Disease State Prediction From Single-Cell Data Using Graph Attention Networks

Neal G. Ravindra^{*†}
Yale University
neal.ravindra@yale.edu

Arijit Sehanobish^{*†}
Yale University
arijit.sehanobish@yale.edu

Jenna L. Pappalardo[‡]
Yale University
jenna.pappalardo@yale.edu

David A. Hafler[‡]
Yale University
david.hafler@yale.edu

David van Dijk[†]
Yale University
david.vandijk@yale.edu

ABSTRACT

Single-cell RNA sequencing (scRNA-seq) has revolutionized biological discovery, providing an unbiased picture of cellular heterogeneity in tissues. While scRNA-seq has been used extensively to provide insight into both healthy systems and diseases, it has not been used for disease prediction or diagnostics. Graph Attention Networks (GAT) have proven to be versatile for a wide range of tasks by learning from both original features and graph structures. Here we present a graph attention model for predicting disease state from single-cell data on a large dataset of Multiple Sclerosis (MS) patients. MS is a disease of the central nervous system that can be difficult to diagnose. We train our model on single-cell data obtained from blood and cerebrospinal fluid (CSF) for a cohort of seven MS patients and six healthy adults (HA), resulting in 66,667 individual cells. We achieve 92 % accuracy in predicting MS, outperforming other state-of-the-art methods such as a graph convolutional network and a random forest classifier. Further, we use the learned graph attention model to get insight into the features (cell types and genes) that are important for this prediction. The graph attention model also allow us to infer a new feature space for the cells that emphasizes the differences between the two conditions. Finally we use the attention weights to learn a new low-dimensional embedding that can be visualized. To the best of our knowledge, this is the first effort to use graph attention, and deep learning in general, to predict disease state from single-cell data. We envision applying this method to single-cell data for other diseases.

^{*}Both authors contributed equally to this research.

[†]Internal Medicine(Cardiology), Computer Science

[‡]Neurology, Immunobiology

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

ACM CHIL '20, April 2–4, 2020, Toronto, ON, Canada

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7046-2/20/04.

CCS CONCEPTS

• **Applied computing** → **Computational transcriptomics**;
• **Computing methodologies** → **Neural networks**; Feature selection.

KEYWORDS

graph attention networks, single cell RNA-seq, disease prediction, multiple sclerosis

ACM Reference Format:

Neal G. Ravindra, Arijit Sehanobish, Jenna L. Pappalardo, David A. Hafler, and David van Dijk. 2020. Disease State Prediction From Single-Cell Data Using Graph Attention Networks. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '20)*, April 2–4, 2020, Toronto, ON, Canada. ACM, New York, NY, USA, 10 pages.

1 INTRODUCTION

Single cell RNA-sequencing (scRNA-seq) technology provides gene expression data for individual cells, yielding omics-scale information at single-cell resolution. The information contained within transcriptomes of individual cells can be used to identify rare cells, genes, and other changes associated with fundamental biological processes and pathological states. However, identifying hidden differences in cell pathophysiological trajectories is challenging. ScRNA-seq data typically contains gene count data for more than 20,000 transcripts across 10,000 - 100,000 cells and has a high-degree of technical noise, and encompasses a large degree of biological variability [8, 24]. While these data promise to provide rich insight into complex systems that determine health and well-being, integrative computational tools are required to comprehensively learn about cell states from scRNA-seq data collected across cell, disease, and temporal modalities [41].

Computational approaches have been developed to use scRNA-seq data to identify bio-markers for potential use in diagnostics [15], to understand genetic factors underlying disease heterogeneity [53], and to predict cellular response to experimental perturbations [29]. In addition, several classifiers, using primarily non-deep learning methods, have been employed to predict cell-types from single-cell samples [1, 55]. However, to the best of our knowledge, no computational approaches have been developed to harness scRNA-seq data to predict disease states for the use of transcriptomic technology in diagnosis [37]. Despite this gap, scRNA-seq data may revolutionize our understanding of the biological processes associated with various disease states in addition to providing personalized insight

into a patient's pathophysiological state. These kinds of omics-scale insights may be useful in tailoring treatments for personalized medicine [22]. We propose a learning approach that harnesses the rich-information content in a scRNA-seq dataset to characterize and predict two pathological states related to neuroinflammation.

Multiple sclerosis (MS) is a disease of the nervous system in which myelin, the insulating substance that surrounds nerve cell axons, is damaged by the body's immune system [30, 42]. This damage causes individuals with MS to have neurological symptoms such as difficulty with coordination, vision impairment, pain, and fatigue, which can manifest acutely in episodic periods or lead to progressive decline of neurological function [30]. Genome-wide association studies have identified some genetic risk-factors associated with MS and highlighted the role of the immune system in the disease but to the best of our knowledge, no one has used transcriptomic data and modeling to study the association of transcriptomic markers with MS [18, 36, 38]. Lack of molecular mechanistic understanding of the pathology of MS caused by an aberrant immune system complicates the ability in diagnosing and characterizing MS; indeed, there is no single test to diagnose MS [28]. Few molecular MS markers are known [11]. However, current indicators of MS are related to general neurological decline rather than the root causes of MS [11, 28]. Given the central role of the immune system in MS, we relied on a dataset composed of cerebrospinal fluid (CSF) and blood, both of which contain a high-proportion of immune cells. The selection of an immune-rich single-cell sample is critical to identifying molecular markers that can be used to predict MS disease-state and to aid in developing a diagnostic test based on single-cell transcriptomic technology.

The existence of disease-modifying treatments highlights the value of this information-rich scRNA-seq data, as it can be used to characterize personalized disease-states, which may be differentially targeted with personalized treatment [28]. Current diagnoses of MS relies on imaging, neurological symptoms observed by a clinician and differential diagnoses, and spinal fluid abnormalities [42]. Currently, personalized approaches using a genetic risk score, based on genome-wide association studies, have poor agreement with clinical and MRI measures [28, 34, 36].

Several classes of bio-molecular networks like transcriptional regulatory networks, protein-protein interaction networks, and metabolic networks are modeled as graphs. Mathematical graph theory is a straightforward way to represent this information, and graph-based models can exploit the global and local characteristics of these networks relevant to cell biology. In recent years, Graph Neural Networks (GNNs) have been widely adopted in various tasks, such as graph classification [13, 26], link prediction [54] and node classification [23, 45, 49]. Various GNNs have been proven to be effective in achieving state-of-the-art performance in a variety of graph datasets such as social networks [49], citation networks [45], recommender systems [23, 49] and protein-protein interaction networks [45]. The underlying graph structure is utilized by GNNs to operate convolution directly on graphs by passing node features to neighbors [49], or to perform convolution in the spectral domain by using the eigenbasis of the Graph Laplacian operator [23]. Since the introduction of the self-attention mechanism in the influential work [44], it has become extremely common to use self-attention for various sequence-sequence tasks like machine translations and

learning representations. Self-attention allows the model to deal with variable sized inputs and focus on the most important part of the inputs. The self-attention mechanism has been introduced to the graph domains by [45] and also developed by [5, 10, 52]. In our work, we chose to utilize the Graph Attention Network (GAT) as introduced by [45] to do *node classification* because of its high performance on complex datasets, its adaptability to unseen datasets, and the model's interpretability. Figure 1 shows the workflow of our paper.

