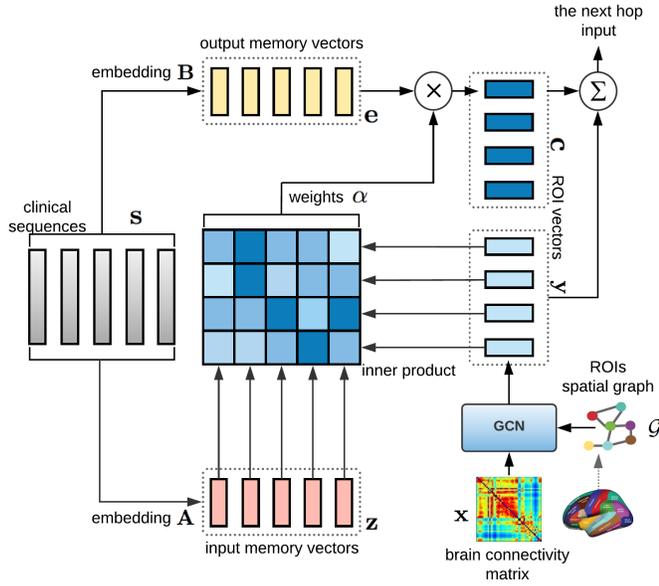


Figure 2: Illustration of memory augmented graph convolution in a single hop (the 1-st hop). See Section 2.3 for the details.

for memory access and integration respectively in the retrieving procedure.

Memory Representation Retrieving

To retrieve memory vectors from the embedding space, we firstly need to decide which vector to choose. Not all records in a sequence contribute equally when it comes to the representation learning for brain graphs. Hence, attentive weights are adopted here to make a soft combination of all memory vectors. Mathematically, the weights are computed by a softmax function on the inner product of the input memory vectors z_j and the learned ROI vectors y_i ,

$$\alpha_{ij} = \text{softmax}(y_i z_j) = \frac{\exp(y_i z_j)}{\sum_{j'=1}^t \exp(y_i z_{j'})} \quad (5)$$

Once the informative memory vectors are indicated by weights α_{ij} , the correspondence strength for attention are shown. As Fig. 2 illustrated, our attention is 2-dimensional that describes similarities between the representations generated from two modality sources. To make this feasible, we assume that both memory and ROI vectors are in the embedding space with same dimension. Next we represent the contextual information by the aggregation of weights and output memory vectors. Specifically,

$$\mathbf{c}_i = \sum_{j=1}^t \alpha_{ij} \mathbf{e}_j \quad (6)$$

where \mathbf{c}_i is a row vector of the context matrix \mathbf{c} , and is aware of a new representation for the ROI.

To integrate the context vectors with feature maps of GCN, element-wise sum is employed as $\hat{\mathbf{y}}_i = \mathbf{y}_i + \mathbf{c}_i$. The intuition of using the sum operator derives from the neural networks [18, 50],

in which the learned features in the next layer would benefit from both components in their networks.

The entire operations in a single hop are shown in Fig. 2, which is regarded as one layer (hop) of our model MemGCN. The output feature matrix of the single hop $\hat{\mathbf{y}} = (\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n)$ is fed into the next hop, and again as an input of the next GCN.

2.4 Multi-hop Layer

Basically, memory mechanism allow the network to read the input sequences multiple times to update the memory contents at each step and then make a final output. Compared to single step attention [3, 47, 54], contextual information from the memory is collected iteratively and cumulatively for feature maps learning. In particular, suppose there are L layer memories for L hop operations, the output feature map $\hat{\mathbf{y}}$ at the l -th hop can be rewritten as

$$\mathbf{y}^{l+1} = \mathbf{H}\mathbf{y}^l + \mathbf{c}^l, l = 1, \dots, L \quad (7)$$

where \mathbf{H} is a linear mapping and can be beneficial to the iteratively updating of \mathbf{y} . Similarly, the computational equations for weights and context vectors in Eq.(5) and Eq.(6) are rewritten as

$$\alpha_{ij}^l = \frac{\exp(y_i^l z_j^l)}{\sum_{j'=1}^t \exp(y_i^l z_{j'}^l)} \quad (8)$$

$$\mathbf{c}_i^l = \sum_{j=1}^t \alpha_{ij}^l \mathbf{e}_j^l \quad (9)$$

In addition, a layer-wise updating strategy [50] for input and output memory vectors at multiple hops are used, which is keeping the same embeddings as $\mathbf{A}^1 = \dots = \mathbf{A}^L$ and $\mathbf{B}^1 = \dots = \mathbf{B}^L$.

Notice that the contextual states \mathbf{c} of the first hop are determined by the two given modalities and then accumulated into the generation of the contextual states in the following hops. Consequently, the final output feature maps \mathbf{y}^L rely on the conditional contextual states $\mathbf{c}^1, \dots, \mathbf{c}^{L-1}$ as well as previous feature maps $\mathbf{y}^1, \dots, \mathbf{y}^{L-1}$, where \mathbf{y}^1 is directly generated from brain connectivity matrix \mathbf{x} through one layer graph convolution. The underlying rationale of the multi-hop design is that it is easier for the model to learn what have already been taken into account in previous hops and capture a fine-grained attendance from memories in the current hop.

