

## **Distribution Agreement**

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Allen Zhang

April 10, 2023

Large scale neuronal type classification using graph neural networks

by

Allen Zhang

Liang Zhao  
Adviser

Computer Science

Liang Zhao  
Adviser

Carl Yang  
Committee Member

Deqiang Qiu  
Committee Member

Giorgio Ascoli  
Committee Member

2023

Large scale neuronal type classification using graph neural networks

By

Allen Zhang

Liang Zhao

Adviser

An abstract of  
a thesis submitted to the Faculty of Emory College of Arts and Sciences  
of Emory University in partial fulfillment  
of the requirements of the degree of  
Bachelor of Arts with Honors

Computer Science

2023

## Abstract

Large scale neuronal type classification using graph neural networks

By Allen Zhang

Neuron classification is an important task in contemporary neuroscience, providing a deeper understanding of the intricate structure-function relationships of the brain. As more and more data are collected with advancements across domains, machine learning has emerged as an essential tool for automated data handling and analysis. While recent machine learning and traditional approaches to neuron classification have often relied on morphological or electrophysiological features, graph neural networks (GNNs) have not been extensively explored despite their effectiveness in analyzing complex and irregularly structured data. In this study, we show that supervised classification with GNNs on primary brain regions and cell types performs remarkably well across four large datasets. Our findings indicate that GNNs offer a distinctive and promising approach to neuron classification, with numerous potential avenues for future research.

Large scale neuronal type classification using graph neural networks

By

Allen Zhang

Liang Zhao

Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences  
of Emory University in partial fulfillment  
of the requirements of the degree of  
Bachelor of Arts with Honors

Computer Science

2023

## Acknowledgements

I would like to thank my adviser, Professor Liang Zhao, for his support throughout this project, as well as my committee members Dr. Carl Yang, Dr. Deqiang Qiu, and Dr. Giorgio Ascoli. I would also like to give extra thanks to Dr. Giorgio Ascoli for his continued assistance with NeuroMorpho.org and domain expertise and to Zheng Zhang for his mentorship and technical help, even as we worked on a rather tight schedule. Lastly, I want to give special thanks to my family for supporting me constantly in these hectic final few months.

# Contents

<b>1</b>	<b>Introduction</b> .....	<b>1</b>
<b>2</b>	<b>Methodology</b> .....	<b>4</b>
2.1	Dataset and Preprocessing.....	4
2.2	GNN Models .....	7
2.3	Experiments and Procedure.....	8
<b>3</b>	<b>Results</b> .....	<b>10</b>
3.1	Accuracy and ROC AUC .....	11
3.2	Confusion Matrices .....	15
3.3	Data Analysis .....	19
<b>4</b>	<b>Discussion</b> .....	<b>21</b>
	<b>Bibliography</b> .....	<b>27</b>

# List of Figures

2.1	Images from Example Neural Cell Reconstruction from NeuroMorpho.org .....	6
-----	---	---

# List of Tables

3.1	Dataset, Training-Validation-Test File Counts.....	10
3.2	Primary Cell Type Class Counts.....	11
3.3	Mean Training Accuracies, ROC AUC, Primary Brain Region.....	13
3.4	Mean Validation Accuracies, ROC AUC, Primary Brain Region.....	13
3.5	Mean Test Accuracies, ROC AUC, Primary Brain Region.....	13
3.6	Mean Training Accuracies, ROC AUC, Primary Cell Type .....	14
3.7	Mean Validation Accuracies, ROC AUC, Primary Cell Type .....	14
3.8	Mean Test Accuracies, ROC AUC, Primary Cell Type .....	14
3.9	Confusion Matrix, Primary Cell Type, GCN, Human, Test .....	16
3.10	Confusion Matrix, Primary Cell Type, SGCN, Human, Test.....	16
3.11	Confusion Matrix, Primary Cell Type, GCN, Drosophila, Test .....	16
3.12	Confusion Matrix, Primary Cell Type, SGCN, Drosophila, Test.....	16
3.13	Confusion Matrix, Primary Cell Type, GCN, Rat, Test .....	17
3.14	Confusion Matrix, Primary Cell Type, SGCN, Rat, Test .....	17
3.15	Confusion Matrix, Primary Cell Type, GCN, Mouse, Test.....	18
3.16	Confusion Matrix, Primary Cell Type, SGCN, Mouse, Test.....	18

# Chapter 1

## Introduction

The task of neuron classification underlies one of the fundamental questions of neuroscience: the connection between structure and function. Understanding neuron types is directly related to a more complete description of how they can affect behavior and cognition, but this is no easy task considering the sheer diversity of cell types in the brain [1]. The continued growth of big data in neuroscience, accelerated through key advances in data collection in recent years [2], has also quickly rendered manual classification efforts impractical.

Computer science and machine learning methodologies are at the forefront to tackle this problem, with many studies in the past decade leveraging different forms of machine learning to develop methods to automate classification of neurons [1] [3]. These methods have varied greatly, but they often revolve around either the comparison of morphometric features or alignment of neuronal structures [4] [5]. These studies have offered crucial information regarding the functionality of neuron morphology, but methodologies that are closely tied to user-selected morphological features carry inherent drawbacks of potential bias or limited power in representing a global structure [6]. As such, it is worth examining the potential graph neural networks (GNNs) could provide in neuronal classification given they are largely feature-agnostic, which is an area that has not seen much investigation despite robust advances in the efficiency and accuracy of such models in recent years [7] [8].

