

Original Paper

ExAD-GNN: Explainable Graph Neural Network for Alzheimer’s Disease State Prediction from Single-cell Data

Ziheng Duan, Cheyu Lee and Jing Zhang*

Department of Computer Science, University of California, Irvine, USA

ABSTRACT

Alzheimer’s disease (AD) is a prevalent neurodegenerative disorder with significant impacts on patients and their families. Therefore, accurate and early diagnosis of AD is crucial for improving patient outcomes and developing effective treatments. However, despite advancements in machine learning for AD diagnosis, current methods lack molecular-level insights and completely ignore the heterogeneity in complex human brains, thus potentially masking crucial disease mechanisms. Here, we present ExAD-GNN, an Explainable Graph Neural Network for predicting AD status from single-cell sequencing data. Leveraging K Nearest Neighbours (KNN) graphs derived from the expression profiles of individual cells, ExAD-GNN achieves two primary goals: predicting AD pathology at a cellular level and identifying cell-type-specific marker genes for AD diagnosis through a unique learnable gene importance metric. Extensive benchmarking on large-scale scRNA-seq data with state-of-the-art methods demonstrates ExAD-GNN’s noticeably improved AD prediction accuracy and robustness across various cell types and samples. Furthermore, an extensive ablation study and literature search confirm the majority of top AD risk genes highlighted by our method, demonstrating the effectiveness of ExAD-GNN’s model interpretation scheme. In summary, we develop ExAD-GNN as a publicly

*Corresponding author: Jing Zhang, zhang.jing@uci.edu

available software for the scientific community to gain molecular insights into AD pathology from scRNA-seq data.

1 Introduction

Alzheimer’s disease (AD) is a neurodegenerative disorder characterized by the progressive loss of memory and cognitive function, affecting millions of individuals worldwide and placing a tremendous burden on patients, their families, and healthcare systems [6]. Therefore, it is crucial to develop accurate and timely AD diagnosis methods to allow appropriate patient care, facilitate targeted interventions, and advance scientific understanding and treatment options for this debilitating disease [27]. Unfortunately, AD is a complex disorder at the molecular level [39]. While decades of efforts have narrowed down a few risk genes, the genetic and molecular mechanisms underlying AD are still largely unknown [3]. We face significant hurdles in developing effective early diagnosis and treatment plans for this devastating disease [32].

Recent advances in machine learning algorithms and initiatives for transparent data access open new avenues for AD diagnosis using advanced computational models [60]. For instance, numerous machine learning models have been proposed to predict AD status using cognitive tests [34], magnetic resonance imaging (MRI) [7, 25, 33], positron emission tomography (PET) images [13, 14], cerebrospinal and blood biomarkers [9, 50]. Additionally, several machine learning models, including support vector machines, artificial neural networks, and deep learning, have been proposed to classify AD status using large-scale neural imaging databases [46, 51, 61]. While promising by demonstrating improved accuracy, it is still difficult to gain molecular-level insights into disease pathology in AD [23]. Later on, as high-throughput sequencing technologies continue to advance, population-scale genetic, transcriptomic, and epigenetic profiling studies further supplement the existing data types to detect AD at the earliest possible stage (pre-dementia) [41]. Combined with the aforementioned clinical data, several machine learning strategies, including univariate associations and multivariate deep learning models, have been used to successfully predict disease status with high accuracy and identify novel biomarkers in AD [1, 59]. However, the human brain is a complex tissue with distinct cell types, each contributing uniquely to AD pathology. Most of the existing studies rely on bulk tissue-level sequencing data, which are profiled from thousands to millions of mixed cells. As a result, they completely overlook the heterogeneity within tissues, potentially masking critical insights into disease mechanisms in AD [65].

Technological developments in single-cell sequencing revolutionize AD research by simultaneously profiling gene expression in individual cells. It

has pushed the investigation of cellular heterogeneity and intricate biological processes with unprecedented resolution – individual cells [8, 19, 24, 35, 36, 38, 40, 42, 49, 57, 71, 73], paving the way for novel insights into the molecular underpinnings of AD. As a result, several single-cell genomic research has been conducted to investigate AD pathology and to provide new molecular insights. For instance, Mathys *et al.* [41] performed population-scale single-cell RNA sequencing (scRNA-seq) in post-mortem human brains from AD patients and healthy controls and revealed both cell-type-specific and cell-type-shared transcription perturbation signatures in AD. Besides, Morabito *et al.* [45] performed single-cell epigenetic and transcriptomic profiling and identified cell-type-specific cis-regulatory elements (CREs) and transcription factors (TF) that may mediate gene-regulatory changes in late-stage AD. However, these studies mainly focus on differences in AD brains and healthy controls at the intermediate phenotype level (e.g., gene expression and regulation), resulting in uncertainties about the direct impact on clinical diagnosis.

To fill this gap, we present an explainable graph neural network, ExAD-GNN, for characterizing and predicting AD pathological states from scRNA-seq data. Specifically, our method builds a cell-to-cell similarity graph based on the uniformly processed cell embeddings from AD patients and healthy controls, with cells and their gene expression profiles representing nodes and node features, respectively. We hypothesize that AD could introduce robust molecular perturbations across different patients in a cell-type-specific manner, which leads to memory impairment and cognitive decline that can ultimately affect behavior, speech, visuospatial orientation, and the motor system. With this straightforward intuition, our model combines cell neighboring information and gene expression profiles to achieve two goals: – predict AD pathology at a cell level and prioritize marker genes for AD diagnosis in each cell type (Figure 1). Importantly, ExAD-GNN’s design incorporates a cell type-specific gene importance score matrix, a significant feature that enhances the model’s interpretability by allowing immediate insights into the importance of different genes and modulating the influence of individual genes on the final AD prediction for each cell type. This inherent interpretability [43] not only deepens our understanding of the disease process at a molecular level but also aids in identifying potential therapeutic targets [44]. Therefore, ExAD-GNN provides not only accurate predictions at cell levels but also affords valuable insights, paving the way for mechanistic exploration and possible therapeutic interventions in AD.

To demonstrate the effectiveness of ExAD-GNN, we applied it to scRNA-seq from 31 post-mortem human prefrontal cortex (PFC) samples. We showed that our model can accurately predict AD conditions at the individual cell level. Moreover, we found that our model could highlight novel and known genes associated with AD in different cell types. Finally, we have made ExAD-GNN freely available as a software package for the community to predict AD

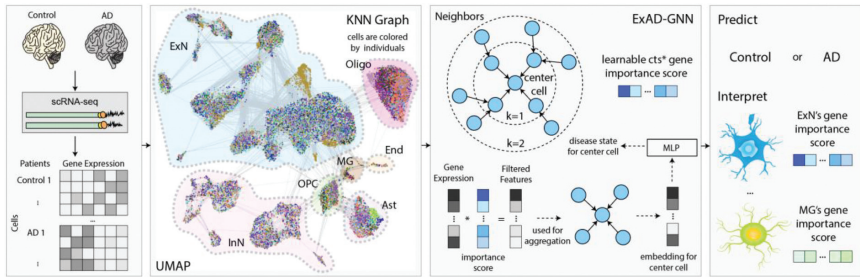


Figure 1: ExAD-GNN’s overview.

Note: The process includes four key steps: (1) Assembly of a cell-by-gene matrix using both AD and control samples, where cells from both conditions are presented along with their gene expression profiles. (2) Construction of a KNN graph based on the cell-by-gene matrix, where connections are made between each cell and its K nearest neighbours according to their gene expression similarities. (3) The detailed architecture of ExAD-GNN, where the model takes the KNN graph as input and learns to predict disease status for each cell by capturing the similarities and differences between AD and control cells in the high-dimensional expression space. (4) Prediction and interpretation, where the model predicts the disease status for each cell and highlights the significant genes influencing the predictions, thereby improving interpretability.

conditions and quantify their changes across various cell types and conditions. With the exponential growth in the availability of single-cell data, we expect that ExAD-GNN can predict disease conditions and highlight disease risk genes with higher accuracy, pushing our understanding of disease pathology to a single-cell resolution.