2.5 Matching Layer

Metric learning for brain connectivity graphs with multiple layers normally involves several non-linearities so that the complex underlying data structure can be captured. To train such a neural network, a large amount of training data are necessary to prevent overfitting [5, 24]. Although large-scale labeled dataset are often limited in clinical practice, metric learning between sample pairs allow us increase the training data significantly because of the possible combination of two samples [30]. In our case, take a brain image acquisition as a sample, the goal of metric learning is to learn discriminative properties to distinguish whether the sample pairs in the same diagnosis class or not.

The basic hypothesis is that, if two samples share the same diagnosis result, the matching score between their high-level feature maps should be high. Here, two sorts of matching function are explored to calculate the similarities between pairs of acquisitions.

Inner Product Matching

Let \mathbf{x}_m and $\mathbf{x}_{m'}$ denote any pair of initial brain connectivity matrices, $\mathbf{y}_{m,i}^L$ and $\mathbf{y}_{m',i}^L$ denote their associated feature vectors learned from the L -th hops by MemGCN, where i is a vertex of ROI. The Euclidean distance computed in the matching layer is a vector with each dimension corresponding to each ROI, which is,

$$d_i(\mathbf{x}_m, \mathbf{x}_{m'}) = \|\mathbf{y}_{m,i}^L - \mathbf{y}_{m',i}^L\|_2, i = 1, \dots, n \quad (10)$$

Thus, $\mathbf{d} = (d_1, \dots, d_n)$ is the output of the matching layer. Instead of computing the distance directly, the feature maps are normalized along with dimension of hidden features and then a inner product is used to get a similarity vector,

$$sim_i(\mathbf{x}_m, \mathbf{x}_{m'}) = (\mathbf{y}_{m,i}^L)^T \mathbf{y}_{m',i}^L, i = 1, \dots, n. \quad (11)$$

where sim_i is the inner product similarity on i -th dimension, and it is equivalent to Euclidean distance if the vectors are normalized.

Bilinear Matching

Above matching function in Eq.(11) only considers the similarity of the corresponding ROI vectors of a given paired brain graphs. The similarities computed by different ROI are not modeled. To the aim, a simple bilinear matching function [7, 55] is used here. The matching score is defined as

$$sim_{i,j}(\mathbf{x}_m, \mathbf{x}_{m'}) = (\mathbf{y}_{m,i}^L)^T \mathbf{M} \mathbf{y}_{m',j}^L, i, j = 1, \dots, n. \quad (12)$$

where $sim_{i,j}$ is the similarity between ROI i and j based on bilinear matching. $\mathbf{M} \in \mathbb{R}^{f_{out}^L \times f_{out}^L}$ is a matrix parameterizing the matching between the paired feature maps. With the matching procedure in Eq.(12), the output of the matching layer is a matrix, with each element suggesting the strength of ROI connections. It is worth to note that if the parameter matrix \mathbf{M} is an identity matrix, the bilinear matching reduces to the inner product matching.

2.6 Model Training

As in MemGCN, our output layer models the probability of each sample pair is matching or non-matching. The similarity representation from matching layer is passed to a fully connected layer and a softmax layer for the eventual classification. For each pair, set the output of fully connected layer is a feature vector \mathbf{r} . We calculate the probability distribution over the binary classes by

$$\mathbf{p} = \text{softmax}(\mathbf{w}_c^T \mathbf{r}) \quad (13)$$

where $\mathbf{w}_c \in \mathbb{R}^2$ is a trainable parameter.

We train our model using a regularized cross-entropy loss function. Let $\mathcal{X} = \{(\mathbf{x}_m, \mathbf{x}_{m'})\}^N$ be the training set of N acquisition pairs. N is the number of total pairwise combination of brain graphs. The number of acquisitions M is much smaller than N . The loss function we minimize is

$$\mathcal{L} = \sum_{m,m'}^N \tilde{y}_{m,m'} \log p_{m,m'} + (1 - \tilde{y}_{m,m'}) \log(1 - p_{m,m'}) + \gamma \|\Theta\|_2 \quad (14)$$

where $\tilde{y}_{m,m'}$ denotes the label for sample pair $(\mathbf{x}_m, \mathbf{x}_{m'})$, Θ is the collection of trainable parameters. The MemGCN is trained on machines with NVIDIA TESLA V100 GPUs by using Adam optimizer [28] with mini-batch.

3 EXPERIMENTS

3.1 Dataset

The data we used to evaluate MemGCN are obtained from the Parkinson Progression Marker Initiative (PPMI) [37] study. PPMI is an ongoing PD study that has meticulously collected various potential PD progression markers that have been conducted for more than six years. Neuroimages and EHRs are considered as two modalities in this work.

To obtain brain connectivity graphs, a series of preprocessing procedures are conducted. For the correction of head motion and eddy current distortions, FSL eddy-correct tool is used to align the raw data to the b0 image. Also, the gradient table is corrected accordingly. To remove the non-brain tissue from the diffusion MRI, the Brain Extraction Tool (BET) from FSL [49] is used. To correct for echo-planar induced (EPI) susceptibility artifacts, which can cause distortions at tissue-fluid interfaces, skull-stripped b0 images are linearly aligned and then elastically registered to their respective preprocessed structural MRI using the Advanced Normalization Tools (ANTs¹) with SyN nonlinear registration algorithm [2]. The resulting 3D deformation fields are then applied to the remaining diffusion-weighted volumes to generate full preprocessed diffusion MRI dataset for the brain network reconstruction. In the meantime, ROIs are parcellated from T1-weighted structural MRI using Freesufer².