GNNs are a type of neural network designed to operate on data structured as graphs. In contrast to traditional neural networks such as multi-layer perceptrons or convolutional neural networks that typically operate on vector or grid-like data, including images and text, GNNs take data represented as a set of nodes and connecting edges. A node represents an entity of interest, and an edge represents a relationship between two nodes. Additionally, nodes may be positioned in dimensional space, enabling graphs to encode spatial information. For example, in the case of neurons, a node can correspond to a neuronal compartment of the cell located in three-dimensional space, while an edge can correspond to the physical connection between two compartments, which together form a three-dimensional graph representation of a neuron. The key idea behind GNNs is to propagate feature information between nodes in the graph by iteratively aggregating information from a node's neighboring nodes and using it to update the node's own representation [9]. This process is repeated multiple times for each node and its neighbors to capture information across the whole graph. Graphs can differ greatly in structure and size, making it challenging to apply traditional neural networks. However, the flexibility of graphs in representing data with a non-Euclidean structure is what also lends GNNs their applicability; indeed, GNNs have been successful in a wide range of settings, including node classification [10] [11], chemical molecule analysis [12], and traffic forecasting [13] [14].

In this work, we examine the performance of supervised deep learning models applied to three-dimensional neuron reconstructions as graphs. Rather than categorizing and clustering neurons based on morphological features, the proposed approach using GNNs aims to predict a neuron's anatomy, such as its cell type or brain region of origin, based solely on its graph structure and information contained in a standardized neuron morphology file. This standardization is part of the public online database of digital neuron reconstructions

NeuroMorpho.org, which consists of a collection of over 240,000 neural cell reconstructions and counting contributed by hundreds of laboratories around the world [15]. Each neuron reconstruction in NeuroMorpho.org has a standardized morphology file in the SWC file format and is associated with detailed metadata, including animal subject, anatomy, experiment details, and source [16]. To investigate the efficacy of GNNs on the supervised classification of neurons, several different models of varying methodologies were trained on over 220,000 of the available neuronal cell reconstructions on NeuroMorpho.org, split across the four main species of animal subjects in the repository. The five models trained to classify the primary brain region and primary cell type of neurons demonstrated impressive classification accuracy across the large, variable datasets, indicating that GNN models hold great promise as a powerful tool for better understanding neurons and the complex biological systems they form.

# Chapter 2

## Methodology

### 2.1 Dataset and Preprocessing

The 3D morphological reconstructions of neuronal cells used in this work were obtained from NeuroMorpho.org. Due to apparent restrictions with downloading extremely large numbers of files using NeuroMorpho.org’s web-based download functionality, the requisite neuron morphology files were instead acquired through direct requests to the URLs of the files. To obtain the list of all neuron reconstructions linked with a specific animal species, NeuroMorpho’s “Browse” feature was first used for filtering based on animal species. The resultant HTML of the filtered list was then scraped to generate a list of all the names of the neuron reconstructions. We developed a Python script to create the URLs of the neuron reconstruction morphology files from this list of names and download the file data. This process was performed for four species—human, *drosophila melanogaster* (*drosophila*), rat, and mouse—which constitute the majority of the neurons in the NeuroMorpho.org database. After excluding data that either could not be downloaded or preprocessed into a usable format, the resulting dataset comprised  $n = 220,433$  total neural cells. Neuron data was separated by the four species:  $n = 10,659$  human,  $n = 35,496$  *drosophila*,  $n = 52,411$  rat, and  $n = 121,867$  mouse neurons. The code used can be accessed at the supplementary [GitHub](#) repository.

NeuroMorpho.org uses a standardized format for representing neuron morphology called SWC [17]. As explained in the FAQ of NeuroMorpho.org:

Each line [in an SWC file] has 7 fields encoding data for a single neuronal compartment:

- an integer number as compartment identifier
- type of neuronal compartment
  - 0 - undefined
  - 1 - soma
  - 2 - axon
  - 3 - basal dendrite
  - 4 - apical dendrite
  - 5 - custom (user-defined preferences)
  - 6 - unspecified neurites
  - 7 - glia processes
- x coordinate of the compartment
- y coordinate of the compartment
- z coordinate of the compartment
- radius of the compartment
- parent compartment

Every compartment has only one parent and the parent compartment for the first point in each file is always -1. [18]

To process these SWC files into a usable representation, we used another Python script, which can also be accessed at [GitHub](#). One neuron reconstruction—that is, one SWC file—is equivalent to one graph, and each line in the SWC file is equivalent to one node in the graph: x,

y, and z coordinates and parent compartment information translate to position and edge connectivity, respectively, and neuronal compartment type and radius are included as node features.