2 Methodology

As shown in Figure 1, our ExAD-GNN is an explainable graph neural network model to predict cellular level disease status and prioritize AD risk genes in a cell type specific manner. To achieve this goal, it first builds a KNN graph based on the cell-cell similarity and combined cell neighbouring (node) information with gene expression profiles (feature) to distinguish cells from AD samples with those from healthy controls. Method details for data process, modeling parameters, and performance benchmarking will be provided in the following sections.

2.1 Detailed scRNA-seq Data Pre-processing and Assembly of the Cell-by-Gene Matrix

We downloaded scRNA-seq data from post-mortem human prefrontal cortex (PFC), which included 19 AD samples and 12 healthy controls. After strict QC, 35,473 nuclei were preserved for downstream analyses. The detailed scRNA-seq data pre-processing pipeline was as follows.

Step 1. Count Matrix Generation and Ambient RNA Clean-up

We first downloaded the raw reads and used CellRanger count v6.0 [74] to generate the cell-by-count matrix for each sample, which was run independently. No aggregation (using CellRanger aggr) was carried out at this stage. To more carefully separate out true cells from empty droplets with ambient RNA, we used the program remove-background¹ from the CellBender package [58]. For efficiency considerations, we employed the program in command-line form, which was wrapped in our Python script and utilized GPUs with default parameters. Specifically, the options included in the program run were: a target false positive rate (`-fpr`) of 0.01; the number of training epochs (`-epochs`) = 150; the rough expected number of cells (`-expected-cells`) = the output in `metrics_summary.csv` from CellRanger count.

Step 2. Per-fastq Set/Sample QC using Pegasus

Next, we performed per sample QC in Pegasus. After filtering cells based on the lower bounds, we removed 1,135 genes included in the MitoCarta v3.0 database [53] such as mitochondrial genes and certain genes highly correlated with RNA sample quality (for example, Hodge *et al.* [22]). The robust genes were identified, and the counts matrix was log-normalized using the default options in Pegasus. Next, doublets were identified using a combination of Scrublet [69] (default mode) and DoubletDetection [70]. The parameters for the DoubletDetection BoostClassifier algorithm included `n_iters` = 25, `use_phenograph` = False, and `standard_scaling` = True. The subsequent predict function employed the parameters `p_thresh` = 1e-16 and `voter_thresh` = 0.3.

Step 3. Sample Merging, Dimension Reduction, Clustering, and Cell Type Annotation

After filtering out samples with fewer than 500 cells, our dataset comprised 35,473 cells, each with 26,495 genes, from 31 samples, including 12 control samples and 19 AD samples. We further processed the data by identifying robust genes and retaining only those genes for subsequent analysis. The data was then log-normalized for dimensionality reduction using 3,000 highly variable genes using Principal Component Analysis (PCA). Then, we removed PCA components with a correlation coefficient greater than 0.05 with gene intensities, resulting in a final set of 41 dimensions. As shown in Figure 2A, most samples were homogeneously mixed.

Next, we used Pegasus' `infer_cell_types` function to associate the Leiden clusters with reference cell types based on the hybrid marker gene sets obtained

¹<https://cellbender.readthedocs.io/en/latest/index.html>.




Figure 2: UMAPs of the scRNA-seq data.

Note: (A) UMAP of the scRNA-seq data coloured by samples. (B) UMAP scatter plot of the scRNA-seq data coloured by cell types.

from merging neuronal subclass markers from the BRAIN Initiative Cell Census Network (BICCN) taxonomy and non-neuronal subclasses from Network [48]. In total, we defined 8 cell types with 24 subclasses. As shown in Figure 2B, there was a clear clustering effect with samples of the same cell type aggregating together.

2.2 Construction of a KNN Graph

From our pre-processed data, we used the 41-dimensional PCA features to construct a KNN graph. For each of the 35,473 cells, K nearest neighbours within this 41-dimensional space were identified, resulting in a graph with 35,473 nodes, each having 26,495 features. The KNN graph encapsulates the intricate relationships between cells based on their gene expression similarities. This robust representation of the local structure of the high-dimensional data is integral to the effective application of our ExAD-GNN model in subsequent stages of analysis.

2.3 The Detailed Architecture of ExAD-GNN

We start by representing our input graph as $G = (V, E)$, where V is the set of nodes (cells), and E represents the edges between them. Each node i in V is associated with a gene feature vector $f_i \in \mathbb{R}^D$, where D is the number of features (genes). Our objective is to output probabilities of the test cells belonging to the control or AD groups. This comprehensive approach enables researchers to predict AD pathology at a cell level and prioritize marker genes for AD diagnosis in each cell type.

ExAD-GNN predicts AD pathology at a cellular level

For each node i in the graph G , we obtain the node’s embedding by aggregating its neighbours’ information [67]. Inspired by GraphSAGE [20], we adopt the

following aggregator function for the information propagation:

$$h_i^{k+1} = \sigma \left(W^k \left(\text{AGGREGATE} \left(h_j^k \mid \forall j \in N(i) \right) \oplus h_i^k \right) \right), \quad (1)$$

where h_i^k and h_i^{k+1} are the hidden representations of node i in the k -th and $(k+1)$ -th layer, respectively; $N(i)$ represents the set of neighbours of node i ; \oplus is the concatenate function and AGGREGATE is a function that combines the hidden representations of neighbours. σ is the activation function (e.g., ReLU) and W^k is the learnable parameter for the k -th layer. We will examine different AGGREGATE functions in the experiment section. The input feature of the first GNN layer is given in the next subsection. After several GNN layers, we compute the output probabilities for the test nodes by applying a softmax function to the final hidden representations H_i :

$$P_i = \text{softmax}(H_i). \quad (2)$$

To train our model, we minimize the cross-entropy loss between the predicted probabilities P_i and the true labels Y_i (0 for control, 1 for AD) for all test nodes:

$$L_{\text{predict}} = - \sum (Y_i * \log(P_i) + (1 - Y_i) * \log(1 - P_i)). \quad (3)$$

ExAD-GNN prioritizes marker genes for AD diagnosis in each cell type

A crucial aspect of our model is the incorporation of the cell type-specific gene importance score matrix. We denote it as $S \in \mathbb{R}^{|C| \times D}$, where $|C|$ is the number of cell types and each element S_{ij} indicating the importance of gene j for cell type i . To incorporate this information, we multiply the gene feature vectors with the sigmoid of the cell type-specific gene importance score matrix:

$$F_i = \text{sigmoid}(S) * f_i. \quad (4)$$

For the first GNN layer, we have $h_i^0 = F_i$. Additionally, we encourage sparsity in the cell type-specific gene importance score matrix by adding an L1 regularization term to the loss function:

$$L_{\text{reg}} = \|\text{sigmoid}(S)\|_1. \quad (5)$$

The total loss function is given by:

$$L_{\text{total}} = L_{\text{predict}} + \alpha * L_{\text{reg}}, \quad (6)$$

where α is the penalty weight. This approach allows our model to focus on the most important genes for each cell type, improving interpretability and performance.

2.4 Prediction and Interpretation with ExAD-GNN

Having detailed the structure of the ExAD-GNN model in the previous section, we now describe how it is used for prediction and interpretation at a cellular level. For a given test cell, the first step is to aggregate the information from its neighbors in the KNN graph. This aggregation is performed according to the (1) detailed earlier, where we incorporate the hidden representations of the cell’s neighbors to update the cell’s own representation. With this updated representation of the test cell, we then apply a softmax function according to (2) to the final hidden representation of the test cell. This yields a probability distribution over the two possible classes: control and AD. For the interpretation phase, we utilize the learned cell type-specific gene importance score matrix S . We retrieve the specific vector S_i corresponding to the cell type of the test cell, where i denotes the cell type of the test cell. This vector $S_i \in \mathbb{R}^D$ (D is the number of genes) contains the importance scores of each gene for the given cell type. The genes with the highest scores in this vector are identified by ExAD-GNN as the most influential genes for the prediction of the test cell.

