

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Segmentation and Detection Methods to Connect Brain Structure and Function

Permalink

<https://escholarship.org/uc/item/9s9189t4>

Author

Zhang, Angela

Publication Date

2021

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Segmentation and Detection Methods to Connect Brain Structure and Function

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Electrical and Computer Engineering

by

Angela Zhang

Committee in charge:

Professor B.S. Manjunath, Chair
Professor Kenneth Rose
Professor Shivkumar Chandrasekaran
Professor William Smith
Professor Jefferson Chen, M.D.

December 2021

The Dissertation of Angela Zhang is approved.

Professor Kenneth Rose

Professor Shivkumar Chandrasekaran

Professor William Smith

Professor Jefferson Chen, M.D.

Professor B.S. Manjunath, Committee Chair

December 2021

Segmentation and Detection Methods to Connect Brain Structure and Function

Copyright © 2021

by

Angela Zhang

Dedicated to my family - Evan, Alicia, Mom, Dad, and Smokey.

Acknowledgements

First of all, thanks to my wonderful friends and family who helped, counseled, and supported me in all aspects of life through the years. Special thanks to Evan for his wonderful cooking, big heart and tireless optimism.

I am grateful to Professor B.S. Manjunath for all of his support - intellectually, financially, and emotionally - throughout the years we have worked together. Without his guidance, this thesis would not be possible. Likewise, the financial support of National Institutes of Health [grant numbers T32-GM08620, 5R01NS103774] and National Science Foundation [grant number 1664172] made this work possible.

Many thanks to Professors Jefferson Chen and William Smith for their generosity in spending countless hours and resources to gather and annotate experimental data, as well as their invaluable guidance, encouragement, and subject matter expertise.

Thanks to Kerriane Ryan for her time and amazing annotations, and to Professors Kenneth Rose and Shivkumar Chandrasekaran for their advice.

Finally, thanks to all of the members of the Vision Research Laboratory whose time at the lab intersected with mine: Amil Khan for believing in our vision and spending nights to optimize deep learning code and implement the algorithm on the cloud; Shailja for her unyielding help and support in all of our projects together; Carlos Torres and Oytun Ulutan for their sage advice and encouragement; Michael Goebel, Raphael Ruschel, and Jiaxiang Jiang for being good sports and teammates; Satish Kumar and Devendra Jangid for saving my computer on multiple occasions; Iftekar Niloy and Austin McEver for their dry humor; Po-yu Kao, Thuyen Ngo, Utkarsh Gaur, Archith John, and Ekta Prashnani.

Curriculum Vitæ

Angela Zhang

Education

- 2021 Ph.D. in Electrical and Computer Engineering, University of California, Santa Barbara.
- 2018 M.S. in Electrical and Computer Engineering, University of California, Santa Barbara.
- 2015 B.S. in Bioengineering, Minors in EECS and Music, University of California, Berkeley.

Experience

- 2020 - Present Senior R&D Engineer, Abbott Laboratories, Westford, MA.
- 2016 - 2020 Graduate Student Researcher, Vision Research Laboratory, University of California, Santa Barbara, Santa Barbara, CA.
- 2017 Software Engineering Intern, Google Inc., Irvine, CA.
- 2016 Algorithms Intern, Toyon Research Corporation, Goleta, CA.