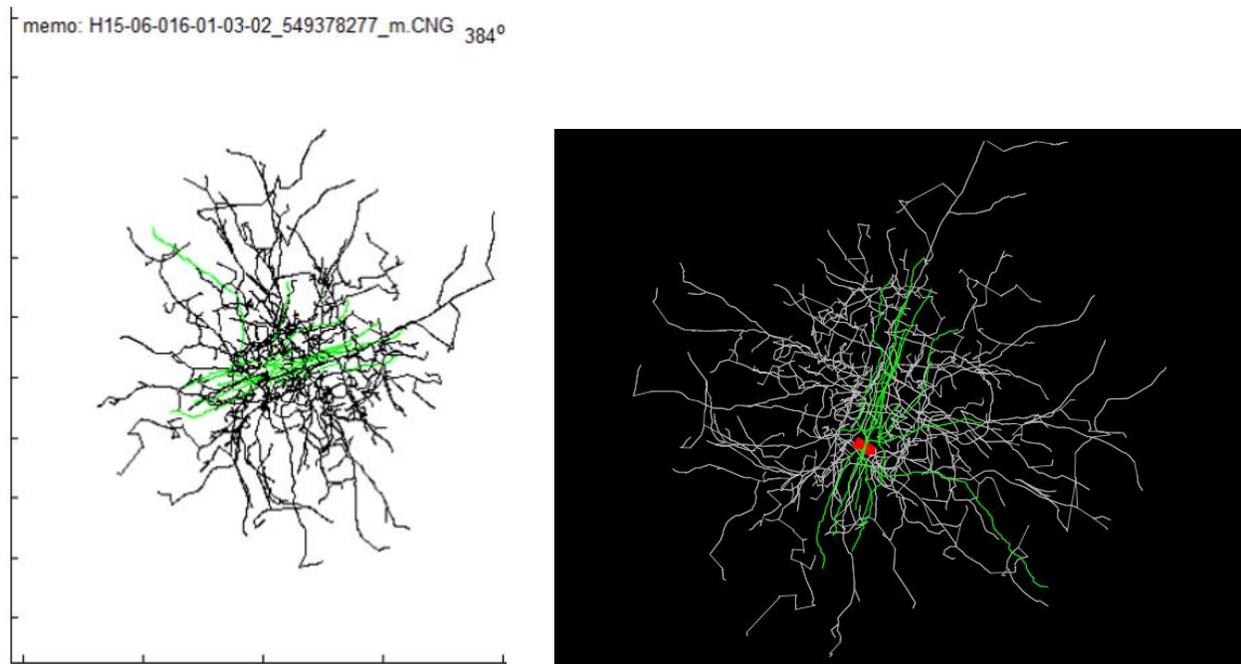


Figure 2.1: Images of example neural cell reconstruction from NeuroMorpho.org, which depict the neural cell reconstruction for the cell with NeuroMorpho.org ID NMO\_86952 [19] [20]. The left image is a screenshot of the cell as viewed in the “Animation” feature on the cell’s corresponding NeuroMorpho.org page, and the right image is the cell’s representative image in the database.

To complete the conversion into a graph representation of a neuron morphology SWC file, first, within each species, the metadata of all files are accessed through the NeuroMorpho.org API [15]. The metadata are parsed iteratively to count the class occurrences, such as “hippocampus” or “neocortex,” of the desired classification category, such as “primary brain region,” throughout each species dataset, resulting in a list of the class counts. The list of  $n$  classes in the dataset is then numbered and set as the graph-level target to train against for each

graph representation, depending on the known class of the neuron. For instance, if “hippocampus” is encoded as “5” in the class list, a neuron with a known primary brain region of “hippocampus” would be assigned the graph-level label “5.”

## 2.2 GNN Models

In this study, we assessed the performance of five different GNN models, each designed to capture distinct types of information for representation learning. Specifically, the models can be broadly categorized into three groups: GNNs that utilize graph information, GNNs that utilize spatial information, and GNNs that leverage both types of information.

Graph Convolutional Networks (GCNs), a type of GNN that uses graph structure information, utilize feature information from a given node’s one-hop neighbors to compute the node’s new representation. GCNs can be considered as learning filters, or convolutions, that transform the feature information of a node and its neighbors to produce a new, more informative node representation [10].

Graph Isomorphism Networks (GINs) also use graph information, but instead of using node feature information, GINs use a learnable graph-level aggregation function to generate node representations [21]. This aggregation function is permutation-invariant, meaning that it processes the neighbors of a node irrespective of their ordering.

In contrast, PointNet is a model that primarily operates on spatial information in the form of point cloud data, which is unstructured data consisting of a set of three-dimensional points in space where each point is represented as a vector of features. PointNet applies a series of layers that operate on each point independently, followed by max pooling over all points to generate a

global feature vector that represents the entire point cloud [7]. Similar to GINs, PointNet is also permutation-invariant.

The remaining two models belong to the category of GNNs that use both graph and spatial information. The first is a modified version of PointNet, referred to as SpatialPoint. SpatialPoint jointly considers graph and spatial information by treating the coordinates of nodes and their neighbors as messages passed along their connecting edges. It is included in Zhang and Zhao [8] as a variation of PointNet.

Finally, Spatial Graph Convolutional Networks (SGCNs) apply a convolution operation, similar to GCNs, to learn spatial and graph information by using the relative coordinates between nodes and their neighbors [22].

## 2.3 Experiments and Procedure

All experiments were performed on a 64-bit Linux machine with NVIDIA GPUs (RTX A4000, 16GB GDDR6). GNN models were implemented using PyTorch Geometric, a Python library for deep learning on graphs, point clouds, and other irregular structure data types [23] [24].

After preprocessing the neuron morphology files into a graph format compatible with the PyTorch Geometric library, the files from each species dataset were aggregated and utilized as input for the GNN models. All four species datasets were subjected to each of the five GNN models. On each run, the datasets were randomly divided into training, validation, and test sets with an 80%:10%:10% ratio, and identical hyperparameters were employed across all trials, except for the random seed that was responsible for the data split. To reduce the variance of random data splits, each GNN model was run five times on each dataset, and metrics were averaged across the five runs. Subsequently, the GNN models were used to preprocess and

analyze each dataset twice: once for “primary brain region” and once for “primary cell type.”

Details on usage can be seen on [GitHub](#).

# Chapter 3

## Results

Table 3.1 summarizes the file counts for the four datasets used, as well as the number of files used for the random training-validation-test splits. Table 3.2 summarizes the class counts for the primary cell types of the four datasets. For the most part, the primary cell type classes are reasonably well-balanced with relatively robust numbers to train on, save for the “long-range non-principal GABAergic” and “null” classes, which are by far the minority classes. Given the small sizes, it is difficult to draw conclusions about these classes, particularly the “null” class, considering this designation is related to a lack of metadata information rather than specific attributes. As a note, detailed class counts for the primary brain regions are not provided here, as they involve many more classes that differ between datasets. These class counts will be provided on [GitHub](#) instead.