2.5 Alternative Models and Parameters

We compared ExAD-GNN with three categories of alternative models, including naïve models (random guess), traditional machine learning methods (e.g., Random Forest and KNN), and deep learning-based approaches (including Multi-Layer Perceptron (MLP), Graph Convolutional Network (GCN) [29], Graph Attention Network (GAT) [62], and GraphSAGE [20]). For each method, we carefully selected the parameters to ensure fair comparisons and consistent evaluation.

For KNN, we explored different values of k in the range (1, 3, 5, 10, 30, 50, 100), selecting the value that yielded the highest validation accuracy for each dataset. In the case of Random Forest, we used the default parameters, allowing trees to grow without height restrictions. For MLP, GCN, GAT, GraphSAGE, and ExAD-GNN, we maintained a uniform two-layer architecture, with a hidden layer dimension of 128. We employed the Adam optimizer [28] with a learning rate of 1e-3 for gradient descent. The batch size for training was set to 256. To prevent overfitting, we terminated the training process if the validation accuracy did not improve for 30 consecutive epochs. In the context of GAT, we set the number of attention heads to 8, enabling the model to capture multiple aspects of node relationships. Both GraphSAGE and ExAD-GNN sampled 4 neighbours for aggregation and utilized the mean aggregation strategy. Furthermore, ExAD-GNN incorporated a penalty weight α of 1e-4 to balance AD predict accuracy and gene selection importance.

2.6 Experimental Settings for Performance Benchmarking

Our benchmark and experimental settings contained two transductive tasks for cellular level AD prediction: (i) a cell-type-specific manner and (ii) a sample-specific manner. For the cell-type-specific manner, we randomly assigned 80% of the nodes from each cell type to serve as training data, with the remaining 20% evenly split between validation and test data. For robustness considerations, we generated ten different train/validation/test splits and computed the average accuracy and standard deviation across these splits for each method. This approach ensured a robust and comprehensive assessment of each model’s performance.

To account for the sample-specific performance disparities, we also conducted sample-specific task by masking all cells within one sample to create a test set. Subsequently, we chose one control and one AD sample from the remaining pool to construct a validation set. The remaining 28 samples constituted the training set. We reported the performance metrics for each method using this splitting scheme. Given that there are 31 potential sample-specific splits, we calculated and reported the average accuracy and standard deviation across all these splits. This approach ensured the generalizability of our results.

3 Experimental Results

We applied our ExAD-GNN model to the uniformly processed and annotated scRNA-seq data to predict AD status at the cellular level. For the AD prediction module, we conducted comprehensive extensive performance benchmarks with seven other models from three categories (details see Method section: Alternative Models). We showed that our model outperformed existing models by a large margin. Additionally, we explored the interpretability of our model by identifying key genes contributing to AD prediction, visualizing latent embeddings to demonstrate the discriminative power of ExAD-GNN, and conducting a detailed investigation of hyperparameters tuning. Details of these results were discussed in the following section.

3.1 ExAD-GNN Outperforms Other Methods for the Cellular Level AD Prediction Task in a Cell-Type-Specific Manner

We first aimed to predict AD status at a cellular level, in other words, whether individual cells were from AD patients or healthy controls. While conducting a transductive task in a cell-type-specific manner (details in the Method section: Experimental Settings), we reported our model’s performance on 3,547 test cells (10% for the whole dataset) and comprehensively compared this prediction performance with seven alternative methods (details in the Method

section: Alternative Models and Parameters). We found that ExAD-GNN noticeably outperformed all the other models by a large margin (Table 1). For instance, ExAD-GNN achieved 94% accuracy when merging cells from different cell types, noticeably higher than other deep learning-based methods (e.g., 69%~92% other GNN models, and 90% for MLP), traditional machine learning algorithms (61%~79%), and naïve baselines (56%).

Next, we separated cells from eight major cell types (their summary and relevance to AD are shown in Table 2) and benchmarked cell-type-specific AD prediction accuracy, as shown in Table 1. In six of eight major cell types, ExAD-GNN showed consistently improved prediction accuracy over other methods, demonstrating its robustness across different cell types. It was worth mentioning that prediction accuracy dropped significantly for all models as the number of available cells decreased. For instance, Endo and VLMCs were two relatively rare cell types in the human brain, with only 22 and 74 cells from 31 samples. Some samples even showed very few cells from these cell types, leading to significantly deteriorated prediction performance for all models. Among all types, ExN, Oligo, and Astro showed the highest accuracy (95%, 94%, and 96% respectively), consistent with their important roles in AD pathology [2, 52, 66].

Furthermore, given the complexity and variety of neuronal functions and their potential varying contributions to AD, we further broke down the ExN and InN into several subclasses (bottom sections in Table 1). This step enabled us to assess our model’s performance across highly specialized neuronal subclasses. As shown in Table 1, ExAD-GNN outperformed other methods in six out of nine excitatory neuron subclasses and five out of nine inhibitory neuron subclasses. Despite the broad range of subclasses and their complexity, ExAD-GNN consistently achieved higher accuracy. This underscored the model’s ability to effectively utilize cell type-specific gene importance score matrices, enhancing its precision in major cell types and subclasses.

3.2 ExAD-GNN Outperforms Other Methods for the Cellular Level AD prediction Task in a Sample-Specific Manner

Next, we focused on the performance of ExAD-GNN and seven other methods in a sample-specific manner by predicting the AD status of all cells sample by sample to test our model’s capability of handling inter-individual variations (details in the Method section: Experimental Settings for Performance Benchmarking). As shown in Table 3, ExAD-GNN showed the highest accuracy (69%) across all cell types, noticeably outperforming all other methods (52%~68% for other deep learning methods, 55%~59% for traditional machine learning algorithms, and 50% for naïve baselines). This result suggested the robustness of ExAD-GNN in handling sample-specific tasks. At the same time, a decrease in accuracy and an increase in standard deviation were noted

Table 1: Cellular level AD prediction accuracy in a cell-type-specific manner. The numbers in the table are mean/std, and the best results under each cell type are bolded.