Figure 1: Schematic of our method depicting the use of a cell by gene count matrix in a graph attention model to predict and visualize disease state.

Our contributions are as follows: (i) we use graph neural networks on single cell data to predict disease state of MS patients and healthy individuals, (ii) we use the attention weights matrix to visualize the new graphs learned by the model, (iii) we extract meaningful features from the model and investigate them in the context of existing biological knowledge on MS.

2 DATASETS USED

We used a pooled dataset composed of blood and cerebrospinal fluid (CSF) samples from 7 MS patients and 6 healthy adults (HA), constituting 26 scRNA-seq samples with an average of 2,500 cells per sample and about 5,000 cells per patient. We performed standard scRNA-seq quality-control filtering steps to the raw data, including removing cells expressing $> 10\%$ mitochondrial genes, normalizing gene counts by the sequenced library size per cell, square-root transforming gene counts, and removing mitochondrial and ribosomal protein transcripts from downstream analysis [48]. After the preprocessing steps our dataset is composed of 66,667 cells with 22,005 gene features per cell. Figure 2(A) shows our full graph (used in *transductive* task) colored by the individual patients and figure 2(B) is colored by various cell types. Variability between samples dominates biological variability in reduced dimensional embeddings and these batch effects complicate model inference of biological features [6]. We treat each cell as a node, so our dataset has 22,005 features per node, with high variability across features. The number of features is an order of magnitude higher than the citation, social-network, and protein-protein interaction datasets previously used by GNNs. This feature matrix was used to compute a k-nearest neighbor graph from euclidean distances in PCA reduced space (10 neighbors and 100 principal components) and clustered by the Louvain algorithm to identify cell types within the pooled dataset [4, 32]. This led to the unsupervised identification of at least 16 cell types within the pooled dataset, including about 75% T cells broken into various lineages and about 15% B

cells, in addition to other immune and peripheral blood mononuclear cells, including macrophages, monocytes, natural-killer cells, and platelets. The nodes are classified as healthy or MS. Thus the classification problem we are interested in is at the *cell level*.

Figure 2: Low-dimensional embeddings of single-cell data used in this paper. A) UMAP scatter plot of the single-cell data used for training our model colored by the 6 healthy individuals (HA*) and 7 MS patients (MS*). Gray lines show edges in a nearest-neighbor graph. B) UMAP scatter plot colored by cell types identified in the blood and CSF.

We use this dataset to perform two tasks: (i) inductive inference, and (ii) transductive inference. For the *transduction* task, we randomly assign 10% of the nodes for validation and 10% for testing while keeping the ratio of healthy and MS cells the same as in the full dataset. Then we randomly mask the labels of 50% of the training, validation and the test set. Thus we can think of our problem as having a large graph where half of the graph is unlabeled and our goal is to predict the labels based on the labels of the nodes in the graph and the features of the node. With the transductive setup in mind, we feed in the features of *all* nodes during training and our goal is to predict the masked labels.

For the *induction* task, we randomly choose a healthy adult and a MS patient each for the validation and test sets and the remaining 5 MS patients and 4 healthy adults for training. We construct separate

knn-graphs for the training, validation, and test sets just like before. It is important to note that the testing graphs are completely unseen during training. Table 1 gives a complete description of the datasets used in various experiments.

Table 1: Characteristics of our single cell dataset

Task		Train	Dev	Test
Inductive	# Nodes	43866	9686	13033
	# Edges	332398	73552	100715
	# Features	22005	22005	22005
	# Classes	2	2	2
	# Graphs	1	1	1
Transductive	# Nodes	54000	6000	6667
	# Features	22005	22005	22005
	# Classes	2	2	2
	# Edges	5007093		

2.1 Related work

Machine learning, and particularly deep learning have been widely used in the medical domain to predict diseases but mostly using medical images or EHR data. Machine learning has also been used on scRNA-seq data for data de-noising, batch correction, data imputation to correct for sparse signal detection or dropout, unsupervised clustering, cell-type prediction [1, 3, 24, 43] and low-dimensional visualization of samples relevant to human health [2, 43]. Recently, transfer learning (using an autoencoder) has been used to extract new gene-gene interactions [47] and random walks have been used to predict gene responses to drugs [20]. In the biomedical domain GNNs have mostly been used for prediction of protein-protein interactions [45] and gene-gene interactions [20, 56]. However, to the best of our knowledge, predicting diseases from single-cell RNA seq data using state-of-the-art methods have not been tried successfully.

3 ARCHITECTURE OF OUR MODEL

We use the Graph Attention Network by [45] and we follow the exposition in [45]. In this section we describe our model in detail. The building block of our network is the graph attention layer. The input to our layer is a set of node features, $\mathbf{h} = \{h_1, h_2, \dots, h_N\}$, where $h_i \in \mathbb{R}^F$, N is the number of nodes, and F is the number of features in each node. The layer produces a new set of node features (of possibly different cardinality F') as its output, $\mathbf{h}' = \{h'_1, h'_2, \dots, h'_N\}$ where $h'_i \in \mathbb{R}^{F'}$.

We apply a linear transformation called the weight matrix, $\mathbb{W} \in \mathbb{R}^{F' \times F}$ to every node. The weight matrix $\mathbb{W}$ is initialized with Glorot initialization [17]. Self-attention is then computed on the nodes, i.e. a shared attention mechanism

$$a : \mathbb{R}^{F'} \times \mathbb{R}^{F'} \rightarrow \mathbb{R} \quad (1)$$

that computes attention coefficients

$$e_{ij} = a(\mathbb{W}h_i, \mathbb{W}h_j) \quad (2)$$

We treat these e_{ij} 's as importance scores of node j 's features to that of node i . This is the main difference between a graph attention network and a graph convolution network [23]. Unlike [23],

different importance scores are given to different nodes in the same neighborhood, which allows for greater adaptability of the model to various complicated datasets like ours.

Furthermore, we normalize these e_{ij} 's across all choices of j , where j lies in some neighborhood of i using the softmax function:

$$\alpha_{ij} = \text{softmax}_j(e_{ij}) = \frac{\exp(e_{ij})}{\sum_{k \in \mathcal{N}_i} \exp(e_{ik})} \quad (3)$$

where $\mathcal{N}_i$ is some neighborhood of the node i . In all our applications, the graphs are the KNN graphs and $\mathcal{N}_i$ is the first order k -neighbors of the node i . Thus in practice, we only compute e_{ij} (and thus α_{ij}) for all nodes j in the first-order neighborhood of the node i , instead of all pairs of nodes i and j .