	Training	Validation	Test	Total
human	8527	1065	1067	10659
drosophila	28396	3549	3551	35496
rat	41928	5241	5242	52411
mouse	97493	12186	12188	121867

Table 3.1: Dataset file counts, training-validation-test split counts

	glia	interneuron	long-range non- principal GABAergic	null	principal cell	sensory	Total
human	490	269	0	3	9385	512	10659
drosophila	160	9005	0	11	21661	4659	35496
rat	8985	5989	0	142	36869	426	52411
mouse	53905	10144	11	105	52215	5487	121867

Table 3.2: Primary cell type class counts

### 3.1 Accuracy and ROC AUC

Tables 3.3-3.8 summarize the mean accuracies and mean ROC AUC scores for each of the GNN model performances on the four datasets when classifying for primary brain region and primary cell type. Each table represents mean accuracies and ROC AUC scores for the training, validation, or test splits.

Accuracy in classification tasks is a simple metric that describes the percentage of correct predictions made and is calculated by dividing the number of correct predictions by the total number of predictions. On the other hand, ROC AUC, short for area under the curve (AUC) of a receiver operating characteristic (ROC) curve, is a measure of how well a model is able to distinguish between positive and negative classes. The ROC curve is a plot of true positive rate (TPR)—correct positive results among all positive samples—against false positive rate (FPR)—incorrect positive results among all negative samples—at different classification thresholds. Calculating the AUC of this curves yields the ROC AUC, which acts as a summary of the curve. An ROC AUC score of 0.5 means the model is no better than random chance guessing and gets just as many incorrect as correct positive and negative classifications.

ROC AUC only works on binary classification problems, but it can be extended to multiclass classification with an approach called “one versus one” (OVO). In OVO, a binary classifier is trained for every pair of classes, with one class designated as the positive class and

the other as the negative class. An ROC AUC score is calculated for every individual pair of classes, and the overall ROC AUC for the multiclass classification is calculated by averaging these ROC AUC values [25]. OVO contrasts with the “one versus rest” (OVR) method, which instead calculates ROC AUC scores using one class as the positive class and all other classes as the negative class, repeating for and averaging over all individual classes. Comparatively, OVO is more computationally expensive, but it can be more accurate in cases where there are many classes or classes are not balanced in number. Given the number of classes and varying class counts in the four datasets used, we calculated the ROC AUC scores using OVO.

GNN Model	human		drosophila		rat		mouse	
	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC
GCN	0.8590	0.8076	0.4679	0.8239	0.8012	0.7254	0.5174	0.7398
GIN	0.8898	0.8343	0.5107	0.8562	0.8637	0.7371	0.6078	0.7378
PointNet	0.9056	0.8181	0.5525	0.8663	0.8429	0.7205	0.5697	0.7476
SpatialPoint	0.9361	0.8577	0.5731	0.8799	0.9306	0.7647	0.6863	0.6966
SGCN	0.9461	0.8801	0.5978	0.8692	0.9494	0.7491	0.7182	0.7218

Table 3.3: Mean training accuracies and ROC AUC scores, primary brain region classification

GNN Model	human		drosophila		rat		mouse	
	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC
GCN	0.8753	0.7762	0.4673	0.7906	0.8078	0.6632	0.5208	0.6607
GIN	0.8935	0.7843	0.5137	0.8426	0.8653	0.6663	0.6054	0.6343
PointNet	0.8918	0.8000	0.5445	0.8390	0.8223	0.6470	0.5623	0.6578
SpatialPoint	0.9102	0.8088	0.5468	0.8432	0.8920	0.6839	0.6603	0.6387
SGCN	0.9459	0.8295	0.5670	0.8458	0.9368	0.6694	0.7084	0.6155

Table 3.4: Mean validation accuracies and ROC AUC scores, primary brain region classification

GNN Model	human		drosophila		rat		mouse	
	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC
GCN	0.8337	0.7270	0.4639	0.8132	0.8019	0.6544	0.5112	0.6264
GIN	0.8905	0.7474	0.5103	0.8406	0.8652	0.6546	0.5932	0.6128
PointNet	0.8722	0.7646	0.5360	0.8361	0.8207	0.6369	0.5534	0.6219
SpatialPoint	0.8982	0.7558	0.5456	0.8407	0.8972	0.6879	0.6518	0.6091
SGCN	0.9275	0.8129	0.5709	0.8466	0.9349	0.6489	0.7071	0.6140

Table 3.5: Mean test accuracies and ROC AUC scores, primary brain region classification

GNN Model	human		drosophila		rat		mouse	
	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC
GCN	0.9784	0.9120	0.7280	0.8881	0.9257	0.9172	0.9254	0.7801
GIN	0.9790	0.9006	0.7278	0.8773	0.9327	0.9124	0.9423	0.8114
PointNet	0.9729	0.8944	0.7901	0.8755	0.9368	0.9417	0.9378	0.8515
SpatialPoint	0.9750	0.8962	0.8014	0.8719	0.9622	0.9631	0.9378	0.8322
SGCN	0.9865	0.9651	0.7966	0.9187	0.9781	0.9766	0.9662	0.9034

Table 3.6: Mean training accuracies and ROC AUC scores, primary cell type classification

GNN Model	human		drosophila		rat		mouse	
	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC
GCN	0.9630	0.9119	0.7200	0.8882	0.9252	0.9095	0.9228	0.7789
GIN	0.9793	0.9151	0.7205	0.8827	0.9313	0.9028	0.9386	0.7963
PointNet	0.9822	0.9148	0.7798	0.7935	0.9255	0.9249	0.9329	0.8501
SpatialPoint	0.9822	0.9224	0.7712	0.8075	0.9506	0.9409	0.9324	0.8211
SGCN	0.8922	0.8850	0.7647	0.9285	0.9665	0.9608	0.9598	0.8918

Table 3.7: Mean validation accuracies and ROC AUC scores, primary cell type classification