The connectivity graphs computed by three whole brain tractography methods [56] for are applied, which is a coverage of the tensor-based deterministic approach (Fiber Assignment by Continuous Tracking [42]), the Orientation Distribution Function (ODF)-based deterministic approach (the 2nd-order Runge-Kutta, RK2 [4]), as well as the probabilistic approach (Hough voting [1]). 84 ROIs are finally obtained. We define each the coordinates for ROIs using the mean coordinate for all voxels in the corresponding regions (see Spatial Graph Construction in Section 2.2 for the details). After preprocessing, we collect a dataset of 754 DTI acquisitions, where 596 of them are brain graphs of Parkinson's Disease (PD) patients and the rest 158 are from Healthy Control (HC) subjects. The spatial graph we constructed has 84 vertices and 527 edges, with each vertex is corresponding to a ROI.

Additionally, sequential EHR records are aligned with corresponding brain connectivity graphs. For each acquisition \mathbf{x} , a sequence of its associated input features (s_1, \dots, s_t) can be used for the external memories. Note that the sequences are chunked at the time points of neuroimaging acquisition, and we only use the subsequences before the time point to make a reasonable experimental design. Usually, the number of timestamps in sequences are different because subjects provide their medical records with distinct frequencies, we set the length of sequence as $t = 12$ according to the statistics of the PPMI study. Padding is utilized for those sequences with fewer timestamps. The specific clinical assessments we study here are motor (MDS-UPDRS Part II-III [20]) and non-motor (MDS-UPDRS Part I [20] and MoCA [44]) symptoms which are crucial for evaluating a disease course of PD. There are 79 discrete clinical features and 331 dimensions after binarization, then we have the original dimensions of clinical feature which is

¹<http://stnava.github.io/ANTs/>

²<https://surfer.nmr.mgh.harvard.edu>

Table 1: Results for classifying matching vs. non-matching brain graphs on the test sets of tensor-FACT, ODF-RK2, and Hough in terms of Accuracy and AUC metrics. Performances without and with extra modalities are shown. (Hop number $L = 3$ for MemGCNs)

Extra Modalities	Methods	tensor-FACT		ODF-RK2		Hough	
		Accuracy	AUC	Accuracy	AUC	Accuracy	AUC
None	Raw Edges	65.94±3.78	58.47±4.05	67.56±4.12	60.93±5.60	67.90±4.09	64.49±3.56
	PCA	69.19±3.13	64.10±2.10	68.38±2.50	60.93±2.63	66.28±4.60	63.46±3.52
	FCN	71.65±3.58	66.17±2.00	70.66±3.79	68.80±2.80	70.01±3.28	61.91±3.42
	FCN-2layer	84.22±2.76	82.36±2.87	82.31±2.68	82.53±4.74	84.27±2.63	81.77±3.74
	GCN-inner	93.69±2.15	92.67±4.94	93.23±2.63	93.04±5.26	92.80±2.51	93.90±5.48
	GCN-bilinear	93.89±1.76	94.77±6.08	94.00±2.65	94.32±5.72	93.34±2.26	93.35±5.14
Motor	AttGCN	93.34±2.68	93.37±5.97	94.48±2.22	94.13±5.33	93.42±5.37	94.44±0.29
	AttLstmGCN	94.97±2.51	93.19±4.32	94.21±1.20	94.11±4.52	93.62±1.03	94.72±4.51
	MemGCN-inner	95.28±1.77	96.29±6.35	94.82±1.94	95.61±6.01	94.51±2.23	95.66±6.11
	MemGCN-bilinear	94.80±1.74	95.49±5.35	95.17±2.54	96.05±6.30	95.20±1.40	96.11±6.22
Non-motor	AttGCN	93.33±2.09	93.35±4.30	93.80±1.17	93.23±4.14	93.61±1.61	93.51±4.88
	AttLstmGCN	93.07±1.58	93.51±4.43	93.09±2.30	93.54±5.00	93.65±1.29	93.21±4.33
	MemGCN-inner	93.92±2.36	94.22±5.58	93.99±2.18	94.30±5.37	94.05±1.29	94.39±5.62
	MemGCN-bilinear	94.11±2.24	94.51±5.61	94.50±2.20	95.12±5.93	94.40±1.26	94.92±5.75
Fusion	AttGCN	93.62±2.99	94.25±5.88	94.76±3.31	94.33±5.23	94.01±1.94	94.74±5.35
	AttLstmGCN	94.70±2.35	94.38±5.41	94.89±2.71	94.87±4.49	94.64±2.02	94.80±5.51
	MemGCN-inner	95.43±2.22	96.42±6.36	95.54±2.98	96.59±6.44	95.48±2.34	96.49±6.41
	MemGCN-bilinear	95.47±2.25	96.48±6.40	95.87±2.56	96.84±6.36	95.64±2.00	96.74±6.51

$D = 331$. At last, an imputation strategy Last Occurrence Carry Forward (LOCF) is adopted, since there are several missing entries at timestamps.