In our experiments, the attention mechanism a is a single-layer feedforward neural network on which we apply the LeakyReLU non-linearity. The normalized attention coefficients are then used to compute a linear combination of the features corresponding to them, to serve as the final output features for every node (after applying a nonlinearity, σ):

$$h'_i = \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbb{W} h_j \right). \quad (4)$$

As is now common in attention based networks, we use multi-head attention, similar to [44, 45]. In this case, each attention head learns a unique set of attention weights independent of the other heads in a given layer. Specifically, K independent attention mechanisms execute the transformation of Equation 4, and then their features are concatenated, resulting in the following output feature representation:

$$h'_i = \left\|_{l=1}^K \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^l \mathbb{W}^l h_j \right) \right\|. \quad (5)$$

where $\|$ denotes the concatenation operation, α_{ij}^l are normalized attention coefficients computed by the l -th attention mechanism and $\mathbb{W}^l$ is the corresponding input linear transformation's weight matrix. Note that, in the final output, each node will have KF' features.

Finally, in the final (prediction) layer of the network, we first employ averaging, and then apply the final nonlinearity (logistic sigmoid for our classification problems). Thus the equation for the final layer is:

$$h'_i = \sigma \left(\frac{1}{K} \sum_{l=1}^K \sum_{j \in \mathcal{N}_i} \alpha_{ij}^l \mathbb{W}^l h_j \right). \quad (6)$$

We also use weight clipping to prevent the gradients from exploding. Table 2 shows all the hyperparameters in our models used for the experiments. Our code is written in PyTorch and is publicly available at <https://github.com/vandijklab/scGAT>.

Another reason of working with Graph Attention Networks is its power of interpretability. It is easy to visualize the attention heads, i.e. the attention scores between various nodes by a head in any given layer. Due to memory issues, we implemented a sparse version of the code of [45].

4 EXPERIMENTS

In this section we will explain in detail our experiments. We compare our work with random forests, a multilayer perceptrons, Graph Convolution Networks as introduced in [23], and a dummy estimator. Table 3 gives complete details of our results for the following tasks. All reported numbers are from the test sets. We also experimented with changing the default hyperparameters like changing the width and breadth of the model, the learning rate, optimizers, losses, and the dropout rate. However the best results were obtained by our default hyperparameters.

Generally all deep learning models are trained by stochastic gradient descent which requires us to minibatch our graphs. To breakdown a large graph for memory purposes and fast efficient training, we follow the algorithm in [9] which is implemented in PyTorch Geometric. For the *inductive* task, we break up our training graph into 4000 roughly equal subgraphs and use a batch size of 256 graphs. The average number of nodes in each batch is 2741. But we do not split the test or the validation graphs. For the *transductive* task, we break our graph as before in 4000 subgraphs and use a batch size of 256 subgraphs. The average number of nodes per subgraph in a batch is 4657.

Table 2: Default hyperparameters used in the experiments

	Graph Attention Network	Graph Convolution Network
Number of layers	2	2
Hidden_size	8	256
Attention Heads	8	N/A
Optimizer	Adagrad	Adagrad
clip_grad_norm	5	2
learning rate scheduler	cosine_decay & 1st decay steps = 1000	cosine_decay & 1st decay steps = 1000
weight_decay	.0005	.005
Batch size	256	256
Dropout	.5	.4
Slope in LeakyRelu	.2	.2
Training Epochs	1000	1000
Early stopping	100	100

4.1 Transductive learning

We use our single cell dataset and we follow the transductive learning setup as in [49]. The transductive setup implies that during training the model has access to all the nodes' feature vectors. The predictive power of the model is then evaluated on the masked test nodes and we use masked validation nodes for validation purposes. Our main task is to predict if a cell is healthy or has MS. We ran our GAT model 8 times.

Table 3: Experimental results

Task	Model	Accuracy
Inductive	Random	51.8
	MLP	56.7
	Random Forest	58.5
	Graph Convolutional Network	72.1
	Graph Attention Network(our)	92.3 ± .7
Transductive	Graph Convolutional Network	82.91
	Graph Attention Network(our)	86 ± .3

Figure 3: Predicted probabilities from induction task per cell. A) PHATE plot of the data colored by ground truth labels. B) PHATE plot colored by the predictions of our graph attention model. C) Model prediction probability (Y-axis) for indicated ground truth label (X-axis).

4.2 Inductive learning

For this task we again use our single cell dataset. We have 13 persons' blood and CSF data and out of them 6 of them have MS and the rest are healthy. We randomly choose a blood and CSF data of a healthy adult and a MS patient each for the validation and test sets and the remaining 5 MS patients and 4 healthy adults for training. We construct separate knn-graphs for the training, validation, and the test sets. We ran our GAT model 16 times.

Figure 4 shows the loss of our model for the induction and the transduction tasks.

Figure 4: Negative log-likelihood losses for training and validation sets in the induction (left) and transduction (right) tasks.

5 DISCUSSION OF OUR RESULTS

Encouraged by the predictive ability of our model in both transductive and inductive tasks, we take a deeper look at the attention heads and attention weights at various layers of our model. In this section, we will extract meaningful biological information from the attention heads.

5.1 Finding interesting gene markers

We look at $\mathbb{W}_h^T$ where T stands for transpose and $\mathbb{W}_h$ stands for the learned weight matrix for head h in layer 1. $\mathbb{W}_h$ in our case is an 8×22005 matrix where 8 is the size of the hidden dimension. For each row i of each $\mathbb{W}_h^T$, we take the max of the absolute value and we define

$$g_i^h = \max_j(|w_{ij}|) \quad (7)$$

and then we pick the highest 10 g_i^h for each h . We can think of these genes as the most important features the model learns to discriminate between healthy and MS cells.

A heatmap of these features for the model trained in the inductive task is shown in figure 5. We are encouraged to find **IL2RG** as a learned feature in discriminating between healthy and MS cells. The interleukin-2 receptor (IL2R) is involved in a signaling pathway that is essential for T cell function and genetic variants of IL-2R are associated with various autoimmune disorders, including MS [7, 18]. In addition, we are encouraged that the network gives a high weight to **CD19**, a marker for B cells, since B cells are targeted by therapeutics and is thought to contribute to MS by secreting autoreactive antibodies that cause neuroinflammation and demyelination [21]. Gene Ontology of the other identified features revealed that 38 out of 58 mapped genes are involved in the regulation of the hormone secretion while a number of other genes in our list are involved in the regulation of the nerve cell development. Lipid metabolism and nerve cell development are critical to normal myelination, so alterations to these processes may lead to a triggering of the autoimmune system and cause inflammation related to development of MS pathologies [14, 42].