GNN Model	human		drosophila		rat		mouse	
	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy	ROC AUC
GCN	0.9816	0.9146	0.7204	0.8318	0.9254	0.9219	0.9255	0.8494
GIN	0.9816	0.9066	0.7213	0.8449	0.9320	0.9052	0.9166	0.8388
PointNet	0.9681	0.8786	0.7752	0.8405	0.9259	0.9332	0.9284	0.8694
SpatialPoint	0.9699	0.8717	0.7647	0.8738	0.9484	0.9378	0.9280	0.8733
SGCN	0.9829	0.9438	0.7639	0.9127	0.9651	0.9619	0.9250	0.8337

Table 3.8: Mean test accuracies and ROC AUC scores, primary cell type classification

## 3.2 Confusion Matrices

Tables 3.9-3.16 provide examples of confusion matrices for runs of the GCN and SGCN models—the generally weakest and strongest performing models out of the five, respectively—on primary cell type classification for the four datasets. A confusion matrix summarizes the predicted and actual classes in a tabular format, where each row corresponds to a true class and each column corresponds to a predicted class. The diagonal of the matrix represents the number of observations that were correctly classified for each class, while the off-diagonal elements represent the number of misclassifications. Since a single confusion matrix can only represent one run of one model on one dataset, we opted to select several runs to demonstrate the observed patterns. Furthermore, in the case of the primary brain region classification task, the high number of classes made a confusion matrix impractical to show in text; however, example confusion matrices are also available on [GitHub](#).

		Predicted class					total
		glia	interneuron	null	principal cell	sensory	
True class	glia	<b>56</b>	0	0	1	0	57
	interneuron	0	<b>12</b>	0	14	0	26
	null	0	0	<b>0</b>	0	0	0
	principal cell	0	1	0	<b>929</b>	0	930
	sensory	0	0	0	5	<b>49</b>	54

Table 3.9: Sample confusion matrix, primary cell type (GCN, human, test split)

		Predicted class					total
		glia	interneuron	null	principal cell	sensory	
True class	glia	<b>57</b>	0	0	0	0	57
	interneuron	0	<b>13</b>	0	12	1	26
	null	0	0	<b>0</b>	0	0	0
	principal cell	0	2	0	<b>927</b>	1	930
	sensory	0	0	0	3	<b>51</b>	54

Table 3.10: Sample confusion matrix, primary cell type (SGCN, human, test split)

		Predicted class					total
		glia	interneuron	null	principal cell	sensory	
True class	glia	<b>21</b>	0	0	2	0	23
	interneuron	0	<b>39</b>	0	863	22	924
	null	0	0	<b>0</b>	0	0	0
	principal cell	0	18	0	<b>2116</b>	20	2154
	sensory	0	1	0	63	<b>386</b>	450

Table 3.11: Sample confusion matrix, primary cell type (GCN, drosophila, test split)

		Predicted class					total
		glia	interneuron	null	principal cell	sensory	
True class	glia	<b>21</b>	0	0	2	0	23
	interneuron	0	<b>307</b>	0	615	2	924
	null	0	0	<b>0</b>	0	0	0
	principal cell	0	172	0	<b>1973</b>	9	2154
	sensory	0	1	0	23	<b>426</b>	450

Table 3.12: Sample confusion matrix, primary cell type (SGCN, drosophila, test split)

		Predicted class					total
		glia	interneuron	null	principal cell	sensory	
True class	glia	<b>905</b>	0	0	9	0	914
	interneuron	0	<b>422</b>	0	162	0	584
	null	0	10	<b>0</b>	11	0	21
	principal cell	34	134	0	<b>3507</b>	2	3676
	sensory	2	0	0	6	<b>38</b>	46

Table 3.13: Sample confusion matrix, primary cell type (GCN, rat, test split)

		Predicted class					total
		glia	interneuron	null	principal cell	sensory	
True class	glia	<b>911</b>	0	0	3	0	914
	interneuron	0	<b>505</b>	0	79	0	584
	null	0	8	<b>1</b>	12	0	21
	principal cell	35	32	0	<b>3609</b>	1	3676
	sensory	2	1	0	3	<b>40</b>	46

Table 3.14: Sample confusion matrix, primary cell type (SGCN, rat, test split)

		Predicted class					total	
		glia	interneuron	GABAergic	null	principal cell		sensory
True class	glia	<b>5445</b>	2	0	0	14	1	5462
	interneuron	0	<b>744</b>	0	0	259	0	1003
	long-range non-principal GABAergic	0	0	<b>0</b>	0	0	0	0
	null	0	8	0	<b>0</b>	4	0	12
	principal cell	12	86	0	0	<b>5603</b>	14	5715
	sensory	2	2	0	0	19	<b>513</b>	536

Table 3.15: Sample confusion matrix, primary cell type (GCN, mouse, test split)

		Predicted class					total	
		glia	interneuron	GABAergic	null	principal cell		sensory
True class	glia	<b>5446</b>	2	0	0	13	1	5462
	interneuron	1	<b>755</b>	0	0	239	8	1003
	long-range non-principal GABAergic	0	0	<b>0</b>	0	0	0	0
	null	0	7	0	<b>1</b>	4	0	12
	principal cell	15	130	0	0	<b>5003</b>	25	5715
	sensory	1	9	0	0	27	<b>499</b>	536

Table 3.16: Sample confusion matrix, primary cell type (SGCN, mouse, test split)

### 3.3 Data Analysis

Broadly speaking, GNN performance on the neuron classification tasks ranged from moderate to exceptionally strong. Generally, the GNN models also showed a pattern of performance compared against each other: GCN usually had the weakest performance in both accuracy and ROC AUC metrics, followed by GIN, PointNet, SpatialPoint, and finally SGCN usually having the best performance. This falls in line with the idea that GNN models that utilize both spatial and graph information can capture a graph's representation more accurately than only spatial or only graph information could.