Method	Major cell types										
	All			Astro	Endo	ExN	InN	MG	OPC	Oligo	VLMC
Random Guess	0.56/0.01	0.51/0.04	0.40/0.20	0.50/0.01	0.52/0.02	0.56/0.15	0.50/0.04	0.50/0.04	0.50/0.01	0.40/0.15	
Random Forest	0.79/0.01	0.85/0.03	0.60/0.37	0.78/0.00	0.69/0.01	0.69/0.11	0.78/0.03	0.78/0.03	0.89/0.01	0.74/0.18	
K-Nearest Neighbors	0.61/0.01	0.60/0.03	0.40/0.37	0.63/0.01	0.55/0.03	0.44/0.09	0.61/0.02	0.61/0.02	0.58/0.03	0.57/0.20	
Multi-Layer Perceptron	0.90/0.01	0.88/0.02	0.60/0.37	0.92/0.00	0.85/0.01	0.70/0.07	0.90/0.02	0.90/0.02	0.89/0.01	0.89/0.09	
Graph Convolutional Network	0.69/0.01	0.70/0.02	0.40/0.37	0.70/0.01	0.59/0.02	0.49/0.10	0.62/0.04	0.62/0.04	0.68/0.01	0.56/0.17	
Graph Attention Network	0.80/0.02	0.74/0.05	0.55/0.42	0.85/0.02	0.65/0.04	0.54/0.10	0.75/0.04	0.75/0.04	0.76/0.05	0.60/0.18	
GraphSAGE	0.92/0.01	0.50/0.04	0.60/0.37	0.50/0.01	0.52/0.03	0.48/0.10	0.61/0.03	0.61/0.03	0.50/0.03	0.61/0.21	
ExAD-GNN	0.94/0.00	0.95/0.02	0.30/0.24	0.95/0.00	0.86/0.02	0.71/0.05	0.94/0.01	0.94/0.01	0.96/0.01	0.80/0.11	
Method	Subclass for ExN										
	L2/3 IT	L4 IT	L5 ET	L5 IT	L5/6 NP	L6 CT	L6 IT	L6 IT	L6 IT	Car3	L6b
Random Guess	0.50/0.01	0.51/0.03	0.52/0.08	0.50/0.04	0.48/0.10	0.51/0.08	0.52/0.03	0.52/0.03	0.52/0.06	0.51/0.07	
Random Forest	0.80/0.01	0.77/0.02	0.76/0.10	0.74/0.02	0.78/0.08	0.70/0.05	0.78/0.03	0.78/0.03	0.72/0.04	0.70/0.07	
K-Nearest Neighbors	0.64/0.01	0.69/0.01	0.59/0.10	0.62/0.03	0.63/0.11	0.57/0.06	0.60/0.04	0.60/0.04	0.58/0.04	0.52/0.08	
Multi-Layer Perceptron	0.93/0.01	0.92/0.01	0.92/0.06	0.91/0.01	0.88/0.09	0.90/0.03	0.92/0.02	0.92/0.02	0.91/0.04	0.90/0.03	
Graph Convolutional Network	0.74/0.02	0.73/0.03	0.57/0.16	0.65/0.02	0.68/0.10	0.62/0.08	0.68/0.04	0.68/0.04	0.58/0.06	0.58/0.08	
Graph Attention Network	0.88/0.02	0.85/0.02	0.66/0.23	0.85/0.02	0.73/0.11	0.79/0.07	0.85/0.03	0.85/0.03	0.75/0.04	0.78/0.06	
GraphSAGE	0.55/0.01	0.36/0.02	0.44/0.13	0.47/0.02	0.34/0.09	0.44/0.10	0.54/0.02	0.54/0.02	0.59/0.05	0.49/0.09	
ExAD-GNN	0.96/0.01	0.96/0.01	0.88/0.07	0.96/0.01	0.82/0.08	0.91/0.04	0.96/0.01	0.96/0.01	0.88/0.05	0.92/0.04	
Method	Subclass for InN										
	Chandelier	Lamp5	Lamps Lhx6	Pax6	Pvalb Sncg	Sst	Sst Chodl	Vip			
Random Guess	0.50/0.13	0.54/0.06	0.48/0.08	0.53/0.13	0.53/0.05	0.49/0.08	0.50/0.05	0.53/0.04			
Random Forest	0.68/0.11	0.73/0.05	0.70/0.06	0.66/0.19	0.71/0.02	0.62/0.06	0.64/0.05	0.68/0.05			
K-Nearest Neighbors	0.56/0.13	0.43/0.06	0.55/0.08	0.66/0.16	0.58/0.04	0.44/0.08	0.53/0.03	0.58/0.05			
Multi-Layer Perceptron	0.82/0.08	0.86/0.06	0.83/0.06	0.89/0.03	0.84/0.03	0.84/0.07	0.84/0.03	0.90/0.30			
Graph Convolutional Network	0.52/0.12	0.60/0.06	0.60/0.11	0.59/0.21	0.61/0.04	0.70/0.09	0.57/0.06	0.60/0.49			
Graph Attention Network	0.62/0.10	0.64/0.07	0.60/0.09	0.70/0.17	0.69/0.03	0.66/0.15	0.62/0.07	0.80/0.40			
GraphSAGE	0.45/0.12	0.42/0.05	0.56/0.08	0.64/0.19	0.52/0.05	0.44/0.06	0.53/0.03	0.60/0.49			
ExAD-GNN	0.81/0.04	0.84/0.05	0.83/0.07	0.69/0.16	0.89/0.01	0.82/0.09	0.86/0.03	0.30/0.46			

Table 2: Summary of major cell types and their relevance to AD.

Cell Types	Abbreviation	Description	Relation with AD
Astrocytes	Astro	Astro cells are star-shaped glial cells in the central nervous system.	Astro cells are known for their roles in neural networking, regulation of neurotransmission, and progression of AD [63].
Endothelial cells	Endo	Endo cells form the inner lining of blood vessels and are crucial for the blood-brain barrier.	Alterations in Endo cells are often observed in AD [55].
Excitatory neurons	ExN	ExN cells are the most prevalent type of neurons in the brain.	ExN cells are responsible for transmitting excitatory signals. They are primary targets of AD pathology [15].
Inhibitory neurons	InN	InN cells regulate the excitability of neural circuits and maintain a balance in the brain’s signaling.	InN cells’ dysfunction is implicated in AD [37].
Microglia	MG	MG cells are the primary immune cells of the brain and spinal cord.	MG cells are involved in the inflammatory response in AD [68].
Oligodendrocyte precursor cells	OPC	OPCs play a critical role in developmental and adult myelinogenesis and can differentiate into oligodendrocytes.	Myelin breakdown from impaired repair of OPCs may be the initiating step in AD pathology [5].
Oligodendrocytes	Oligo	Oligo cells are responsible for producing the myelin sheath that insulates neuronal axons, thereby facilitating efficient signal transmission.	Damage to Oligo cells can result in impaired neural communication, a condition that contributes to the progression of AD [54].
Vascular and leptomeningeal cells	VLMC	VLMC cells are involved in the vascular system of the brain and meninges	VLMC cells have been shown to undergo changes in AD [26].

compared to the previous cell-type-specific manner. This was mainly due to the inherent complexity and variability across different samples, which encapsulates a broader range of gene expression profiles and introduces additional challenges in achieving precise classification.

Despite these challenges, ExAD-GNN outperformed six of the eight major cell types. Meanwhile, accuracy for under-represented cell types like Endo and VLMC was still low due to their sparse presence. In a sample-specific manner, we also divided ExN and InN cell types into subclasses, despite the increased complexity due to inter-sample heterogeneity. As shown in Table 3, ExAD-GNN still outperformed other methods in six out of nine excitatory neuron subclasses and five out of nine inhibitory neuron subclasses, demonstrating ExAD-GNN’s robustness in managing complex disease prediction scenarios.

Table 3: Architecture of our proposed one branch. Both branches contain the same structure and parameter settings. Here T denotes the number of time frames.