Thus it is encouraging to see that representations learned by our model align with the medical knowledge. Given the paucity of molecular markers that can characterize MS disease states, other than genes identified through genome-wide association studies and metabolites indicative of neurological degeneration, we presume that our model and this novel list of gene features will guide further

experimental work into the pathophysiology of MS [28]. Broadly, this feature selection approach promises to reveal biomarkers of MS disease pathologies, which may be used for more cost-effective diagnostic testing than the single-cell technologies affords like DNA microarray screening or measuring gene expression by flow cytometry [14].

5.2 Finding cell types most prototypical of MS

In both the tasks, we computed the top 100 cells that have the highest predicted probability of being MS. These are the cells that are strongly indicative of MS. We noted that all of these cells are from patients' CSF. In contrast, the 100 most prototypical healthy cells are approximately evenly split between being sourced from blood or CSF samples. This confirms the common medical hypothesis that patients' CSF has a higher degree of information relevant to central nervous system inflammation than blood cell samples [31]. In addition, these observations support that single-cell technology may supplement existing neuroinflammation diagnostic criteria relying on detection of spinal fluid abnormalities [28, 31].

We assess the predicted probability of being MS, stratified by cell type in figure 6A. Cells involved in innate immunity, such as macrophages and natural-killer cells, have a relatively lower average predicted probability of being MS compared to cells involved in adaptive immunity, such as B cells and T cells. This reflects the current medical understanding of the immune system's role in MS and the key role that adaptive immunity plays in autoimmune disorders and neuroinflammation more broadly [16, 28, 39]. This suggests that our model is tuned to recapitulate biological knowledge.

We assess the predicted probability of being MS, stratified by patient in figure 6B. MS is known to present in different forms and the severity of the disease can vary across patients [30]. Our model performs a node classification task, yielding the predicted probability of being MS for every cell drawn from a single patient. Thus, a single patient can have cells that are more prototypically healthy than MS. The average predicted probability is as low as 0.23 for healthy patients and as high as 0.88 for MS patients in the transduction task. Both the transduction and induction tasks have models that achieve good discriminability between disease-states, with a Pearson's correlation coefficient of 0.76 and 0.82, respectively.

5.3 Visualizing Attention heads

In this subsection, we will take a deeper look at the features that our GAT model learned and visualize the latent space using [32] and [48]. We can construct a matrix whose (i, j) entry is α_{ij} as defined by equation 3 for each head in the first layer. We can think of these attention matrices as normalized adjacency matrix of a directed graph (i.e. the stochastic matrix one obtains from normalizing the adjacency matrix of a graph).

And we can feed it into PHATE (figure 7), allowing us to visualize the new graphs that our model has learned. We can easily see that head 4 and head 8 learn very different graphs. On the other hand, we can look at the learned $\mathbf{h}'$ as in equation 4 by each head in each layer which we view as the new features or the latent space. Just as before, we can feed these new features into a manifold learning algorithm and can construct new KNN graphs (with 10

Figure 5: Heatmap showing the top 10 gene features for each head colored by learned weights in the induction task and clustered by gene.

neighbors) for visualization with UMAP, as shown in the figure 8 for a 64 dimensional feature space constructed by concatenating the hidden units for each each head per node. Here we see that

Figure 6: Predicted probabilities from transduction task for MS class label stratified by cell type and individual patient. Violin plot shows kernel density estimate and inner box plot shows interquartile range and median. A) Probabilities of having MS predicted by our model for each cell type. B) Probabilities of having MS predicted by our model for each patient.

the learned latent space separates healthy and MS cells in the low-dimensional embedding. This further supports that our model learns to discriminate between healthy and MS labels.

6 FUTURE WORK

We are excited by the improvements that a Graph Attention Networks bring to single cell analysis. We would like to continue this direction further by incorporating in multidimensional edge level features. Another direction of our research is to come up with new semi/unsupervised ways of creating new graphs from single cell data by completely bypassing PCA and KNN graphs. And then we can apply the models in this paper along with the edge level modifications to our new graphs.

We are going to scale up our work to a much larger dataset which requires us not only to experiment with the depth and the breadth of our model but also introduces new computational challenges like fitting large graphs in memory and doing computations more effectively.

One of the striking results we found that conforms to the medical knowledge is that the CSF cells are strong indicators of MS, whereas the blood cell are not very good indicators. However getting the CSF data is an expensive and invasive process. We are in process of acquiring only the blood samples of MS patients. Thus predicting

Figure 7: New low-dimensional embeddings learned by the PHATE algorithm after inputting graphs learned by various attention heads in our model for the induction task.

Figure 8: Visualization of the latent space learned by our Graph Attention Model during the induction task. The UMAP plot of the original graph (left) and the UMAP plot of the latent space (right).

MS from only blood samples becomes an extremely difficult task. To alleviate the problem, we would like to do **transfer learning** from our model that is trained on both blood and CSF samples and the train/test **only** on the new blood samples. We hope that the transfer learning will achieve better results than an untrained graph neural network.

7 CONCLUSION

Multiple sclerosis (MS) is an unpredictable disease of the central nervous system that disrupts the flow of information within the brain, and between the brain and body. The molecular markers and characteristics of MS are still actively studied and there remains no single test for diagnosing MS. Single-cell sequencing has revolutionized biological discovery, providing an in-depth picture of cellular differences and the interactions between various cells in their microenvironments. However, scRNA-seq has mostly been used for basic sciences but not for clinical or diagnostic use due to its cost and the complexity of the data. In this work we use a graph attention model on single cell data to predict MS. The model is trained on a cohort of seven MS patients and six healthy adults (HA) from the blood and cerebral spinal fluid (CSF) of each patient. We perform transductive and inductive tasks and we achieve 86 % and 92 % accuracy in predicting MS beating other methods such as graph convolutional networks, random forest and MLP. Further, we use the learned graph attention model to get insight into the signatures (transcriptomic and genetic features) and the cells that are important for this prediction. Using the attention weights we extracted a list of 100 genes and we were encouraged to find **CD19** and **IL2RG** in that list. We envision that the feature selection approach to produce a minimal gene set can function as biomarkers in a clinical setting. We also computed the top 100 cells that have the highest predicted probability of being MS and all of these cells came from the patients' CSF. In contrast, the 100 most prototypical healthy cells are approximately evenly split between being sourced from blood or CSF samples. This confirms the common medical hypothesis that patients' CSF has a higher degree of information relevant to central nervous system inflammation than blood cell samples. Finally the graph attention model also allow us to infer a new feature space for the cells that emphasizes the difference between the two conditions which we visualize via PHATE and UMAP. There has been more emphasis on the identification of determinants of disease progression and on how individual information can be used to personalize treatment and we hope that this graph neural networks will play in important role in predicting disease and in identifying key determinant of those diseases.