3.2 Experimental Setup

Implementation Details

To learn similarities between brain connectivity matrices, acquisitions in the same group (PD or HC) are labeled as matching pairs while those from different groups are labeled as non-matching pairs. Hence, we have 283, 881 pairs in total, with 189, 713 matching pairs and 94, 168 non-matching pairs.

We selected hyperparameter values through random search [6]. Batch size is 32. Initial learning rate is $5e-3$, and early stop is used once the model stops improving. The L2-regularization weight is $1e-2$. For each graph convolution operation, the order of Chebyshev polynomials and the feature map dimension are respectively set as $r = 30$ and $f_{out} = 32$. For the memory network, memory size and dimension of embedding are respectively set as $t = 12$ and $d = 32$. The MemGCN code are available at <https://github.com/sheryl-ai/MemGCN>.

Baselines

To test the performance of MemGCN, we report the empirical results of comparisons with a set of baselines. Here are the methods that classify brain graphs without any other modalities.

- **Raw Edges.** It is one simple approach that is to directly use the numerical values from the connectivity matrix to represent the brain network. The feature space is a 84×84 vector.
 - **PCA.** Principal Component Analysis (PCA) is used to reduce the data dimensionality. After forming a sample-by-feature input matrix, PCA is performed by keeping the first 100 principal components, which is an optimal setting in practice.
 - **FCN.** Fully Connected Network (FCN) is employed as feature extractor for raw connectivity matrix. The output dimension of the model is set as $f_{out} = 1024$. FCN is used on brain networks in [43].
 - **FCN-2layer.** 2-layers FCN has the same number of parameters in its first layer with FCN and reduces the dimension down to 64 in layer 2.
 - **GCN-inner.** Metric Learning for brain networks using GCN is first introduced in [31], where a global loss function is used to supervise pairwise similarities. The cross-entropy loss is adopted in our experiments to be consistent with other models.
 - **GCN-bilinear.** The bilinear matching layer proposed in Section 2.5 is added on the basis of GCN to conduct metric learning. It is a version of MemGCN-bilinear without memory.
- Furthermore, neural networks without memory mechanism that can embed the clinical data are built as baselines.
- **AttGCN.** Instead of using input and output memories on sequences, only one embedding matrix for the computation of attentive weights is incorporated with GCN via the sum operation.

Table 2: Comparisons for MemGCN by various setting of the number of hops and the matching methods.

# of hops	Matching	tensor-FACT		ODF-RK2		Hough	
		Accuracy	AUC	Accuracy	AUC	Accuracy	AUC
1	inner product	94.20±2.42	94.07±5.19	94.03±2.32	94.39±5.15	94.61±2.05	95.72±5.50
2	inner product	95.36±2.60	96.35±6.36	95.40±2.27	96.39±6.35	95.21±2.92	96.10±6.30
3	inner product	95.43±2.22	96.42±6.36	95.54±2.98	96.59±6.44	95.48±2.34	96.49±6.41
1	bilinear	94.68±2.04	95.32±5.98	93.88±2.22	94.17±5.49	94.37±1.89	95.17±4.95
2	bilinear	95.19±2.14	96.06±6.18	94.61±2.91	95.27±5.89	95.23±2.50	96.17±5.26
3	bilinear	95.47±2.25	96.48±6.40	95.87±2.56	96.84±6.36	95.64±2.00	96.74±6.51

- **AttLstmGCN.** A standard bi-directional LSTM with attention [3] is established for sequential EHR data and then its context states are combined with GCN feature maps.

Finally, two variants of our model are given.

- **MemGCN-inner.** The proposed MemGCN with the inner product matching layer.

- **MemGCN-bilinear.** The proposed MemGCN with the bilinear matching layer.

For a fair comparison, the reported models are built under the same pairwise matching architecture for metric learning. Inner product matching are employed in the baseline model if it is not stated as a bilinear version. 5-fold cross validation are conducted in all of our experiments.

3.3 Results

Matching vs. Non-matching Classification

Table 1 reports the performance for the binary classification task. Three sorts of memory augmentation are configured as motor sequences, non-motor sequences and a fusion of them. The metrics for evaluation are Accuracy and Area Under the Curve (AUC).

From the results we can observe that, the Raw Edges, and simple feature extraction approach such as PCA and 1-layer FCN cannot predict a reliable distance for sample pairs and correspondingly achieve a promising results on the matching classification task. More layers with extra non-linearities have a good influence on the fully connected networks to capture the complicated patterns from acquisitions. All the GCN based methods can largely improve both of Accuracy and AUC performance in three DTI sets generated by Tensor-FACT, ODF-RK2, and Hough tractography algorithms, which demonstrates the effectiveness of graph convolution on the brain connectivity graphs. Overall, the bilinear matching strategy is outperform the inner product matching strategy slightly on both GCN and MemGCN. The best AUC performance is 96.48, 96.84, and 96.74, which are accomplish by MemGCN-bilinear with fusion clinical sequences as the external modality. With attention mechanism, AttGCN and AttLstmGCN also perform well in the given circumstances. However, they cannot boost the results significantly compared to the vanilla GCN. The reason that MemGCN behaves better than them is probably separate memories for reading and retrieving are employed in a multi-hop network.

Table 3 shows the concrete effects of increasing the number of hops on inner product and bilinear matchings. The number of hops is tuned from 1 to 3. The results on Accuracy and AUC

metrics illustrate that our multi-hop framework indeed improves performance constantly.