Publications

- A. Zhang**, A. Khan, S. Majeti, J. Pham, C. Nguyen, P. Tran, V. Iyer, A. Shelat, J. Chen, and B. Manjunath. *Automated Segmentation and Connectivity Analysis for Normal Pressure Hydrocephalus*. In review at Biomedical Engineering Frontiers.
- A. Zhang**, S. Shailja, C. Borba, Y. Miao, M. Goebel, R. Ruschel, K. Ryan, W. Smith, and B. Manjunath. *Automatic Detection and Neurotransmitter Prediction of Synapses in Electron Microscopy*. To be submitted to BMC Bioinformatics.
- S. Shailja, **A. Zhang**, and B. Manjunath. *A computational geometry approach for modeling neuronal fiber pathways*. the International Conference on Medical Image Computing and Computer Assisted Intervention, 2021.
- P. Kao, S. Shailja, J. Jiang, **A. Zhang**, A. Khan, J. Chen, and B. Manjunath. *Improving Patch-Based Convolutional Neural Networks for MRI Brain Tumor Segmentation by Leveraging Location Information*. Frontiers in Neuroscience, 2020.
- M. Kourakis, C. Borba, **A. Zhang**, E. Newman-Smith, P. Salas, B. Manjunath, and W. Smith. *Parallel visual circuitry in a basal chordate*. eLife, 2019.
- P. Kao, T. Ngo, **A. Zhang**, J. Chen and B.S. Manjunath. *Brain Tumor Segmentation and Tractographic Feature Extraction from Structural MR Images for Overall Survival Prediction*. Proceedings of the 21st International Conference on Medical Image Computing and Computer Assisted Intervention, 2018.
- P. Kao, E. Rojas, J. Chen, **A. Zhang** and B.S. Manjunath. *Unsupervised 3-D Feature Learning for Mild Traumatic Brain Injury*. Proceedings of the 19th International Conference on Medical Image Computing and Computer Assisted Intervention, 2016.

Abstract

Segmentation and Detection Methods to Connect Brain Structure and Function

by

Angela Zhang

Exploring the connection between structure and function is an important piece of the puzzle in understanding the brain. The structure of the brain refers to the physical connections and morphological features of the neuronal cells, while function refers to the behaviors and symptoms manifesting in the individual.

In this research we investigate the structure-function relationships at different spatial scales, from single neuron connections to tissue level connectomics. We present methods to detect and identify the neurotransmitter types of synapses in a small organism, *Ciona*, from confocal and electron microscopy images. *Ciona* is a chordate closely related to humans on the evolutionary tree, but with a much simpler neuronal structure. Understanding the *Ciona* connectome is an important step towards understanding the human brain. This is one of the first attempts to discover the relationship of the synaptic structures with their function in *Ciona* that includes mapping the confocal data with the EM data, detecting synaptic regions, and developing a classification model for neurotransmitter types.

At the tissue level, we consider the problem of predicting Normal Pressure Hydrocephalus (NPH) from the brain CT scans of human subjects. NPH is a highly treatable condition which is often misdiagnosed due to confounding factors, such as Parkinson's. Here we develop methods for segmenting the brain scans and using the human connectome data to improve the overall prediction accuracy. Detailed experimental results validating the methods are presented.

Contents

Curriculum Vitae	vi
Abstract	vii
1 Introduction	1
1.1 Motivation	2
1.2 Challenges and Contributions	6
1.3 Organization	8
2 Tissue Level Analysis: Predicting NPH in the Human Brain	10
2.1 Introduction	10
2.2 Background	11
2.3 Contributions	16
2.4 Data	17
2.5 Approach	18
2.6 Experiments and Results	25
2.7 Summary	30
3 Neuron Level Analysis: Ciona Connectome	39
3.1 Significance of Ciona Connectome	40
3.2 Dataset	43
3.3 Methods	50
3.4 Results	53
3.5 Summary	56
4 Synapse Level Analysis: Ciona Connetome	58
4.1 Prior Work	59
4.2 Synapse Detection and Classification	60
4.3 Results	66
4.4 Summary	77

5 Discussion	79
5.1 Broader Impacts	80
5.2 Future Directions	81
Bibliography	83

Chapter 1

Introduction

The brain is a complex web of conscious and sub-conscious thought, where processes are generated, organized, and executed. It is an organ which we are just beginning to understand. Two things we can determine about the brain are 1) its physical structure and 2) the behavior of the individual who possesses the brain in question. While we can obtain these two pieces of information through imaging techniques, little is known about the link between structure of the brain and its behavior, or function. The motivation of this dissertation is to shed more light on this connection on multiple levels of granularity.

We observe the brain through imaging techniques on the tissue level of granularity, where groups of cells with similar function can be viewed. In humans, these cells form specialized tissue in the brain, which can be imaged in a clinical setting with Computed Tomography (CT) and Magnetic Resonance Imaging (MRI).