REFERENCES

- [1] Jose Alquicira-Hernandez, Anuja Sathe, Hanlee P Ji, Quan Nguyen, and Joseph E Powell. 2019. scPred: accurate supervised method for cell-type classification from single-cell RNA-seq data. *Genome Biology* 20, 1 (2019), 264. <https://doi.org/10.1186/s13059-019-1862-5>
- [2] Matthew Amodio, David van Dijk, Krishnan Srinivasan, William S Chen, Hussein Mohsen, Kevin R Moon, Allison Campbell, Yujiao Zhao, Xiaomei Wang, Manjunatha Venkataswamy, Anita Desai, V Ravi, Priti Kumar, Ruth Montgomery, Guy Wolf, and Smita Krishnaswamy. 2019. Exploring single-cell data with deep multitasking neural networks. *Nature Methods* 16, 11 (2019), 1139–1145. <https://doi.org/10.1038/s41592-019-0576-7>
- [3] Cédric Arisdakessian, Olivier Poirion, Breck Yunits, Xun Zhu, and Lana X Garmire. 2019. DeepImpute: an accurate, fast, and scalable deep neural network method to impute single-cell RNA-seq data. *Genome Biology* 20, 1 (2019), 211. <https://doi.org/10.1186/s13059-019-1837-6>
- [4] Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (Oct 2008), P10008. <https://doi.org/10.1088/1742-5468/2008/10/p10008>
- [5] Léon Bottou and Yann LeCun. 2005. Graph transformer networks for image recognition. <http://yann.lecun.com/exdb/publis/pdf/bottou-05.pdf>.
- [6] Maren Büttner, Zhichao Miao, F. Alexander Wolf, Sarah A. Teichmann, and Fabian J. Theis. 2019. A test metric for assessing single-cell RNA-seq batch correction. *Nature Methods* 16, 1 (Jan 2019), 43–49. <https://doi.org/10.1038/s41592-018-0254-1>
- [7] Karen Cerosaletti, Anya Schneider, Katharine Schwedhelm, Ian Frank, Megan Tatum, Shan Wei, Elizabeth Whalen, Carla Greenbaum, Mariko Kita, Jane Buckner, and S. Alice Long. 2013. Multiple Autoimmune-Associated Variants Confer Decreased IL-2R Signaling in CD4+CD25hi T Cells of Type 1 Diabetic and Multiple Sclerosis Patients. *Plos* 8, 12 (12 2013), 1–9. <https://doi.org/10.1371/journal.pone.0083811>
- [8] Geng Chen, Baitang Ning, and Tielu Shi. 2019. Single-Cell RNA-Seq Technologies and Related Computational Data Analysis. *Frontiers in Genetics* 10 (2019), 317. <https://doi.org/10.3389/fgene.2019.00317>
- [9] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. 2019. Cluster-GCN: An Efficient Algorithm for Training Deep and Large Graph Convolutional Networks. , 10 pages. <https://doi.org/10.1145/3292500.3330925>
- [10] Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F. Stewart, and Jimeng Sun. 2017. GRAM: Graph-Based Attention Model for Healthcare Representation Learning (KDD '17). Association for Computing Machinery, New York, NY, USA, 787–795. <https://doi.org/10.1145/3097983.3098126>
- [11] Olga Ciccarelli, Frederik Barkhof, Benedetta Bodini, Nicola De Stefano, Xavier Go-Lay, Klaas Nicolay, Daniel Pelletier, Petra J.W. Pouwels, Seth A. Smith, Claudia A.M. Wheeler-Kingshott, Bruno Stankoff, Tarek Youssry, and David H. Miller. 2014. Pathogenesis of multiple sclerosis: Insights from molecular and metabolic imaging. *The Lancet Neurology* 13, 8 (2014), 807–822. [https://doi.org/10.1016/S1474-4422\(14\)70101-2](https://doi.org/10.1016/S1474-4422(14)70101-2)
- [12] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2016. Fast and accurate deep network learning by exponential linear units (elus). <https://arxiv.org/abs/1511.07289>
- [13] David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. 2015. Convolutional Networks on Graphs for Learning Molecular Fingerprints (NIPS'15). MIT Press, Cambridge, MA, USA, 2224–2232.
- [14] Arie Gafson, Matt J Craner, and Paul M Matthews. 2017. Personalised medicine for multiple sclerosis care. *Multiple Sclerosis Journal* 23, 3 (2017), 362–369. <https://doi.org/10.1177/1352458516672017>
- [15] Danuta R Gawel, Jordi Serra-Musach, Sandra Lilja, Jesper Aagesen, Alex Arenas, Bengt Asking, Malin Bengtner, Janne Björkander, Sophie Biggs, Jan Ernerudh, Henrik Hjortswang, Jan-Erik Karlsson, Mattias Köpsen, Eun Jung Lee, Antonio Lentini, Xinxu Li, Mattias Magnusson, David Martinez-Enguita, Andreas Matussek, Colm E Nestor, Samuel Schäfer, Oliver Seifert, Ceylan Sonmez, Henrik Stjernman, Andreas Tjärnberg, Simon Wu, Karin Åkesson, Alex K Shalek, Margaretha Stenmarker, Huan Zhang, Mika Gustafsson, and Mikael Benson. 2019. A validated single-cell-based strategy to identify diagnostic and therapeutic targets in complex diseases. *Genome Medicine* 11, 1 (2019), 47. <https://doi.org/10.1186/s13073-019-0657-3>
- [16] F. W. Gay, T. J. Drye, G. W.A. Dick, and M. M. Esiri. 1997. The application of multifactorial cluster analysis in the staging of plaques in early multiple sclerosis: Identification and characterization of the primary demyelinating lesion. *Brain* 120, 8 (Aug 1997), 1461–1483. <https://doi.org/10.1093/brain/120.8.1461>
- [17] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. , 249–256 pages.
- [18] David A. Hafler, Alastair Compston, Stephen Sawcer, Eric S. Lander, Mark J. Daly, Philip L. De Jager, Paul I.W. De Bakker, Stacey B. Gabriel, Daniel B. Mirel, Adrian J. Ivins, Margaret A. Pericak-Vance, Simon G. Gregory, John D. Rioux, Jacob L. McCauley, Jonathan L. Haines, Lisa F. Barcellos, Bruce Cree, Jorge R. Oksenberg, and Stephen L. Hauser. 2007. Risk alleles for multiple sclerosis identified by a genomewide study. *New England Journal of Medicine* 357, 9 (Aug 2007), 851–862. <https://doi.org/10.1056/NEJMoa073493>
- [19] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 1025–1035.
- [20] Haripriya Harikumar, Thomas P. Quinn, Santu Rana, Sunil Gupta, and Svetha Venkatesh. 2019. A random walk down personalized single-cell networks: predicting the response of any gene to any drug for any patient. <https://doi.org/10.1101/837807> arXiv:https://www.biorxiv.org/content/early/2019/11/11/837807.full.pdf
- [21] Stephen L. Hauser, Amit Bar-Or, Giancarlo Comi, Gavin Giovannoni, Hans-Peter Hartung, Bernhard Hemmer, Fred Lublin, Xavier Montalban, Kottil W. Rammohan, Krzysztof Selmaj, Anthony Traboulsee, Jerry S. Wolinsky, Douglas L. Arnold, Gaëlle Klingelschmitt, Donna Masterman, Paulo Fontoura, Shibeshih Belachew, Peter Chin, Nicole Mairon, Hideki Garren, and Ludwig Kappos. 2017. Ocrelizumab versus Interferon Beta-1a in Relapsing Multiple Sclerosis. *New England Journal of Medicine* 376, 3 (2017), 221–234. <https://doi.org/10.1056/NEJMoa1601277> PMID: 28002679.
- [22] Konrad J. Karczewski and Michael P. Snyder. 2018. Integrative omics for health and disease. *Nature Reviews Genetics* 19, 5 (May 2018), 299–310. <https://doi.org/10.1038/nrg.2018.4>
- [23] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. <https://openreview.net/pdf?id=SJU4ayYgl>