Identical ROIs vs. Discriminative ROIs

The interpretability of MemGCN is investigated. As the representation learned in the inner product matching layer can be explained as pairwise similarities at 84 ROI dimensions, it describes the significance of each ROI in metric learning. Therefore, we compute the average similarities for all the PD-PD pairs and the average similarities for all the PD-HC pairs. ROIs with the highest scores in the PD group could be considered as the identical ROIs for PD, while those with the lowest scores in the PD versus HC group are regarded as the discriminative ROIs.

The interpretable results depends on memory augmentation of motor, non-motor, and the fusion data are presented in Table 3. While the whole functions of the human brain regions are still unclear, it is quite intriguing that MemGCN can locate some of the modality-related ROIs, which might be critical for PD study. For instance, The most identical ROI for PD with motor features as augmentation is the Thalamus with one of its major role as motor control. Also, lingual gyrus discovered by non-motor features is linked to processing vision, especially related to letters. On the other hand, MemGCN can help us to find which ROI is sufficiently discriminative to distinguish PD patients with healthy controls. Several important ROIs belongs to the current research of clinicians and domain experts are detected, i.e., Caudate and Putamen areas.

To show the representation generated from the bilinear matching layer, we draw the edges between ROIs with high similarities in Fig.3. Similar to Table 3, the most identical edges for PD group and the most discriminative edges between PD and HC groups are depicted. The interesting patterns we found might be deserved to the further exploration in clinical scenarios.

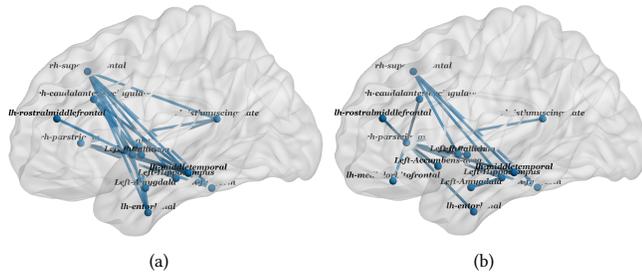
Longitudinal Alignment: Case Study

From Fig. 4(a) and 4(b) we observe that though the structures of three hops of memory layer are same, the values of the attention weights they learned are quite different in typical cases. The matrices we draw in terms of colormaps in Fig.4 indicate the attentive weights α for one PD case and one healthy control case. Here we abandon the first 2 padding dimensions of the shown cases and give 10 memory positions (rows in the matrices). The attendance of all the 84 ROI vertices are depicted (columns in the matrices). A darker color indicates where MemGCN is attending during the multi-hop updating for representations. Basically, given a specific

Table 3: The interpretability of the output representation of MemGCN’s inner product matching layer. Top-5 identical ROIs in PD group and discriminative ROIs between PD and HC groups are listed. Similarity scores are given.

	Motor		Non-motor		Fusion	
	ROI Name	Score	ROI Name	Score	ROI Name	Score
Identical ROIs (PD Group)	Right Thalamus Proper	0.9258	Rh Paracentral	0.8563	Rh Pars Opercularis	0.9344
	Lh Insula	0.9253	Rh Lingual	0.8180	Rh Lateral Occipital	0.8372
	Right Pallidum	0.9226	Right Pallidum	0.8091	Left Accumbens Area	0.7887
	Lh Rostral Middle Frontal	0.9210	Lh Parsorbitalis	0.6554	Rh Parahippocampal	0.7827
	Parahippocampal	0.9206	Left Thalamus Proper	0.6387	Rh Frontalpole	0.7742
Discriminative ROIs (PD vs. HC Group)	Right Putamen	-0.9134	Left Putamen	-0.7423	Right Thalamus Proper	-0.8960
	Right Accumbens Area	-0.9075	Lh Frontal Pole	-0.5754	Left Caudate	-0.8439
	Left Hippocampus	-0.9059	Lh Supramarginal	-0.5731	Lh Paracentral	-0.8227
	Right VentralDC	-0.9058	Lh Inferior Parietal	-0.5693	Lh Middle Temporal	-0.7865
	Left Caudate	-0.9014	Lh Paracentral	-0.4851	Lh Cuneus	-0.7528

* Lh and Rh are the abbreviations of Left Hemisphere and Right Hemisphere respectively.

**Figure 3: The connectivity patterns learned by the bilinear matching layer. (a) The top identical edges for PD group; (b) The top discriminative edges between PD and HC groups.**

case, which time point has more influences on his/her PD progression and which ROI is more important according to the clinical evidences can be analyzed through this longitudinal alignment between DTIs and EHRs.

In general, the first hop attention appears to be primarily concerned with identifying the salient interaction between time-aware sequences and ROIs’ feature maps. In this hop, the majority values are close to zero and only a few values are close to one, such that a sketch of key ROIs and timestamps are signified. The second and the third hops are then responsible for the fine-grained interactions that are relevant to optimizing the representation for the distance learning task.

Another important observation is that the PD case has different interaction patterns compared to the healthy control. At each hop, PD has a relatively narrow attention and fewer responses across memory positions. Consider the PD case shown in Fig. 4(a), longitudinal alignments occur at timestamps 2, 4, and 6 after 3-hop updating, meanwhile a series of ROIs might function on the disease progression. By the Desikan-Killiany Atlas, the darker ROI dimensions from 76 to 79 are Rh Insula, Right Thalamus Proper,

Right Caudate, and Right Putamen, respectively, which matches our expectation for the PD case.