Method	All									
	Astro		Endo		ExN	Major cell types			Oligo	VLMC
	L2/3 IT	L4 IT	L5 ET	L5 IT	L5/6 NP	L6 CT	L6 IT	L6 IT Car3		
Random Guess	0.50/0.01	0.49/0.04	0.15/0.30	0.50/0.02	0.50/0.06	0.46/0.29	0.54/0.09	0.49/0.04	0.19/0.27	
Random Forest	0.55/0.17	0.67/0.26	0.18/0.36	0.55/0.21	0.49/0.17	0.48/0.32	0.63/0.24	0.60/0.18	0.27/0.39	
K-Nearest Neighbors	0.59/0.32	0.61/0.46	0.24/0.42	0.61/0.44	0.61/0.48	0.52/0.46	0.61/0.48	0.40/0.42	0.26/0.43	
Multi-Layer Perceptron	0.52/0.16	0.53/0.24	0.16/0.32	0.52/0.12	0.48/0.16	0.34/0.31	0.56/0.26	0.55/0.20	0.16/0.35	
Graph Convolutional Network	0.54/0.25	0.62/0.31	0.18/0.37	0.59/0.27	0.50/0.27	0.39/0.37	0.53/0.36	0.58/0.30	0.16/0.35	
Graph Attention Network	0.60/0.21	0.62/0.30	0.10/0.30	0.60/0.25	0.59/0.24	0.38/0.34	0.61/0.30	0.56/0.29	0.11/0.29	
GraphSAGE	0.68/0.26	0.75/0.22	0.23/0.38	0.65/0.30	0.65/0.27	0.61/0.32	0.78/0.23	0.79/0.20	0.28/0.38	
ExAD-GNN	0.69/0.23	0.84/0.17	0.14/0.28	0.66/0.29	0.66/0.26	0.54/0.32	0.82/0.21	0.83/0.21	0.33/0.40	
Method	Subclass for ExN									
	L2/3 IT	L4 IT	L5 ET	L5 IT	L5/6 NP	L6 CT	L6 IT	L6 IT Car3	L6b	
Random Guess	0.50/0.03	0.50/0.08	0.46/0.27	0.49/0.06	0.43/0.23	0.51/0.16	0.51/0.08	0.51/0.13	0.52/0.17	
Random Forest	0.56/0.23	0.56/0.31	0.31/0.34	0.54/0.28	0.49/0.27	0.48/0.26	0.50/0.23	0.52/0.27	0.49/0.21	
K-Nearest Neighbors	0.62/0.46	0.61/0.48	0.52/0.32	0.61/0.44	0.62/0.47	0.61/0.48	0.57/0.35	0.56/0.40	0.60/0.45	
Multi-Layer Perceptron	0.53/0.20	0.53/0.23	0.41/0.32	0.52/0.22	0.44/0.29	0.42/0.20	0.50/0.18	0.43/0.24	0.51/0.22	
Graph Convolutional Network	0.58/0.29	0.56/0.35	0.42/0.43	0.50/0.31	0.56/0.36	0.51/0.37	0.46/0.32	0.41/0.44	0.44/0.36	
Graph Attention Network	0.63/0.26	0.64/0.31	0.35/0.34	0.59/0.30	0.46/0.37	0.53/0.30	0.54/0.28	0.51/0.32	0.51/0.31	
GraphSAGE	0.66/0.30	0.64/0.33	0.66/0.36	0.64/0.31	0.56/0.39	0.59/0.32	0.63/0.33	0.62/0.36	0.62/0.31	
ExAD-GNN	0.67/0.29	0.64/0.32	0.61/0.37	0.64/0.32	0.61/0.36	0.59/0.31	0.62/0.33	0.55/0.34	0.63/0.31	
Method	Subclass for InN									
	Chandelier	Lamp5	Lamp5 Lhx6	Pax6	Pvalb	Sucg	Sst	Sst Chodl	Vip	
Random Guess	0.57/0.21	0.45/0.22	0.52/0.22	0.36/0.29	0.51/0.07	0.47/0.27	0.45/0.20	0.16/0.37	0.51/0.13	
Random Forest	0.47/0.30	0.53/0.28	0.47/0.35	0.35/0.36	0.47/0.19	0.54/0.29	0.46/0.27	0.06/0.25	0.47/0.19	
K-Nearest Neighbors	0.61/0.49	0.61/0.49	0.61/0.49	0.51/0.49	0.61/0.47	0.61/0.49	0.61/0.48	0.19/0.40	0.61/0.48	
Multi-Layer Perceptron	0.45/0.34	0.43/0.23	0.39/0.24	0.45/0.38	0.47/0.18	0.47/0.29	0.47/0.22	0.06/0.25	0.50/0.18	
Graph Convolutional Network	0.50/0.46	0.43/0.35	0.52/0.36	0.44/0.46	0.52/0.34	0.51/0.32	0.47/0.39	0.16/0.37	0.48/0.35	
Graph Attention Network	0.58/0.40	0.54/0.28	0.59/0.37	0.42/0.40	0.60/0.27	0.61/0.29	0.59/0.31	0.18/0.37	0.58/0.29	
GraphSAGE	0.66/0.29	0.65/0.33	0.59/0.33	0.51/0.39	0.64/0.27	0.64/0.34	0.64/0.29	0.18/0.37	0.67/0.26	
ExAD-GNN	0.64/0.31	0.65/0.32	0.62/0.33	0.48/0.39	0.63/0.29	0.67/0.34	0.68/0.27	0.10/0.30	0.67/0.27	

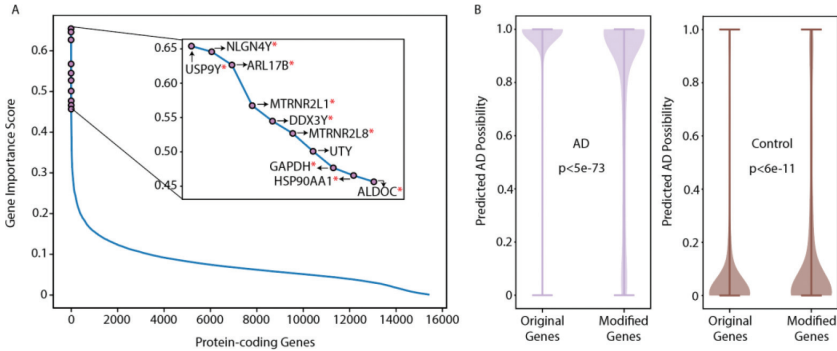


Figure 3: Identifying importance genes.

Note: (A) Gene importance score for ExN. (B) Predicted AD possibility for AD and control with original genes and modified genes.

3.3 The Model Interpretation Module of ExAD-GNN Prioritizes Key Genes Contributing to AD Prediction

In this section, we demonstrated how the model interpretation module of ExAD-GNN can identify AD risk genes in a cell-type-specific manner. With its cellular-level AD prediction, ExAD-GNN enabled an end-to-end training fashion of the gene importance score matrices specific to cell types. This inherent interpretability allowed it to prioritize key genes that significantly contribute to AD prediction (details in the Method section: The Detailed Architecture of ExAD-GNN). As shown in Figure 3A, we started from the fully trained ExAD-GNN module and plotted the sorted genes impact scores in the L2/3 IT cells, a subclass of excitatory neurons. Remarkably, nine of the top ten prioritized genes were reported as AD risk genes previously (stared genes in Figure 3A) [4, 17, 18, 21, 30, 31, 47, 56, 72]. For instance, our third highest ranked gene, ARL17B, has been associated with a decreased risk of AD in the APOE4 negative population [64]. Another gene, MTRNR2L1, also known as Humanin like-1, was among the top upregulated mRNAs in AD [16]. These results demonstrated the effectiveness of ExAD-GNN’s interpretation module.

Next, we performed an ablation study to further validate and visualize the effect of our prioritized risk genes. Specifically, we randomly selected 10% of all cells in the test data set and calculated the difference in AD risk (Δp) by changing the observed gene expression levels (details in the Method section: The Detailed Architecture of ExAD-GNN). As shown in Figure 3B, the top 10 risk genes introduced significantly larger Δp values than the rest of expressed genes (an average of 93.8% and 78.7% for top risk genes and the other genes, $p < 5e-73$ by independent samples t -test). This further testified to ExAD-GNN’s ability to accurately reflected the relationship between gene expression and AD prediction.

3.4 ExAD-GNN Shows Discriminative Power in AD Prediction by Visualizing Latent Embeddings

Cell representation learning is essential to project cells from the high dimensional, sparse cell by gene matrix to a low dimensional, dense space suitable for downstream analyses, such as clustering. Several methods have been proposed to learn joint cell representation by forcing cells from different biological conditions to be perfectly aligned. Here, we tested a different approach from a by-product of our ExAD-GNN - to learn better latent cells by capturing both cell type and disease information. We extracted the latent embeddings after the first GNN layer of a trained ExAD-GNN as the final cell representations (details in Method section: The Detailed Architecture of ExAD-GNN). For a fair comparison, we plotted embeddings of 10% of the cells using traditional PCA and ExAD-GNN. As shown in Figure 4A, the UMAP from PCA showed significant overlap between AD and control cells, leaving them difficult to distinguish. In contrast, the UMAP from the ExAD-GNN embeddings effectively separated AD from control cells while retaining cell-type characteristics (Figure 4B), proving ExAD-GNN’s capacity to learn meaningful representations that distinguish disease states while preserving cell-type-specific features.

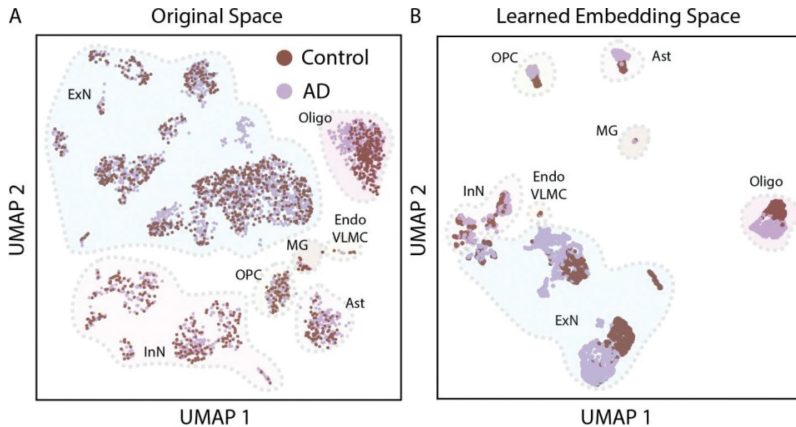


Figure 4: Visualization of embedding space.