- [24] Vladimir Yu Kiselev, Tallulah S. Andrews, and Martin Hemberg. 2019. Challenges in unsupervised clustering of single-cell RNA-seq data. *Nature Reviews Genetics* 20, 5 (May 2019), 273–282. <https://doi.org/10.1038/s41576-018-0088-9>
- [25] Alexandra Kutzelnigg, Claudia F. Lucchinetti, Christine Stadelmann, Wolfgang Brück, Helmut Rauschka, Markus Bergmann, Manfred Schmidbauer, Joseph E. Parisi, and Hans Lassmann. 2005. Cortical demyelination and diffuse white matter injury in multiple sclerosis. *Brain* 128, 11 (10 2005), 2705–2712. <https://doi.org/10.1093/brain/awh641> arXiv:<https://academic.oup.com/brain/article-pdf/128/11/2705/971736/awh641.pdf>
- [26] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. 2019. Self-Attention Graph Pooling. arXiv:1904.08082 pages. arXiv:cs.LG/1904.08082
- [27] John Boaz Lee, Ryan A. Rossi, Sunghul Kim, Nesreen K. Ahmed, and Eunye Koh. 2019. Attention Models in Graphs: A Survey. *ACM Trans. Knowl. Discov. Data* 13, 6, Article Article 62 (Nov. 2019), 25 pages. <https://doi.org/10.1145/3363574>
- [28] Johannes Lorscheider, Vilija G. Jokubaitis, Tim Spelman, Guillermo Izquierdo, Alessandra Lugaes, Eva Havrdova, Dana Horakova, Maria Trojano, Pierre Duquette, Marc Girard, Alexandre Prat, François Grand'Maison, Pierre Grammond, Eugenio Pucci, Cavit Boz, Patrizia Sola, Diana Ferraro, Daniele Spitaleri, Jeanette Lechner-Scott, Murat Terzi, Vincent Van Pesch, Gerardo Iuliano, Roberto Bergamaschi, Cristina Ramo-Tello, Franco Granella, Celia Oreja-Guevara, Helmut Butzkueven, and Tomas Kalincik. 2017. Anti-inflammatory disease-modifying treatment and short-term disability progression in SPMS. *Neurology* 89, 10 (2017), 1050–1059. <https://doi.org/10.1212/WNL.0000000000004330> arXiv:<https://n.neurology.org/content/89/10/1050.full.pdf>
- [29] Mohammad Lotfollahi, F Alexander Wolf, and Fabian J Theis. 2019. scGen predicts single-cell perturbation responses. *Nature Methods* 16, 8 (2019), 715–721. <https://doi.org/10.1038/s41592-019-0494-8>
- [30] Fred D. Lublin, Stephen C. Reingold, and . 1996. Defining the clinical course of multiple sclerosis. *Neurology* 46, 4 (1996), 907–911. <https://doi.org/10.1212/WNL.46.4.907> arXiv:<https://n.neurology.org/content/46/4/907.full.pdf>
- [31] Jan Lycke and Henrik Zetterberg. 2017. The role of blood and CSF biomarkers in the evaluation of new treatments against multiple sclerosis. *Expert Review of Clinical Immunology* 13, 12 (2017), 1143–1153. <https://doi.org/10.1080/1744666X.2017.1400380>
- [32] Leland McInnes, John Healy, and James Melville. 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv:arXiv:1802.03426
- [33] Kevin R. Moon, David van Dijk, Zheng Wang, Scott Gigante, Daniel B. Burkhardt, William S. Chen, Kristina Yim, Antonia van den Elzen, Matthew J. Hirn, Ronald R. Coifman, Natalia B. Ivanova, Guy Wolf, and Smita Krishnaswamy. 2019. Visualizing Structure and Transitions for Biological Data Exploration. <https://www.biorxiv.org/content/early/2019/04/04/120378.full.pdf>. <https://doi.org/10.1101/120378>
- [34] Mark Mühlau, Till F. M. Andlauer, and Bernhard Hemmer. 2016. HLA Genetic Risk Burden in Multiple Sclerosis. *JAMA Neurology* 73, 12 (2016), 1500–1501. <https://doi.org/10.1001/jamaneurol.2016.4329>
- [35] Tomas Olsson, Lisa F. Barcellos, and Lars Alfredsson. 2017. Interactions between genetic, lifestyle and environmental risk factors for multiple sclerosis. *Nature Reviews Neurology* 13, 1 (2017), 25–36. <https://doi.org/10.1038/nrn.2016.187>
- [36] Nikolaos A. Patsopoulos, Sergio E. Baranzini, Adam Santaniello, Parisa Shoostari, Chris Cotsapas, Garrett Wong, Ashley H. Beecham, Tojo James, Joseph Replogle, Ioannis S. Vlachos, Cristin McCabe, Tune H. Pers, Aaron Brandes, Charles White, Brendan Keenan, Maria Cimpan, Phoebe Winn, Ioannis-Pavlos Panteliadis, Allison Robbins, Till F. M. Andlauer, Onigiusz Zarzycki, Bénédicte Dubois, An Goris, Helle Bach Søndergaard, Finn Sellebjerg, Per Soelberg Sørensen, Henrik Ullum, Lise Wegner Thømer, Janna Saarela, Isabelle Cournu-Rebeix, Vincent Damotte, Bertrand Fontaine, Lena Guillot-Noel, Mark Lathrop, Sandra Vukusic, Achim Berthele, Viola Pongratz, Dorothea Buck, Christiane Gasperi, Christiane Graetz, Verena Grummel, Bernhard Hemmer, Muni Hoshi, Benjamin Knier, Thomas Korn, Christina M. Lill, Felix Luessi, Mark Mühlau, Frauke Zipp, Efthimios Dardiotis, Cristina Agliardi, Antonio Amoroso, Nadia Barizzone, Maria D. Benedetti, Luisa Bernardinelli, Paola Cavalla, Ferdinando Clarelli, Giancarlo Comi, Daniele Cusi, Federica Esposito, Laura Ferré, Daniela Galimberti, Clara Guaschino, Maurizio A. Leone, Vittorio Martinelli, Lucia Moiola, Marco Salvetti, Melissa Sorosina, Domizia Vecchio, Andrea Zauli, Silvia Santoro, Nicasio Mancini, Miriam Zuccalà, Julia Mescheriakova, Cornelia van Duijn, Steffan D. Bos, Elisabeth G. Celius, Anne Spurkland, Manuel Comabella, Xavier Montalban, Lars Alfredsson, Izaura L. Bomfim, David Gomez-Cabrero, Jan Hillert, Maja Jagodic, Magdalena Lindén, Fredrik Piehl, Ilijas Jelčić, Roland Martin, Mirela Sospedra, Amie Baker, Maria Ban, Clive Hawkins, Pirro Hysi, Seema Kalra, Fredrik Karpe, Jyoti Khadake, Genevieve Lachance, Paul Molyneux, Matthew Neville, John Thorpe, Elizabeth Bradshaw, Stacy J. Caillier, Peter Calabresi, Bruce A. C. Cree, Anne Cross, Mary Davis, Paul W. I. de Bakker, Silvia Delgado, Marieme Dembele, Keith Edwards, Kate Fitzgerald, Irene Y. Frohlich, Pierre-Antoine Gourraud, Jonathan L. Haines, Hakon Hakonarson, Dorlan Kimbrough, Noriko Isobe, Ioanna Konidari, Ellen Lathi, Michelle H. Lee, Taibo Li, David An, Andrew Zimmer, Lohith Madiredy, Clara P. Manrique, Mitja Mitrovic, Marta Olah, Ellis Patrick, Margaret A. Pericak-Vance, Laura Piccio, Cathy Schaefer, Howard Weiner, Kasper Lage, , , Alastair Compston, David Hafler, Hanne F. Harbo, Stephen L. Hauser, Graeme Stewart, Sandra D'Alfonso, Georgios Hadjigeorgiou, Bruce Taylor, Lisa F. Barcellos, David Booth, Rogier Hintzen, Ingrid Kockum, Filippo Martinelli-Boneschi, Jacob L. McCauley, Jorge R. Oksenberg, Annette Oturai, Stephen Sawcer, Adrian J. Ivanson, Tomas Olsson, and Philip L. De Jager. 2019. Multiple sclerosis genomic map implicates peripheral immune cells and microglia in susceptibility. *Science* 365, 6460 (2019), 1–10. <https://doi.org/10.1126/science.aav7188> arXiv:<https://science.sciencemag.org/content/365/6460/eaav7188.full.pdf>