4 RELATED WORK

We briefly review the existing research that is closely related to the framework proposed in this paper.

EHR Mining. In recent years many algorithms have been proposed to mine insights from patient EHRs. Initially those methods were static in the sense that they first construct patient vectors by aggregating their EHR with in a certain observational time window and then build learning approaches (e.g., predictive models and clustering methods) on top of those vectors [11, 19, 51]. Most of these methods are shallow except the DeepPatient work which applied AutoEncoder to further compress the patient vectors and obtain better representations [40]. Recently researchers have also been exploring CNN and RNN type of approaches to incorporate the temporal information in patient EHRs into the modeling process [12, 13, 36]. However, these methods compressed the patient EHRs to a vector before it was fed to the final model, which is not as flexible as the memory network we adopted.

GCN for Neuroimage Analysis. Many data mining approaches have been developed to perform neuroimage analysis in recent years [38], among which deep learning models are very popular because of their huge success in various computer vision problems [34]. Recently, Ktena *et al.* [32] propose to learn a metric from patients’ neuroimages on top of the features constructed using GCN (where the graph is basically the patients’ brain network constructed on the ROIs), which can discriminate the cases versus controls with autism. Zhang *et al.* [57] extend such approach to handle the multiple modalities of the brain networks (e.g., constructed from different tractography algorithms on DTI images). However, none of them incorporated any clinical records from the patients. Our work is the first step towards filling the gap.

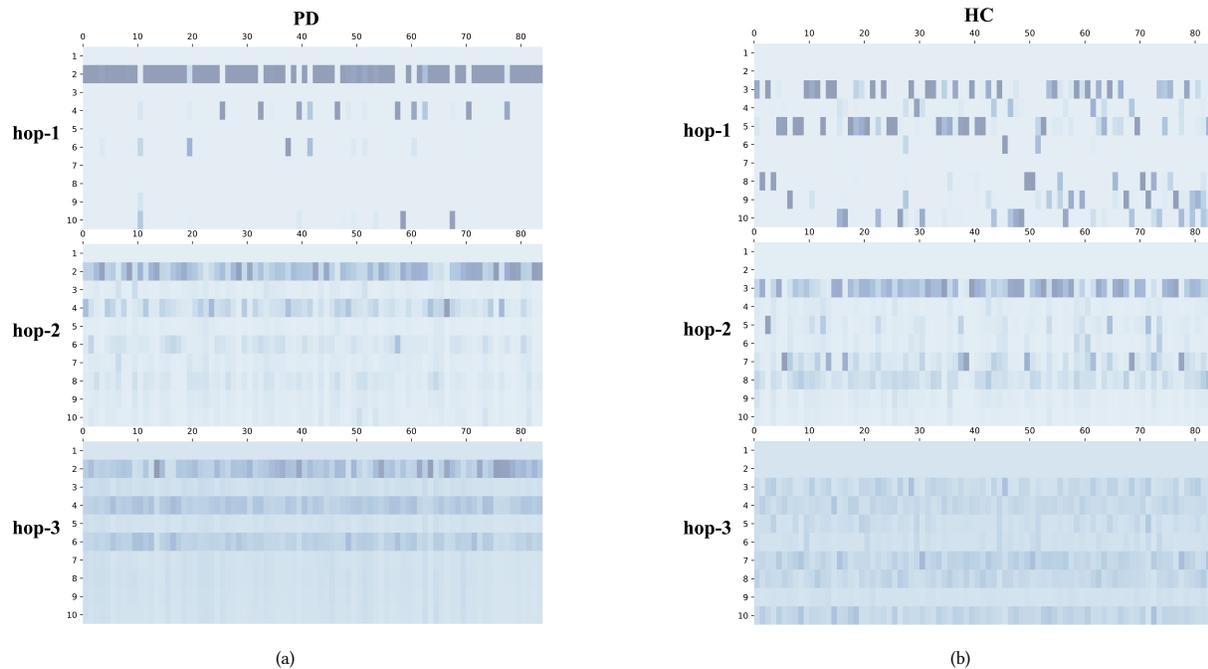


Figure 4: Visualizations of attention interaction matrices of MemGCN for one PD and one HC case during 3 memory hops. The rows and columns of the matrices respectively denote memory positions and ROI vertices. The darker color in the colormaps means a larger value which is close to 1, and the lighter color means a smaller value which is close to 0.

5 CONCLUSION

We propose a novel framework, memory-based graph convolution network (MemGCN), to perform integrative analysis of patient clinical records and neuroimages. On the one hand, our experiments on classification of Parkinson’s Disease case patients with healthy controls demonstrate the superiority of MemGCN over conventional approaches. On the other hand, the interpretable high-level representations extracted from the inner product or bilinear matching layers are capable of indicating group patterns of brain connectivity via ROI nodes or their edges for PD subjects and healthy controls.

Here we explored the operators of the graph convolution via ChebNet and the embedding via memory mechanism as feature extractors for neuroimages and patient health records respectively. The pairwise distance under metric learning setting in our framework makes a progress in modeling a small cohort data such as PPMI. An important future direction is to design deep architectures that can lower the amount of training data meanwhile learn meaningful representations. We are especially interested in continuing to develop more general end-to-end learnable models in the space of boosting system performance on small data.