Note: Comparison of the original space (A) and the embedding space learned by ExAD-GNN (B).

3.5 ExAD-GNN Demonstrates Robustness and Efficiency through Detailed Training and Parameter Sensitivity Analysis

Hyperparameter tuning is an important step that affects model performance. To test the robustness of our result to such hyperparameters, we investigated




Figure 5: Loss and accuracy during training.

Note: (A) Training loss w.r.t epoch while we mask 10% for each cell types. (B) Train accuracy and valid accuracy w.r.t epoch while we mask 10% for each cell types. (C) Training loss w.r.t epoch while we mask one individual. (D) Train accuracy and valid accuracy w.r.t epoch while we mask one individual.

the effect of different parameters in our training process. First, we evaluated the efficiency of our model as follows. During the training process for both cell-type-specific and sample-specific tasks, we found that our model showed rapid convergence, achieving near-optimal performance within just 50 epochs as shown in Figure 5. We chose to halt the training process when the accuracy of the validation set failed to increase for 30 consecutive epochs. These findings provided strong evidence for ExAD-GNN’s efficiency across diverse tasks.

Next, we analysed the effects of various parameters and model settings on ExAD-GNN’s performance, including the dimension of latent embeddings, the penalty weight α , the number of GNN layers, the number of sampled neighbours, aggregation types and interpretation modules for ExAD-GNN [10–12]. Our findings were as follows. First, cellular level AD prediction was robust to the latent embedding dimensionality, with accuracy changing only slightly from 92% to 94% for embeddings sized 32 to 512 (Figure 6A). Besides, the penalty weight α , which balances disease prediction and risk gene prioritization, was vital. Intuitively, a larger α usually allowed ExAD-GNN to shift its focus towards gene selection by sacrificing the prediction accuracy, while a smaller α would make ExAD-GNN gradually degenerate into an ordinary GNN, no longer filtering for important genes. We hoped to identify genes with the most outstanding contribution to AD prediction while ensuring accuracy. Therefore, choosing an appropriate α was of utmost importance for ExAD-GNN. In this experiment, we found the optimal value for α to be $1e-4$ through rigorous testing and validation. Furthermore, we tested the impact of GNN layers on ExAD-GNN’s predictive accuracy. As shown in Figure 6D, we selected two layers of GNN to calculate the prediction accuracy. Next, we found that the number of neighbours selected for each aggregation process didn’t significantly affect the accuracy (93% to 94% for different numbers of neighbours). Therefore, we chose four neighbours for aggregation in our model. These analyses underscored the robustness of ExAD-GNN to various parameter settings.

In addition, we found that different aggregation strategies had minimal impact on the prediction accuracy (92% \sim 94%, Figure 6E), so we selected

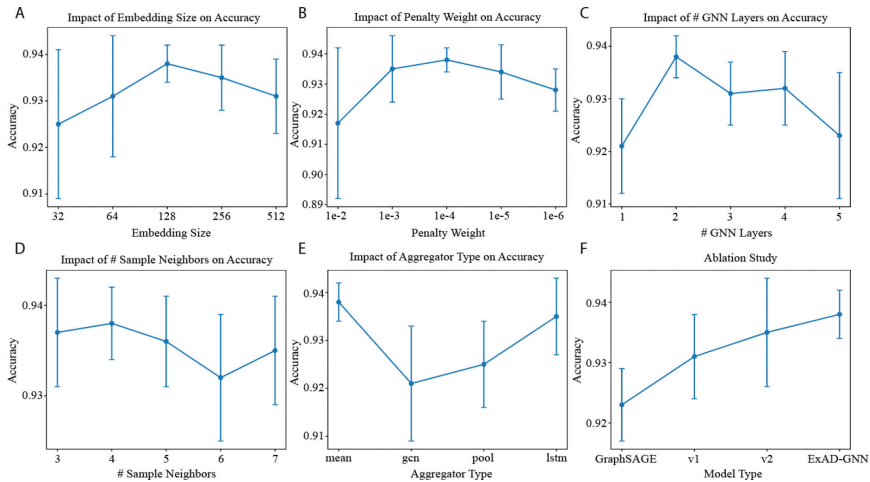


Figure 6: Parameter sensitivity analysis (A-D) and ablation study (E, F).

Note: (A) Test accuracy with different embedding sizes. (B) Test accuracy with different penalty weights. (C) Test accuracy with different number of GNN layers. (D) Test accuracy with different number of sample neighbours. (E) Test accuracy with different aggregator types. (F) Test accuracy with different model types.

the highest “mean” strategy in our model. Further, to comprehend the crucial role of cell-type-specific gene importance score, we explored two versions in ExAD-GNN. In the first version (v1), a single gene importance score was shared among all cell types, which means $S_{v1} \in \mathbb{R}^{1 \times D}$ (details in the Method section: The Detailed Architecture of ExAD-GNN). The second version (v2) assigned unique gene importance scores to each major cell type, but without further differentiation between subclasses of ExN and InN, meaning $S_{v2} \in \mathbb{R}^{|\text{MC}| \times D}$, where $|\text{MC}|$ denotes the number of major cell types (eight in our experiments). As demonstrated in Figure 6F, providing more specific gene importance scores for finer cell types incrementally improved AD prediction accuracy (92% for GraphSAGE, and increasing accuracy for v1 and v2, with ExAD-GNN reaching the highest at 94%). This emphasized the significance of providing learnable gene importance scores for each cell type in our model.

4 Discussion and Conclusion

We developed ExAD-GNN, an accurate and interpretable model for AD diagnosis and risk gene prioritization. Distinct from existing models, it takes advantage of the recent technological developments in the single-cell revolution by predicting AD pathology at the finest possible resolution from the heterogeneous human brains – individual cells. Both across-cell and across-sample

analyses demonstrate ExAD-GNN’s accuracy and effectiveness in capturing sample-specific and cell-type-specific information at the molecular level for AD status prediction. Besides, ExAD-GNN leverages the recent methodology advances in explainable GNNs and uses learnable parameters to highlight key features that strongly impact cellular decision-making processing, enabling us to prioritize AD risk genes in a cell-type-specific manner. To test its effectiveness, we applied it to 31 scRNA-seq data from the human prefrontal cortex of AD patients and healthy controls. As a result, nine of the top ten prioritized genes in excitatory neurons have been previously reported to play critical roles in AD, demonstrating its ability to pick up AD-associated genes. Our detailed investigation of the training process and parameter sensitivity analysis emphasized ExAD-GNN’s robustness to varying parameter settings. This adaptability is crucial considering the complex and diverse nature of cell-specific gene expression profiles.

While the current results are promising, there are certain limitations to our ExAD-GNN model. Specifically, our model’s predictive accuracy is limited in highly underrepresented cell types. For instance, ExAD-GNN’s AD prediction accuracy in rare cell types, such as endothelial and VLMC, significantly deteriorated (94% on average vs. 30% and 80%) due to very little training data (22 cells for Endo and 74 cells for VLMC). However, as single-cell sequencing technology advances, we anticipate the sequencing depth and cell number per sample will significantly increase, allowing an improved performance of ExAD-GNN to capture the intricacies of these cell types.

In summary, we developed ExAD-GNN into publicly available software to tackle the prediction and interpretation of AD status at a single-cell resolution. As research into AD progresses towards a more personalized understanding of disease mechanisms, we anticipate that ExAD-GNN will serve as an invaluable tool in predicting disease outcomes and identifying the key determinants of AD in a cell-type-specific manner. Ultimately, the model’s ability to enhance predictive accuracy and offer insights into critical genes has the potential to contribute to the field of AD research significantly.