Regarding the primary brain region classification task, performance across the four datasets varied greatly. The human and rat datasets showed strong performance with 83.37% to 92.75% and 80.19% to 93.49% mean test accuracy respectively, as well as respective mean test ROC AUC scores ranging from 0.7270 to 0.8129 and 0.6544 to 0.6879. The drosophila and mouse datasets, however, had significantly worse performance, with 46.39% to 57.09% and 51.12% to 70.71% mean test accuracy respectively, and respective mean test ROC AUC scores from 0.8132 to 0.8466 and 0.6128 to 0.6264.

Comparatively, the primary cell type classification task showed better performance across the board for the four datasets. The human, rat, and mouse datasets all had remarkable mean test accuracies in the 90% range: 96.81% to 98.29% for human, 92.54% to 96.51% for rat, and 91.66% to 92.84% for mouse. This was paired with likewise extremely strong mean test ROC AUC scores: 0.8717 to 0.9438 for human, 0.9052 to 0.9619 for rat, and 0.8337 to 0.8733 for mouse. Drosophila was the outlier, with lower mean test accuracies ranging from 72.04% to 77.52% but comparable mean test ROC AUC scores from 0.8318 to 0.9127.

Observing the exemplar confusion matrices for the primary cell type classification task shows several of the same patterns seen previously, such as SGCN's generally better performance than GCN. Of note is that among the six classes seen in the primary cell type classification task, interneurons seem to result in the largest proportional number of mistakes. On the human, drosophila, rat, and mouse datasets, respectively, SGCN incorrectly identified interneurons 13 times out of 26, 617 out of 924, 79 out of 584, and 239 out of 1003, while GCN incorrectly identified interneurons 14 out of 26, a whopping 885 out of 924, 162 out of 584, and 259 out of 1003.

# Chapter 4

## Discussion

The field of neuroscience has long been interested in understanding the complex interactions between the brain's cells and networks and how they give rise to cognitive functions and behavior. Classification of neurons is one of the fundamental steps in progressing towards this goal, and with the increasing availability of high-throughput data acquisition techniques, the amount of data and information that is being created to further our understanding is always increasing. To extract meaningful insights from this growing amount of data, however, there is also a greater need for advanced computational techniques such as machine learning to aid and handle analysis where manual efforts would be far too time-consuming and labor-intensive. Indeed, machine learning has shown immense potential in neuroscience for the classification of neurons based on their morphology or electrophysiology, such as the supervised and unsupervised classification of neurons from morphological features in Bijari et al. [26] and the classification of mouse visual cortex neurons based on electrophysiological and morphological properties in Gouwens et al. [4].

A relatively underexplored avenue of machine learning in neuron classification, however, lies in graph neural networks, which offer powerful potential for analyzing irregular and complex data structures. The branching nature of neurons naturally lends itself to a graph representation, which can be effectively analyzed using GNNs. The present study aimed to explore the efficacy

of GNNs in the context of neuron classification, utilizing the publicly available database at NeuroMorpho.org to construct datasets comprising over 220,000 neurons. The results demonstrate that supervised classification with GNNs hold considerable potential for identification of neurons, with generally high classification accuracies and ROC AUC scores using OVO across the four datasets used. On the high end, the models were able to achieve 91.66% to 98.29% mean test accuracies and 0.8337 to 0.9619 mean test ROC AUC scores on the human, rat, and mouse datasets for the primary cell type classification task. On the lower end, however, the models were able to garner 46.39% to 70.71% mean test accuracies and 0.6091 to 0.8466 mean test ROC AUC scores on the drosophila and mouse datasets for the primary brain region classification task.

From the results, it seems apparent that primary brain region classification seems to be a harder task than primary cell type classification, at least through the analysis of spatial and graph information through GNNs. For each of the four datasets, performance on both accuracy and ROC AUC metrics was worse for the primary brain region classification task than the primary cell type classification task. This could possibly be explained by the increased number of classes in each dataset for primary brain region classification, which ranged from 25 to 45 classes, compared to the 5 or 6 classes for primary cell type classification. A sizable portion of these primary brain region classes do not have many samples and are thus difficult to classify from the start, yet deciding whether to prune or otherwise modify these classifications poses operational complexities.

That being said, the performance of the GNNs on the human and rat datasets for the primary brain region classification task is still worth noting. Despite potential unknown difficulties with the primary brain region classification task, the GNN models on average were

still able to consistently achieve mean test accuracies above 85% and ROC AUC scores from 0.63 in the worst case to 0.81 in the best. This provides evidence that GNNs are capable of classifying primary brain region of origin for neural cells, though more targeted investigation of any patterns in the classification by the GNN models is needed to elucidate findings.

While the performance of GNNs on primary brain region classification ranged merely from moderate to strong, performance on primary cell type classification was exceptional. As highlighted, the GNN models were consistently able to achieve greater than 90% mean test accuracies and greater than 0.85 mean test ROC AUC scores on the human, rat, and mouse datasets for primary cell type classification. Even on the drosophila dataset, which saw comparatively lower mean test accuracies from 72% to 77%, ROC AUC scores were still high, from 0.83 to 0.91, indicating the lower accuracies may be due to some specific problem classes, such as the interneuron class, as evidenced in the sample confusion matrices.