REFERENCES

- [1] Iman Aganj, Christophe Lenglet, Neda Jahanshad, Essa Yacoub, Noam Harel, Paul M Thompson, and Guillermo Sapiro. 2011. A Hough transform global probabilistic approach to multiple-subject diffusion MRI tractography. *Medical image analysis* 15, 4 (2011), 414–425.
- [2] Brian B Avants, Charles L Epstein, Murray Grossman, and James C Gee. 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* 12, 1 (2008), 26–41.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- [4] Peter J Basser, Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. 2000. In vivo fiber tractography using DT-MRI data. *Magnetic resonance in medicine* 44, 4 (2000), 625–632.
- [5] Yoshua Bengio et al. 2009. Learning deep architectures for AI. *Foundations and trends® in Machine Learning* 2, 1 (2009), 1–127.
- [6] James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research* 13, Feb (2012), 281–305.
- [7] Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2014. A semantic matching energy function for learning with multi-relational data. *Machine Learning* 94, 2 (2014), 233–259.
- [8] Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. Large-scale simple question answering with memory networks. *arXiv preprint arXiv:1506.02075* (2015).
- [9] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* 34, 4 (2017), 18–42.
- [10] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann Lecun. 2014. Spectral networks and locally connected networks on graphs. In *International Conference on Learning Representations*.
- [11] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1721–1730.
- [12] Chao Che, Cao Xiao, Jian Liang, Bo Jin, Jiayu Zhu, and Fei Wang. 2017. An RNN Architecture with Dynamic Temporal Matching for Personalized Predictions of Parkinson’s Disease. In *Proceedings of the 2017 SIAM International Conference on Data Mining*. SIAM, 198–206.
- [13] Yu Cheng, Fei Wang, Ping Zhang, and Jianying Hu. 2016. Risk prediction with electronic health records: A deep learning approach. In *Proceedings of the 2016 SIAM International Conference on Data Mining*. SIAM, 432–440.
- [14] Fan RK Chung. 1997. *Spectral graph theory*. Number 92.
- [15] Francis S Collins and Harold Varmus. 2015. A new initiative on precision medicine. *New England Journal of Medicine* 372, 9 (2015), 793–795.
- [16] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*. 3844–3852.