Acknowledgement

We thank the OIT of the UCI ICS department for supporting computing resources.

Financial Support

This work has been supported by National Institutes of Health under award number R01NS128523.

Biographies

Ziheng Duan received his B.S. in July 2020 at Zhejiang University, College of Control Science and Engineering. He is now a Ph.D. student at the University of California, Irvine. His main research interests are bioinformatics and machine learning.

Cheyu Lee received his B.S. in June 2021 at the University of California, Berkeley, College of Chemistry. He is now a M.S. student at the University of California, Irvine. His main research interests are bioinformatics and machine learning.

Jing Zhang received her Ph.D. in December 2013 at the University of Southern California, Electrical Engineering Department. She is now an Assistant Professor at the University of California, Irvine. His main research interests are bioinformatics and machine learning.

References

- [1] A. Aljovic, A. Badnjevic, and L. Gurbeta, “Artificial Neural Networks in the Discrimination of Alzheimer’s disease Using Biomarkers Data,” in *2016 5th Mediterranean Conference on Embedded Computing (Meco)*, 2016, 286–9.
- [2] S. Bandyopadhyay, “Role of Neuron and Glia in Alzheimer’s Disease and Associated Vascular Dysfunction,” *Frontiers in Aging Neuroscience*, 13, 2021.
- [3] Z. Breijyeh and R. Karaman, “Comprehensive Review on Alzheimer’s Disease: Causes and Treatment,” *Molecules*, 25(24), 2020.
- [4] D. A. Butterfield, S. S. Hardas, and M. L. Lange, “Oxidatively modified glyceraldehyde-3-phosphate dehydrogenase (GAPDH) and Alzheimer’s disease: many pathways to neurodegeneration,” *J Alzheimers Dis*, 20(2), 2010, 369–93.
- [5] Z. Cai and M. Xiao, “Oligodendrocytes and Alzheimer’s disease,” *Int J Neurosci*, 126(2), 2016, 97–104.
- [6] R. J. Caselli *et al.*, “Alzheimer Disease: Scientific Breakthroughs and Translational Challenges,” *Mayo Clin Proc*, 92(6), 2017, 978–94.
- [7] S. Chaplot, L. M. Patnaik, and N. R. Jagannathan, “Classification of magnetic resonance brain images using wavelets as input to support vector machine and neural network,” *Biomedical Signal Processing and Control*, 1(1), 2006, 86–92.
- [8] T. Denyer and M. C. P. Timmermans, “Crafting a blueprint for single-cell RNA sequencing,” *Trends Plant Sci*, 27(1), 2022, 92–103.

- [9] J. D. Doecke et al., “Blood-Based Protein Biomarkers for Diagnosis of Alzheimer Disease,” *Archives of Neurology*, 69(10), 2012, 1318–25.
- [10] Z. Duan et al., “Multivariate time-series classification with hierarchical variational graph pooling,” *Neural Netw*, 154, 2022, 481–90.
- [11] Z. Duan et al., “Connecting latent relationships over heterogeneous attributed network for recommendation,” *Applied Intelligence*, 52(14), 2022, 16214–32.
- [12] Z. Duan et al., “Multivariate Time Series Forecasting with Transfer Entropy Graph,” *Tsinghua Science and Technology*, 28(1), 2023, 141–9.
- [13] J. Dukart et al., “Meta-analysis based SVM classification enables accurate detection of Alzheimer’s disease across different clinical centers using FDG-PET and MRI,” *Psychiatry Research-Neuroimaging*, 212(3), 2013, 230–6.
- [14] Y. Fan et al., “Structural and functional biomarkers of prodromal Alzheimer’s disease: A high-dimensional pattern classification study,” *Neuroimage*, 41(2), 2008, 277–85.
- [15] D. Franjic et al., “Transcriptomic taxonomy and neurogenic trajectories of adult human, macaque, and pig hippocampal and entorhinal cells,” *Neuron*, 110(3), 2022, 452–469 e14.
- [16] X. Fu et al., “A blood mRNA panel that differentiates Alzheimer’s disease from other dementia types,” *Journal of Neurology*, 270(4), 2023, 2117–27.
- [17] M. Garofalo et al., “Alzheimer’s, Parkinson’s Disease and Amyotrophic Lateral Sclerosis Gene Expression Patterns Divergence Reveals Different Grade of RNA Metabolism Involvement,” *Int J Mol Sci*, 21(24), 2020.
- [18] M. Gonzalez-Rodriguez et al., “Neurodegeneration and Astrogliosis in the Human CA1 Hippocampal Subfield Are Related to hsp90ab1 and bag3 in Alzheimer’s Disease,” *Int J Mol Sci*, 23(1), 2021.
- [19] G. S. Gulati et al., “Single-cell transcriptional diversity is a hallmark of developmental potential,” *Science*, 367(6476), 2020, 405–+.
- [20] W. L. Hamilton, R. Ying, and J. Leskovec, “Inductive Representation Learning on Large Graphs,” *Advances in Neural Information Processing Systems 30 (Nips 2017)*, 30, 2017.
- [21] L. He et al., “Exome-wide age-of-onset analysis reveals exonic variants in ERN1 and SPPL2C associated with Alzheimer’s disease,” *Transl Psychiatry*, 11(1), 2021, 146.
- [22] R. D. Hodge et al., “Conserved cell types with divergent features in human versus mouse cortex,” *Nature*, 573(7772), 2019, 61–8.
- [23] D. M. Holtzman, J. C. Morris, and A. M. Goate, “Alzheimer’s Disease: The Challenge of the Second Century,” *Science Translational Medicine*, 3(77), 2011.
- [24] B. Hwang, J. H. Lee, and D. Bang, “Single-cell RNA sequencing technologies and bioinformatics pipelines,” *Exp Mol Med*, 50(8), 2018, 1–14.

- [25] R. S. Kamathe and K. R. Joshi, "A novel method based on independent component analysis for brain MR image tissue classification into CSF, WM and GM for atrophy detection in Alzheimer's disease," *Biomedical Signal Processing and Control*, 40, 2018, 41–8.
- [26] S. Kiani Shabestari *et al.*, "Absence of microglia promotes diverse pathologies and early lethality in Alzheimer's disease mice," *Cell Rep*, 39(11), 2022, 110961.
- [27] J. Kim *et al.*, "Neuroimaging Modalities in Alzheimer's Disease: Diagnosis and Clinical Features," *International Journal of Molecular Sciences*, 23(11), 2022.
- [28] D. P. B. J. Kingma, "Adam: A method for stochastic optimization," *arXiv preprint*, 2014.
- [29] W. Kipf T N, "Semi-supervised classification with graph convolutional networks," *arXiv preprint*, 2016.
- [30] S. Koglsberger *et al.*, "Gender-Specific Expression of Ubiquitin-Specific Peptidase 9 Modulates Tau Expression and Phosphorylation: Possible Implications for Tauopathies," *Mol Neurobiol*, 54(10), 2017, 7979–93.
- [31] L. Kong *et al.*, "Single-cell transcriptomic profiles reveal changes associated with BCG-induced trained immunity and protective effects in circulating monocytes," *Cell Rep*, 37(7), 2021, 110028.
- [32] N. Kozauer and R. Katz, "Regulatory Innovation and Drug Development for Early-Stage Alzheimer's Disease," *New England Journal of Medicine*, 368(13), 2013, 1169–71.
- [33] S. Lahmirmi and M. Boukadoum, "New approach for automatic classification of Alzheimer's disease, mild cognitive impairment and healthy brain magnetic resonance images," *Healthcare Technology Letters*, 1(1), 2014, 32–6.
- [34] S. Lahmirmi and A. Shmuel, "Performance of machine learning methods applied to structural MRI and ADAS cognitive scores in diagnosing Alzheimer's disease," *Biomedical Signal Processing and Control*, 52, 2019, 414–9.
- [35] D. Lahnemann *et al.*, "Eleven grand challenges in single-cell data science," *Genome Biol*, 21(1), 2020, 31.
- [36] C. Y. Lee *et al.*, "Venus: An efficient virus infection detection and fusion site discovery method using single-cell and bulk RNA-seq data," *Plos Computational Biology*, 18(10), 2022.
- [37] K. Leng *et al.*, "Molecular characterization of selectively vulnerable neurons in Alzheimer's disease," *Nat Neurosci*, 24(2), 2021, 276–87.
- [38] X. Li and C. Y. Wang, "From bulk, single-cell to spatial RNA sequencing," *Int J Oral Sci*, 13(1), 2021, 36.
- [39] J. M. Long and D. M. Holtzman, "Alzheimer Disease: An Update on Pathobiology and Treatment Strategies," *Cell*, 179(2), 2019, 312–39.