On the cellular level, neurons connect with other neurons in an intricate web of communication. These neurons can be observed with a variety of microscopy techniques. The parts of neuronal cells which send signals to other cells are axons and dendrites, which contain synapses, which are small sub-cellular structures visible thorough electron microscopy. We study the cell-to-cell interactions in the brain of an organism called *Ciona*

intestinalis, or Ciona for short. This organism is a good proxy for the human brain, but is much simpler and smaller, allowing us to gain a more complete understanding of each cell in its brain.

1.1 Motivation

There remains much to be discovered about the structure of the brain and how it relates to behavior and pathology. Currently, brain disorders such as dementia are diagnosed and treatment determined primarily through observing the outwards symptoms and behavior of the patient. Similarly, treatment efficacy is approximated through the administration of cognitive tests and behavioral observation. Analytical tools to empower doctors and researchers to measure structural features of the brain through medical image analysis would enable more accurate disease diagnosis and further in depth analysis of treatment efficacy.

Brain imaging techniques for the human brain consist primarily of Computed Tomography (CT) scans and Magnetic Resonance Imaging (MRI), which capture brain structures on the tissue level. In addition to these modalities, it is also important to understand the brain on the cellular and sub-cellular level. Microscopy techniques allow us to observe individual neurons and neuron clusters at high resolution. Understanding the inter-cellular interactions between neurons will ultimately allow us to reach a more thorough understanding of the complex workings of the brain.

The following sections go into more detail about the different levels of granularity and their benefits and challenges.

1.1.1 Imaging Modalities at Different Scales

Tissue Level Imaging

Tissues are groups of cells that have similar structure and function together as a unit. The human brain consists of grey matter, which are the cell bodies and processing units, and white matter, which are the axons that transport signals between processing hubs. In addition to these tissue types, there is also cerebrospinal fluid in pockets throughout the brain, which transports nutrients and removes waste from brain tissue, in addition to providing a cushion of fluid around the brain. The use of medical imaging techniques such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) allows us to see these structures inside the brain.

Diffusion MRI allows us to see the direction of water movement in the brain, which enables us to map out the location and direction of white matter structures in the brain. These white matter structures are the pathways for communication between brain processing regions (gray matter), and their properties can help us understand the state of the human brain. While diffusion MRI is not as commonly used in a clinical setting, especially with respect to time-sensitive situations and on patients with high levels of trauma, there is a wealth of diffusion MRI data collected by various research efforts such as the Human Connectome Project (HCP) and the Alzheimer’s Disease Neuroimaging Initiative (ADNI), which provide insight into the appearance of white matter structures of neurotypical human brains of various age groups as well as those with Alzheimer’s-caused dementia. The data gathered using diffusion MRI can be used to create a tractograph of a patient, which is a model of the major white matter tracts running through the regions of the brain and can be computed with a myriad of methods, from deterministic fiber tracking to probabilistic computations.

Compared to MRI, CT scans are faster, cheaper and used commonly in a clinical

setting to assess patients with possible brain injury or disease. Though CT scans do not differentiate between white matter and gray matter as well as MRI, they reveal the location and shape of internal brain structures containing cerebrospinal fluid, such as the ventricles and subarachnoid space. In a clinical setting, CT scans are often collected for patients with abnormal physical and psychological symptoms and behaviors requiring diagnosis. This provides valuable data rich with potential for computational analysis. Combining the information contained in clinical CT scans and symptoms with average human brain structure information can provide valuable, multi-modal insight into how certain disorders affect the brain in multiple ways.

Cellular and Sub-Cellular Level Imaging

Cellular and sub-cellular imaging allows us to see the cell-to-cell interactions taking place between neurons. Neurons consist of sub-cellular structures, including notably synapses, which are junctions of communication between one neuron to another.