This performance pairs well with the results in Akram et al. [3], which used a novel morphometric parameter to distinguish neurons from glia with an overall high accuracy of around 97%, depending on certain parameters. If we isolate the classification results seen from the GNN models in the primary cell type classification task to just glia, the performance is on par, if not slightly better: as seen in the confusion matrices in Tables 3.9, 3.11, 3.13, and 3.15, even GCN, the consistently weakest model among the five used, was still able to correctly identify glia in the testing samples with 98.25% (56 out of 57), 91.30% (21 out of 23), 99.01% (905 out of 914), and 99.69% (5445 out of 5462) accuracy for the human, drosophila, rat, and mouse datasets respectively. Other classes, such as principal cells, which form the majority of the four datasets, likewise show remarkable performance, with interneurons generally posing the largest issue, as mentioned previously. Considering these metrics are from 10% testing samples,

it stands to reason that these impressive metrics would extend on average to the full datasets, which also collectively make up a much larger pool of data than, to our knowledge, has ever been seen before. The performance of the GNNs on primary cell type classification provide compelling evidence that GNNs are a robust way to automatically categorize cell types based on structure and spatial information and are worth using to further explore the distinctions between different neural cell types.

With regards to the lower performance observed on both classification tasks for the drosophila dataset, the reason for the discrepancy compared to the other datasets remains unclear. Judging from the confusion matrix in Table 11, drosophila interneurons may pose a particular issue for classification, but this does not necessarily explain why performance is significantly worse on primary brain region classification as well. It is possible that the rapid development of drosophila brains after hatching results in substantially different neural cell structures [27], which might render classification efforts without more specific curation difficult. However, more investigation is needed to substantiate this conjecture.

The findings of this study contribute to the existing literature on neuronal classification techniques, which have thus far primarily relied on explicit morphological or electrophysiological attributes. For example, Guerra et al. [28] distinguished between 128 pyramidal cells and 199 interneurons using 65 measured morphological features and multiple traditional machine learning models, such as hierarchical clustering, Naïve Bayes, and multilayer perceptrons; the best model was able to obtain mean accuracies up to about 91% in the best case. The aforementioned work Akram et al. [3] achieved exceptionally high accuracy in distinguishing neurons from glia using a novel morphometric parameter across around 23,000 neural cells. Li et al. [5] employed persistence diagram summaries, which, like GNNs, use graph

and spatial information of neurons, but are generated with descriptor functions that are still based on morphological or electrophysiological properties. These persistence diagram summaries attained strong classification accuracy across three datasets of 379, 114, and 1268 neuron cells, ranging from 67 to 98% accuracy in the best cases.

Compared to many of these existing studies, the results of the current study with GNNs demonstrate comparable, if not better, performance on a significantly larger pool of data, though more so with certain datasets and tasks than others, as exemplified in the wide range of mean test accuracies from 46% with the worst case worst model to 98% with the best case best model. That being said, drawing direct numbers comparisons between the results of these past studies and those of the current study is not entirely informative, as much of the existing work tends to be more narrow in scope with both tasks and datasets than the current study, which aims to serve as an exploratory investigation into the application of an underutilized tool in neuron classification.

Perhaps more importantly, on a more fundamental level, the GNN methodology is distinct in that it does not heavily rely on user-selected features, but instead utilizes almost exclusively spatial and graph information with limited node features—neuronal compartment type and radius—that are present in standardized morphology files. Not only does this make the GNN methodology remarkably accessible in terms of ease of understanding and use, but this agnostic approach may also reduce the potential for selection bias or other model selection issues while also achieving solid performance that can match known methods on certain tasks like cell type classification.

Finally, a noteworthy aspect to consider is the inherent variability present in the large dataset utilized, such as differences in anatomical details and developmental stages among species, which is a trait that is inevitable in a database as large as NeuroMorpho.org with

contributions from hundreds of labs with differing research interests. The utilization of, to the best of our knowledge, one of the largest datasets in neuron classification research thus far renders careful manual adjustment and cleaning extremely difficult without extensive expertise and time. Despite these challenges, the robust results obtained from the GNN models in this current study are highly compelling, suggesting that performance on a more tailored dataset, such as those selected in other studies for specific purposes, could be even more impressive.

Akram et al. [3] notes that a “breadth-than-depth” approach to neuron classification, by exploring large datasets first then examining potentially useful observed patterns or phenomena after, may be a particularly powerful framework for finding new areas of interest in the hunt to connect structure with function. In this work, GNNs have shown significant promise in successfully classifying neurons, and the future directions for research using GNNs in this domain are manifold, ranging from further classification of secondary and tertiary brain regions and cell types to a closer examination of the results of classification for fundamental structural differences among neuron types. Moreover, the field of GNNs itself is continuously evolving as well with new techniques and architectures that may offer even more powerful and efficient methods. The future for neuroscience is bright, with new avenues ever advancing our understanding of the complex structure-function relationships of the brain.

## Bibliography

- [1] R. Armañanzas and G. A. Ascoli, “Towards the automatic classification of neurons,” *Trends in Neurosciences*, vol. 38, no. 5, pp. 307–318, May 2015, doi: <https://doi.org/10.1016/j.tins.2015.02.004>.
- [2] E. Litvina *et al.*, “BRAIN Initiative: Cutting-Edge Tools and Resources for the Community,” *The Journal of Neuroscience*, vol. 39, no. 42, pp. 8275–8284, Oct. 2019, doi: <https://doi.org/10.1523/jneurosci.1169-19.2019>.
- [3] M. A. Akram, Q. Wei, and G. A. Ascoli, “Machine learning classification reveals robust morphometric biomarker of glial and neuronal arbors,” *Journal of Neuroscience Research*, vol. 101, no. 1, Oct. 2022, doi: <https://doi.org/10.1002/jnr.25131>.
- [4] N. W. Gouwens *et al.*, “Classification of electrophysiological and morphological neuron types in the mouse visual cortex,” *Nature Neuroscience*, vol. 22, no. 7, pp. 1182–1195, Jun. 2019, doi: <https://doi.org/10.1038/s41593-019-0417-0>.
- [5] Y. Li, D. Wang, G. A. Ascoli, P. Mitra, and Y. Wang, “Metrics for comparing neuronal tree shapes based on persistent homology,” *PLoS ONE*, vol. 12, no. 8, p. e0182184, Aug. 2017, doi: <https://doi.org/10.1371/journal.pone.0182184>.
- [6] G. Colombo *et al.*, “A tool for mapping microglial morphology, morphOMICs, reveals brain-region and sex-dependent phenotypes,” *Nature Neuroscience*, vol. 25, no. 10, pp. 1379–1393, Sep. 2022, doi: <https://doi.org/10.1038/s41593-022-01167-6>.
- [7] C. R. Qi, H. Su, M. Kaichun, and Guibas, Leonidas J, “PointNet: Deep learning on point