- [40] V. Marx, “Method of the Year: spatially resolved transcriptomics,” *Nat Methods*, 18(1), 2021, 9–14.
- [41] H. Mathys et al., “Single-cell transcriptomic analysis of Alzheimer’s disease,” *Nature*, 570(7761), 2019, 332–+.
- [42] E. Mereu et al., “Benchmarking single-cell RNA-sequencing protocols for cell atlas projects,” *Nat Biotechnol*, 38(6), 2020, 747–55.
- [43] S. Miao et al., “Interpretable Geometric Deep Learning via Learnable Randomness Injection,” 2022, arXiv:2210.16966, DOI: [10.48550/arXiv.2210.16966](https://doi.org/10.48550/arXiv.2210.16966).
- [44] S. Miao, M. Liu, and P. J. a. e.-p. Li, “Interpretable and Generalizable Graph Learning via Stochastic Attention Mechanism,” 2022, arXiv:2201.12987, DOI: [10.48550/arXiv.2201.12987](https://doi.org/10.48550/arXiv.2201.12987).
- [45] S. Morabito et al., “Single-nucleus chromatin accessibility and transcriptomic characterization of Alzheimer’s disease,” *Nat Genet*, 53(8), 2021, 1143–55.
- [46] E. Moradi et al., “Machine learning framework for early MRI-based Alzheimer’s conversion prediction in MCI subjects,” *Neuroimage*, 104, 2015, 398–412.
- [47] M. S. Nadeem et al., “Symptomatic, Genetic, and Mechanistic Overlaps between Autism and Alzheimer’s Disease,” *Biomolecules*, 11(11), 2021.
- [48] B. I. C. C. Network, “A multimodal cell census and atlas of the mammalian primary motor cortex,” *Nature*, 598(7879), 2021, 86–102.
- [49] E. Papalexi and R. Satija, “Single-cell RNA sequencing to explore immune cell heterogeneity,” *Nat Rev Immunol*, 18(1), 2018, 35–45.
- [50] M. Paraskevaidi et al., “Raman Spectroscopy to Diagnose Alzheimer’s Disease and Dementia with Lewy Bodies in Blood,” *Acs Chemical Neuroscience*, 9(11), 2018, 2786–94.
- [51] E. Pellegrini et al., “Machine learning of neuroimaging for assisted diagnosis of cognitive impairment and dementia: A systematic review,” *Alzheimers Dement (Amst)*, 10, 2018, 519–35.
- [52] P. Preman et al., “Astrocytes in Alzheimer’s Disease: Pathological Significance and Molecular Pathways,” *Cells*, 10(3), 2021.
- [53] S. Rath et al., “MitoCarta3.0: an updated mitochondrial proteome now with sub-organelle localization and pathway annotations,” *Nucleic Acids Res*, 49(D1), 2021, D1541–D1547.
- [54] J. S. Sadick et al., “Astrocytes and oligodendrocytes undergo subtype-specific transcriptional changes in Alzheimer’s disease,” *Neuron*, 110(11), 2022, 1788–1805 e10.
- [55] H. S. Sharma et al., “The blood-brain barrier in Alzheimer’s disease: novel therapeutic targets and nanodrug delivery,” *Int Rev Neurobiol*, 102, 2012, 47–90.

- [56] H. Shen *et al.*, “Sexually dimorphic RNA helicases DDX3X and DDX3Y differentially regulate RNA metabolism through phase separation,” *Mol Cell*, 82(14), 2022, 2588–2603 e9.
- [57] M. Slyper *et al.*, “A single-cell and single-nucleus RNA-Seq toolbox for fresh and frozen human tumors,” *Nat Med*, 26(5), 2020, 792–802.
- [58] J. Stephen, P. C. M. Fleming, and M. Babadi, “CellBender remove-background: a deep generative model for unsupervised removal of background noise from scRNA-seq datasets,” *bioRxiv*, 2019.
- [59] H. I. Suk *et al.*, “State-space model with deep learning for functional dynamics estimation in resting-state fMRI,” *Neuroimage*, 129, 2016, 292–307.
- [60] M. Tanveer *et al.*, “Machine Learning Techniques for the Diagnosis of Alzheimer’s Disease: A Review,” *Acm Transactions on Multimedia Computing Communications and Applications*, 16(1), 2020.
- [61] M. Termenon *et al.*, “Lattice independent component analysis feature selection on diffusion weighted imaging for Alzheimer’s disease classification,” *Neurocomputing*, 114, 2013, 132–41.
- [62] P. C. G. Veličković, A. Casanova, *et al.*, “Graph attention networks,” *arXiv preprint*, 2017.
- [63] A. J. Vincent *et al.*, “Astrocytes in Alzheimer’s disease: emerging roles in calcium dysregulation and synaptic plasticity,” *J Alzheimers Dis*, 22(3), 2010, 699–714.
- [64] D. Vogrinc, K. Goricar, and V. Dolzan, “Genetic Variability in Molecular Pathways Implicated in Alzheimer’s Disease: A Comprehensive Review,” *Front Aging Neurosci*, 13, 2021, 646901.
- [65] M. H. Wang *et al.*, “Guidelines for bioinformatics of single-cell sequencing data analysis in Alzheimer’s disease: review, recommendation, implementation and application,” *Molecular Neurodegeneration*, 17(1), 2022.
- [66] X. L. Wang and L. J. Li, “Cell type-specific potential pathogenic genes and functional pathways in Alzheimer’s Disease,” *Bmc Neurology*, 21, 2021.
- [67] Y. Y. Wang *et al.*, “Heterogeneous Attributed Network Embedding with Graph Convolutional Networks,” in *Thirty-Third Aaai Conference on Artificial Intelligence/Thirty-First Innovative Applications of Artificial Intelligence Conference / Ninth Aaai Symposium on Educational Advances in Artificial Intelligence*, 2019, 10061–2.
- [68] O. Wirths *et al.*, “Inflammatory changes are tightly associated with neurodegeneration in the brain and spinal cord of the APP/PS1KI mouse model of Alzheimer’s disease,” *Neurobiol Aging*, 31(5), 2010, 747–57.
- [69] S. L. Wolock, R. Lopez, and A. M. Klein, “Scrublet: Computational Identification of Cell Doublets in Single-Cell Transcriptomic Data,” *Cell Syst*, 8(4), 2019, 281–291 e9.

- [70] N. M. Xi and J. J. Li, “Benchmarking Computational Doublet-Detection Methods for Single-Cell RNA Sequencing Data,” *Cell Syst*, 12(2), 2021, 176–194 e6.
- [71] A. T. Yeo *et al.*, “Single-cell RNA sequencing reveals evolution of immune landscape during glioblastoma progression,” *Nat Immunol*, 23(6), 2022, 971–84.
- [72] N. S. S. Zakeri, S. Pashazadeh, and H. MotieGhader, “Gene biomarker discovery at different stages of Alzheimer using gene co-expression network approach,” *Scientific Reports*, 10(1), 2020.
- [73] J. Zhao *et al.*, “Detection of differentially abundant cell subpopulations in scRNA-seq data,” *Proc Natl Acad Sci USA*, 118(22), 2021.
- [74] G. X. Zheng *et al.*, “Massively parallel digital transcriptional profiling of single cells,” *Nat Commun*, 8, 2017, 14049.