- [17] Andre Esteva, Brett Kuperl, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542, 7639 (2017), 115.
- [18] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional Sequence to Sequence Learning. In *International Conference on Machine Learning*. 1243–1252.
- [19] Marzyeh Ghassemi, Tristan Naumann, Finale Doshi-Velez, Nicole Brimmer, Rohit Joshi, Anna Rumshisky, and Peter Szolovits. 2014. Unfolding physiological state: Mortality modelling in intensive care units. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 75–84.
- [20] Christopher G Goetz, Barbara C Tilley, Stephanie R Shaftman, Glenn T Stebbins, Stanley Fahn, Pablo Martinez-Martin, Werner Poewe, Cristina Sampaio, Matthew B Stern, Richard Dodel, et al. 2008. Movement Disorder Society-sponsored revision of the Unified Parkinson's Disease Rating Scale (MDS-UPDRS): Scale presentation and clinimetric testing results. *Movement disorders* 23, 15 (2008), 2129–2170.
- [21] Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural turing machines. *arXiv preprint arXiv:1410.5401* (2014).
- [22] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, et al. 2016. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama* 316, 22 (2016), 2402–2410.
- [23] Mikael Henaff, Joan Bruna, and Yann LeCun. 2015. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163* (2015).
- [24] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation* 18, 7 (2006), 1527–1554.
- [25] Joyce C Ho, Joydeep Ghosh, Steve R Steinhilber, Walter F Stewart, Joshua C Denny, Bradley A Malin, and Jimeng Sun. 2014. Limestone: High-throughput candidate phenotype generation via tensor factorization. *Journal of biomedical informatics* 52 (2014), 199–211.
- [26] Joyce C Ho, Joydeep Ghosh, and Jimeng Sun. 2014. Marble: high-throughput phenotyping from electronic health records via sparse nonnegative tensor factorization. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 115–124.
- [27] Jeremy Kawahara, Colin J Brown, Steven P Miller, Brian G Booth, Vann Chau, Ruth E Grunau, Jill G Zwicker, and Ghassan Hamarneh. 2017. BrainNetCNN: convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* 146 (2017), 1038–1049.
- [28] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [29] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [30] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, Vol. 2.
- [31] Sofia Ira Ktena, Sarah Parisot, Enzo Ferrante, Martin Rajchl, Matthew Lee, Ben Glocker, and Daniel Rueckert. 2017. Distance metric learning using graph convolutional networks: Application to functional brain networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 469–477.
- [32] Sofia Ira Ktena, Sarah Parisot, Enzo Ferrante, Martin Rajchl, Matthew Lee, Ben Glocker, and Daniel Rueckert. 2017. Metric learning with spectral graph convolutions on brain connectivity networks. *NeuroImage* (2017).
- [33] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436.
- [34] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen AWM van der Laak, Bram van Ginneken, and Clara I Sánchez. 2017. A survey on deep learning in medical image analysis. *Medical image analysis* 42 (2017), 60–88.
- [35] Yuan Luo, Yu Xin, Ephraim Hochberg, Rohit Joshi, Ozlem Uzuner, and Peter Szolovits. 2015. Subgraph augmented non-negative tensor factorization (SANTF) for modeling clinical narrative text. *Journal of the American Medical Informatics Association* 22, 5 (2015), 1009–1019.
- [36] Tengfei Ma, Cao Xiao, and Fei Wang. 2018. Health-ATM: A Deep Architecture for Multifaceted Patient Health Record Representation and Risk Prediction. In *Proceedings of the 2018 SIAM International Conference on Data Mining*. SIAM, 261–269.
- [37] Kenneth Marek, Danna Jennings, Shirley Lasch, Andrew Siderowf, Caroline Tanner, Tanya Simuni, Chris Coffey, Karl Kieburtz, Emily Flagg, Sohini Chowdhury, et al. 2011. The parkinson progression marker initiative (PPMI). *Progress in neurobiology* 95, 4 (2011), 629–635.
- [38] Vasileios Megalooikonomou, James Ford, Li Shen, Fillia Makedon, and Andrew Saykin. 2000. Data mining in brain imaging. *Statistical Methods in Medical Research* 9, 4 (2000), 359–394.
- [39] Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-Value Memory Networks for Directly Reading Documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 1400–1409.
- [40] 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 (2016), 26094.
- [41] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. 2017. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *Proc. CVPR*, Vol. 1. 3.
- [42] Susumu Mori, Barbara J Crain, Vadappuram P Chacko, and Peter Van Zijl. 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of neurology* 45, 2 (1999), 265–269.
- [43] Brent C Munsell, Chong-Yaw Wee, Simon S Keller, Bernd Weber, Christian Elger, Laura Angelica Tomaz da Silva, Travis Nesland, Martin Styner, Dinggang Shen, and Leonardo Bonilha. 2015. Evaluation of machine learning algorithms for treatment outcome prediction in patients with epilepsy based on structural connectome data. *Neuroimage* 118 (2015), 219–230.
- [44] Ziad S Nasreddine, Natalie A Phillips, Valérie Bédirian, Simon Charbonneau, Victor Whitehead, Isabelle Collin, Jeffrey L Cummings, and Howard Chertkow. 2005. The Montreal Cognitive Assessment, MoCA: a brief screening tool for mild cognitive impairment. *Journal of the American Geriatrics Society* 53, 4 (2005), 695–699.
- [45] Alvin Rajkumar, Eyal Oren, Kai Chen, Andrew M Dai, Nissan Hajaj, Michaela Hardt, Peter J Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, et al. 2018. Scalable and accurate deep learning with electronic health records. *npj Digital Medicine* 1, 1 (2018), 18.
- [46] Pranav Rajpurkar, Awni Y Hannun, Masoumeh Haghighpanahi, Codie Bourn, and Andrew Y Ng. 2017. Cardiologist-level arrhythmia detection with convolutional neural networks. *arXiv preprint arXiv:1707.01836* (2017).
- [47] Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A Neural Attention Model for Abstractive Sentence Summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 379–389.
- [48] Robin Tibor Schirrmeyer, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin Glasstetter, Katharina Eggensberger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and Tonio Ball. 2017. Deep learning with convolutional neural networks for EEG decoding and visualization. *Human brain mapping* 38, 11 (2017), 5391–5420.
- [49] Stephen M Smith. 2002. Fast robust automated brain extraction. *Human brain mapping* 17, 3 (2002), 143–155.
- [50] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*. 2440–2448.
- [51] Jimeng Sun, Jianying Hu, Dijun Luo, Marianthi Markatou, Fei Wang, Shahram Edabollahi, Steven E Steinhilber, Zahra Daar, and Walter F Stewart. 2012. Combining knowledge and data driven insights for identifying risk factors using electronic health records. In *AMIA Annual Symposium Proceedings*, Vol. 2012. American Medical Informatics Association, 901.
- [52] Yichen Wang, Robert Chen, Joydeep Ghosh, Joshua C Denny, Abel Kho, You Chen, Bradley A Malin, and Jimeng Sun. 2015. Rubik: Knowledge guided tensor factorization and completion for health data analytics. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1265–1274.
- [53] Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *International Conference on Learning Representations*.
- [54] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*. 2048–2057.
- [55] Jun Yin, Xin Jiang, Zhengdong Lu, Lifeng Shang, Hang Li, and Xiaoming Li. 2016. Neural generative question answering. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 2972–2978.
- [56] Liang Zhan, Jiayu Zhou, Yalin Wang, Yan Jin, Neda Jahanshad, Gautam Prasad, Talia M Nir, Cassandra D Leonardo, Jieping Ye, Paul M Thompson, et al. 2015. Comparison of nine tractography algorithms for detecting abnormal structural brain networks in Alzheimer's disease. *Frontiers in aging neuroscience* 7 (2015), 48.
- [57] Xi Zhang, Lifang He, Kun Chen, Yuan Luo, Jiayu Zhou, and Fei Wang. 2018. Multi-View Graph Convolutional Network and Its Applications on Neuroimage Analysis for Parkinson's Disease. *arXiv preprint arXiv:1805.08801* (2018).
- [58] Jiayu Zhou, Fei Wang, Jianying Hu, and Jieping Ye. 2014. From micro to macro: data driven phenotyping by densification of longitudinal electronic records. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 135–144.