- [37] Olivier B. Poirion, Xun Zhu, Travers Ching, and Lana Garmire. 2016. Single-Cell Transcriptomics Bioinformatics and Computational Challenges. *Frontiers in Genetics* 7 (Sep 2016), eCollection.
- [38] Stephen Sawcer, Robin J.M. Franklin, and Maria Ban. 2014. Multiple sclerosis genetics. , 700–709 pages. [https://doi.org/10.1016/S1474-4422\(14\)70041-9](https://doi.org/10.1016/S1474-4422(14)70041-9)
- [39] Stephen Sawcer, Garrett Hellenthal, Matti Pirinen, Chris C.A. Spencer, Nikolaos A. Patsopoulos, Loukas Moutsianas, Alexander Dillthey, Zhan Su, Colin Freeman, Sarah E. Hunt, Sarah Edkins, Emma Gray, David R. Booth, Simon C. Potter, An Goris, Gavin Band, Annette Bang Oturai, Amy Strange, Janna Saarela, Céline Bellenguez, Bertrand Fontaine, Matthew Gillman, Bernhard Hemmer, Rhian Gwilliam, Frauke Zipp, Alagurevathi Jayakumar, Roland Martin, Stephen Leslie, Stanley Hawkins, Eleni Giannoulata, Sandra D'Alfonso, Hannah Blackburn, Filippo Martinelli Boneschi, Jennifer Liddle, Hanne F. Harbo, Marc L. Perez, Anne Spurkland, Matthew J. Waller, Marcin P. Mycko, Michelle Ricketts, Manuel Comabella, Naomi Hammond, Ingrid Kockum, Owen T. McCann, Maria Ban, Pamela Whittaker, Anu Kempainen, Paul Weston, Clive Hawkins, Sara Widaa, John Zajicek, Serge Dronov, Neil Robertson, Suzannah J. Bumpstead, Lisa F. Barcellos, Rathi Ravindrarajah, Roby Abraham, Lars Alfredsson, Kristin Ardlie, Cristin Aubin, Amie Baker, Katharine Baker, Sergio E. Baranzini, Laura Bergamaschi, Roberto Bergamaschi, Allan Bernstein, Achim Berthele, Mike Boggild, Jonathan P. Bradfield, David Brassat, Simon A. Broadley, Dorothea Buck, Helmut Butzkueven, Ruggero Capra, William M. Carroll, Paola Cavalla, Elisabeth G. Celius, Sabine Cepok, Rosetta Chiavacci, Françoise Clerget-Darpoux, Kathleen Clysters, Giancarlo Comi, Mark Cossburn, Isabelle Cournu-Rebeix, Mathew B. Cox, Wendy Cozen, Bruce A.C. Cree, Anne H. Cross, Daniele Cusi, Mark J. Daly, Emma Davis, Paul I.W. De Bakker, Marc Debouverie, Marie Beatrice D'Hooghe, Katherine Dixon, Rita Dobosi, Bénédicte Dubois, David Ellinghaus, Irina Elovaaara, Federica Esposito, Claire Fontenille, Simon Foote, Andre Franke, Daniela Galimberti, Angelo Ghezzi, Joseph Glessner, Refujia Gomez, Olivier Gout, Colin Graham, Struan F.A. Grant, Franca Rosa Guerini, Hakon Hakonarson, Per Hall, Anders Hamsten, Hans Peter Hartung, Rob N. Heard, Simon Heath, Jeremy Hobart, Muna Hoshi, Carmen Infante-Duarte, Gillian Ingram, Wendy Ingram, Talat Islam, Maja Jagodic, Michael Kabisch, Allan G. Kermode, Trevor J. Kilpatrick, Cecilia Kim, Norman Klopp, Keijo Koivisto, Malin Larsson, Mark Lathrop, Jeannette S. Lechner-Scott, Maurizio A. Leone, Virpi Leppä, Ulrika Liljedahl, Izaura Lima Bomfim, Robin R. Lincoln, Jenny Link, Jianjun Liu, Aslaug R. Lorentzen, Sara Lupoli, Fabio MacCiarri, Thomas MacK, Mark Marriott, Vittorio Martinelli, Deborah Mason, Jacob L. McCauley, Frank Mentch, Inger Lise Mero, Tania Mihalova, Xavier Montalban, John Mottershead, Kjell Morten Myhr, Paola Naldi, William Ollier, Alison Page, Aarno Palotie, Jean Pelletier, Laura Piccio, Trevor Pickersgill, Fredrik Piehl, Susan Pobywajlo, Hong L. Quach, Patricia P. Ramsay, Mauri Reunanen, Richard Reynolds, John D. Rioux, Mariaemma Rodegher, Sabine Roessner, Justin P. Rubio, Ina Maria Rückert, Marco Salvetti, Erika Salvi, Adam Santaniello, Catherine A. Schaefer, Stefan Schreiber, Christian Schulze, Rodney J. Scott, Finn Sellebjerg, Krzysztof W. Selmaj, David Sexton, Ling Shen, Bridgit Simms-Acuna, Sheila Skidmore, Patrick M.A. Sleiman, Cathrine Smestad, Per Soelberg Sørensen, Helle Bach Søndergaard, Jim Stankovich, Richard C. Strange, Anna Maija Sulonen, Emilie Sundqvist, Ann Christine Syvänen, Francesca Taddeo, Bruce Taylor, Jenefer M. Blackwell, Pentti Tienari, Elvira Bramon, Ayman Tourbah, Matthew A. Brown, Ewa Tronczynska, Juan P. Casas, Niall Tubridy, Aiden Corvin, Jane Vickery, Janusz Jankowski, Pablo Villoslada, Hugh S. Markus, Kai Wang, Christopher G. Mathew, James Wason, Colin N.A. Palmer, Erich Wichmann, Robert Plomin, Ernest Willoughby, Anna Rautanen, Juliane Winkelmann, Michael Wittig, Richard C. Trembath, Jacqueline Yaouanq, Ananth C. Viswanathan, Haitao Zhang, Nicholas W. Wood, Rebecca Zuvich, Panos Deloukas, Cordelia Langford, Audrey Duncanson, Jorge R. Oksenberg, Margaret A. Pericak-Vance, Jonathan L. Haines, Tomas Olsson, Jan Hillert, Adrian J. Ivanson, Philip L. De Jager, Leena Peltonen, Graeme J. Stewart, David A. Hafler, Stephen L. Hauser, Gil McVean, Peter Donnelly, and Alastair Compston. 2011. Genetic risk and a primary role for cell-mediated immune mechanisms in multiple sclerosis. *Nature* 476, 7359 (Aug 2011), 214–219. <https://doi.org/10.1038/nature10251>
- [40] Thomas Stoeger, Martin Gerlach, Richard I. Morimoto, and Luís A. Nunes Amaral. 2018. Large-scale investigation of the reasons why potentially important genes are ignored. *PLOS Biology* 16 (09 2018), 1–25. <https://doi.org/10.1371/journal.pbio.2006643>
- [41] Tim Stuart and Rahul Satija. 2019. Integrative single-cell analysis. *Nature Reviews Genetics* 20, 5 (May 2019), 257–272.
- [42] Alan J. Thompson, Sergio E. Baranzini, Jeroen Geurts, Bernhard Hemmer, and Olga Ciccarelli. 2018. Multiple sclerosis. *Lancet* 391, 10130 (2018), 1622–1636. [https://doi.org/10.1016/S0140-6736\(18\)30481-1](https://doi.org/10.1016/S0140-6736(18)30481-1)