Confocal light microscopy and electron microscopy (EM) can help us see these cells and sub-cellular structures. Due to the time-intensive, costly, and invasive nature of these techniques, it is not currently possible to use them on a living human brain. However, a simpler organism closely related to humans on the evolutionary tree can provide a starting point to understanding the nature of inter-neuron communication. One such organism is the *Ciona intestinalis* larva, a chordate with approximately 177 neurons. *Ciona* is a microscopic organism closely related to humans on the evolutionary tree, more so than insects such as *Drosophila*, another commonly studied organism. The full neuronal map of the *Ciona* has already been mapped [1], but the map did not include the types of messages being sent between neurons, nor are connections drawn between the neuronal map and the behavior of the organism.

One crucial piece of information to understand the inner workings of the brain is the

type of neurotransmitters being sent between neurons. Two major types of neurotransmitters exist - excitatory and inhibitory. In simplified terms, excitatory neurotransmitters sent from one neuron to another cause the receiving neuron to activate, while inhibitory neurotransmitters cause the receiving neuron to deactivate. A neuron can receive multiple conflicting neurotransmitters from different neurons at the same time, leading to complex reactions and obscure action pathways. Fluorescent microscopy of neurons that have undergone in-situ hybridization can reveal the neurotransmitter expression of neuronal types. Specific fluorescent markers are designed to bind to the messenger ribonucleic acid (RNA) of neurons, which are then visible through light microscopy.

In addition to the neurotransmitter expressed by neurons, the complementary piece of the puzzle to understanding neuronal behavior is the connection pathways, or synapses, between neurons. Analogous to the white matter structures in the human brain, synapses between neurons allow one neuron to communicate by sending packets of information (neurotransmitters) to the receiving neuron. Knowledge of the location and size of the synapses in the brain is essential to understanding the circuitry of the brain as a whole. Synapses, comprised of vesicles (which contain packages of neurotransmitters) and post-synaptic density (the region which receives the neurotransmitters), are approximately 20-40 nanometers (nm) in length, while the wavelength of light is approximately 400-700 nm. Because synapses are smaller than the wavelength of light, electron microscopy is used to reveal these structures, as electrons are much smaller than protons (approximately 1.23 nm) and can therefore image at a higher resolution. By imaging all of the neurons and their synapses, we can build a complete connectome mapping the communication pathways of each neuron to all other neurons in the system. This mapping, along with neurotransmitter expression information and behavioral experiments, can then be used to understand the basis for behavior.

1.2 Challenges and Contributions

The main challenges involve gleaning insight from incomplete and low resolution unstructured data, developing a framework which is useful and practical for clinicians in the field, and combining multi-modal data without any available ground truth due to the exploratory nature of the study.

1.2.1 Combining multi-modal medical imaging data

Data captured in hospitals are often low resolution and incomplete. This is due to several reasons, including the need for cost efficiency and concerns for the health of the patient, who is often not in stable condition and cannot be imaged for long periods of time or in restrictive environments such as a Magnetic Resonance Imaging (MRI) device. As such, Computed Tomography (CT) scans are often used due to their lower costs, faster imaging time, and adaptability to different configurations.

In order to gain more information about the patient's status, we combine CT scans with more detailed diffusion MRI in order to get a more complete picture.

Our contributions towards overcoming these challenges are:

- The use of aggregated information from ground truth annotations to provide more context for the segmentation of low-resolution CT scans.
- Bringing in additional white matter structure information through the creation of an atlas based on data from an alternative source and modifying the connectome created from the atlas using generated segmentations specific to each patient.

1.2.2 Adapting algorithms for clinical purposes

In order for an algorithm to be useful in the medical field, it must be adapted to the needs of the clinician. We present our work with our collaborators in neurosurgery at the University of California, Irvine Medical Center to make our algorithms as useful as possible in a clinical setting.

- With our collaborators at the University of California, Irvine (UCI) Medical Center, we have developed a rigorous annotation protocol to fully annotate the important regions of the brain relevant to NPH on CT scans. A valuable dataset was created through this process that can be used for training a deep network.
- In conjunction with the annotation process, we have worked to create a robust system for clinicians to upload and segment each scan on their own, through a cloud-based computational framework. It is currently being used by our collaborators at UCI.