- sets for 3D classification and segmentation,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 77–85. doi: <https://doi.org/10.1109/CVPR.2017.16>.
- [8] Z. Zhang and L. Zhao, “Representation learning on spatial networks,” in *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, 2021, vol. 34, pp. 2303–2318. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/12e35d9186dd72fe62fd039385890b9c-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/12e35d9186dd72fe62fd039385890b9c-Paper.pdf)
- [9] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, “The Graph Neural Network Model,” *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 61–80, Jan. 2009, doi: <https://doi.org/10.1109/tnn.2008.2005605>.
- [10] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *arXiv preprint arXiv:1609.02907*, 2016.
- [11] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, 2017, vol. 30.
- [12] J. Gilmer, S. S. Schoenholz, P. F. Riley, Oriol Vinyals, and G. E. Dahl, “Neural message passing for quantum chemistry,” in *Proceedings of the 34th International Conference on Machine Learning*, 2017, vol. 70, pp. 1263–1272. Available: <https://proceedings.mlr.press/v70/gilmer17a.html>
- [13] B. Yu, H. Yin, and Z. Zhu, “Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting,” in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 2018. doi:

- <https://doi.org/10.24963/ijcai.2018/505>.
- [14] Y. Li, R. Yu, C. Shahabi, and Y. Liu, “Diffusion convolutional recurrent neural network: Data-driven traffic forecasting,” *arXiv preprint arXiv:1707.01926*, 2018.
- [15] M. A. Akram, S. Nanda, P. Maraver, R. Armañanzas, and G. A. Ascoli, “An open repository for single-cell reconstructions of the brain forest,” *Scientific Data*, vol. 5, no. 1, Feb. 2018, doi: <https://doi.org/10.1038/sdata.2018.6>.
- [16] K. Bijari, M. A. Akram, and G. A. Ascoli, “An open-source framework for neuroscience metadata management applied to digital reconstructions of neuronal morphology,” *Brain Informatics*, vol. 7, no. 1, Mar. 2020, doi: <https://doi.org/10.1186/s40708-020-00103-3>.
- [17] R. C. Cannon, D. A. Turner, G. K. Pyapali, and H. V. Wheal, “An on-line archive of reconstructed hippocampal neurons,” *Journal of Neuroscience Methods*, vol. 84, no. 1–2, pp. 49–54, Oct. 1998, doi: [https://doi.org/10.1016/s0165-0270\(98\)00091-0](https://doi.org/10.1016/s0165-0270(98)00091-0).
- [18] “Frequently Asked Questions,” *neuromorpho.org*. <https://neuromorpho.org/myfaq.jsp> (accessed Mar. 23, 2023).
- [19] C. Koch and A. Jones, “Big Science, Team Science, and Open Science for Neuroscience,” *Neuron*, vol. 92, no. 3, pp. 612–616, Nov. 2016, doi: <https://doi.org/10.1016/j.neuron.2016.10.019>.
- [20] G. A. Ascoli, D. E. Donohue, and M. Halavi, “NeuroMorpho.org: A Central Resource for Neuronal Morphologies,” *Journal of Neuroscience*, vol. 27, no. 35, pp. 9247–9251, Aug. 2007, doi: <https://doi.org/10.1523/jneurosci.2055-07.2007>.
- [21] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How powerful are graph neural networks?,” *arXiv preprint arXiv:1810.00826*, 2019.
- [22] T. Danel *et al.*, “Spatial Graph Convolutional Networks,” *Communications in Computer*

- and Information Science*, pp. 668–675, 2020, doi: [https://doi.org/10.1007/978-3-030-63823-8\\_76](https://doi.org/10.1007/978-3-030-63823-8_76).
- [23] A. Paszke *et al.*, “PyTorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, 2019, vol. 32. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf)
- [24] M. Fey and J. E. Lenssen, “Fast graph representation learning with PyTorch Geometric,” *arXiv preprint arXiv:1903.02428*, 2019.
- [25] D. J. Hand and R. J. Till, “A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems,” *Machine Learning*, vol. 45, no. 2, pp. 171–186, 2001, doi: <https://doi.org/10.1023/a:1010920819831>.
- [26] K. Bijari, G. Valera, H. López-Schier, and G. A. Ascoli, “Quantitative neuronal morphometry by supervised and unsupervised learning,” *STAR Protocols*, vol. 2, no. 4, p. 100867, Dec. 2021, doi: <https://doi.org/10.1016/j.xpro.2021.100867>.
- [27] V. Hartenstein, S. Spindler, W. Poreanu, and S. Fung, “The Development of the *Drosophila* Larval Brain,” *Advances in Experimental Medicine and Biology*, vol. 628, pp. 1–31, 2008, doi: [https://doi.org/10.1007/978-0-387-78261-4\\_1](https://doi.org/10.1007/978-0-387-78261-4_1).
- [28] L. Guerra, L. M. McGarry, V. Robles, C. Bielza, P. Larrañaga, and R. Yuste, “Comparison between supervised and unsupervised classifications of neuronal cell types: A case study,” *Developmental Neurobiology*, vol. 71, no. 1, pp. 71–82, Dec. 2010, doi: <https://doi.org/10.1002/dneu.20809>.