- [43] Carlos Torroja and Fatima Sanchez-Cabo. 2019. Digitaldsorter: Deep-Learning on scRNA-Seq to Deconvolute Gene Expression Data. *Frontiers in Genetics* 10 (2019), 978. <https://doi.org/10.3389/fgene.2019.00978>
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need (*NIPS'17*). Curran Associates Inc., Red Hook, NY, USA, 6000–6010.
- [45] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. <https://openreview.net/forum?id=rjXmpikCZ>
- [46] S.V.N. Vishwanathan, Karsten Borgwardt, Nic Schraudolph, and Risi Kondor. 2010. Graph Kernels. *Journal of Machine Learning Research* 11 (2010), 1201–1242.
- [47] Jingshu Wang, Divyansh Agarwal, Mo Huang, Gang Hu, Zilu Zhou, Chengzhong Ye, and Nancy R. Zhang. 2019. Data denoising with transfer learning in single-cell transcriptomics. *Nature Methods* 16, 9 (SEP 2019), 875+.
- [48] F. Alexander Wolf, Philipp Angerer, and Theis Fabian J. 2018. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biology* 19, 15 (2018), e-print. <https://doi.org/10.1186/s13059-017-1382-0>
- [49] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. , 40–48 pages.
- [50] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems (*KDD '18*). Association for Computing Machinery, New York, NY, USA, 974–983. <https://doi.org/10.1145/3219819.3219890>
- [51] Xiang Yue, Zhen Wang, Jingong Huang, Srinivasan Parthasarathy, Soheil Moosavinasab, Yungui Huang, Simon M Lin, Wen Zhang, Ping Zhang, and Huan Sun. 2019. Graph embedding on biomedical networks: methods, applications and evaluations. *Bioinformatics* 36, 4 (Oct 2019), 1241–1251. <https://doi.org/10.1093/bioinformatics/btz718> arXiv:<https://academic.oup.com/bioinformatics/advance-article-pdf/doi/10.1093/bioinformatics/btz718/30160037/btz718.pdf> btz718.
- [52] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J. Kim. 2019. Graph Transformer Networks. <https://papers.nips.cc/paper/9367-graph-transformer-networks.pdf>
- [53] Tao Zeng and Hao Dai. 2019. Single-Cell RNA Sequencing-Based Computational Analysis to Describe Disease Heterogeneity. *Frontiers in Genetics* 10 (2019), 629. <https://doi.org/10.3389/fgene.2019.00629>
- [54] Muhan Zhang and Yixin Chen. 2018. Link Prediction Based on Graph Neural Networks (*NIPS'18*). Curran Associates Inc., Red Hook, NY, USA, 5171–5181.
- [55] Jie Zheng and Ke Wang. 2019. Emerging deep learning methods for single-cell RNA-seq data analysis. *Quantitative Biology* 7, 4 (2019), 247–254. <https://doi.org/10.1007/s40484-019-0189-2>
- [56] Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2018. Graph Neural Networks: A Review of Methods and Applications. , arXiv:1812.08434 pages. arXiv:cs.LG/1812.08434