1.2.3 Gaining insight from multi-modal microscopy data

For cellular and sub-cellular level analysis, we have disparate data from confocal light microscopy (cellular) and electron microscopy (sub-cellular). On the cellular level, we have information about the type of signals each neuron is capable of sending via in-situ hybridization experimentation, but not where they are sending the signals to. On the sub-cellular level, we have a map of the connections between each neuron, but cannot discern what kind of signals they are sending.

While highly detailed annotations of electron microscopy data exists

In order to put together the two pieces of information, we have contributed the following.

- We have used point cloud registration techniques to match cell nuclei locations in light microscopy data to cell nuclei locations in electron microscopy data. Statistical analysis was done to determine the consistency of the registration process on different regions of the brain.
- From the conclusions we are able to draw using point cloud registration, we trained a deep network to differentiate between excitatory and inhibitory synapses. We used the trained network to extrapolate on possible neurotransmitter expressions of cells which we could not differentiate using point cloud registration.

In summary, the goal of the work presented in this dissertation is to better understand how the structure of the brain affects its function on multiple levels of granularity, from sub-cellular structures to large regions of brain tissue.

1.3 Organization

This dissertation is organized in the following manner:

- Chapter 2 describes the development of novel algorithms to segment brain CT scans and predict whether a patient has Normal Pressure Hydrocephalus. The chapter discusses the connection drawn between tissue-level segmentation, patient-specific tractography, and diagnosis prediction.
- Chapter 3 provides the background of in-situ hybridization, confocal light microscopy, and electron microscopy for imaging individual neurons and sub-cellular neuronal structures, as well as a brief description of the *Ciona* larva and its neuronal structure. A point cloud matching algorithm is presented to match in-situ data with electron microscopy data, and results and shortcomings are explained.

- Chapter 4 discusses another method to take advantage of both modalities of microscopy data and gain insights about the Ciona connectome. Deep learning for synapse detection and classification are the major algorithms used, and the results are analyzed in detail to assess their rigor and applicability.
- Chapter 5 summarizes the main concepts and contributions of this dissertation, and discusses potential future work which would benefit the field and which follows from the work presented.

Chapter 2

Tissue Level Analysis: Predicting NPH in the Human Brain

2.1 Introduction

We propose an automated method of predicting Normal Pressure Hydrocephalus (NPH) from CT scans. A deep convolutional network is developed to segment regions of interest from the CT Scans. These regions are combined with information derived from MRI to predict NPH. To our knowledge, this is the first method which automatically predicts NPH from CT scans, and also the first method to incorporate diffusion tractography information for prediction.

Due to their low cost and high versatility, CT scans have long been used to help diagnose intragray-white anomalies such as NPH. However, no well-defined and effective protocol currently exists for the analysis of CT scans for NPH. The Evans' Index, an approximation of ventricle to brain volume using one 2D image slice, has been proposed but is not robust to varying conditions and imaging angles in a clinical setting [2]. In contrast, the proposed approach is an effective way to quantify the regions of interest and

offers a computationally reproducible method for predicting NPH conditions in subjects.

We propose a novel method to predict NPH by combining regions of interest segmented from CT scans with connectome data to compute connectivity features. These features capture the impact of enlarged ventricles, a typical NPH condition, by excluding fiber tracts that pass through these regions. The segmentation and network features are then used to train a Support Vector Machine for NPH prediction. Our method outperforms the current state-of-the-art by 9 precision points and 29 recall points. Our segmentation model outperforms the current state-of-the-art in segmenting the ventricle, gray-white matter, and subarachnoid space in CT scans. Our experimental results demonstrate that fast and accurate volumetric segmentation of CT brain scans can help improve the NPH diagnosis process, and that network properties can increase NPH prediction accuracy.

2.2 Background

More than an estimated 700,000 Americans have NPH, but due to the presence of confounding co-morbidities and the lack of a rigorous diagnosis protocol, the majority of cases are under- or misdiagnosed [3] as other forms of often co-morbid dementia, leading to a delay in treatment that would significantly improve neurologic function. NPH is one of few reversible causes of dementia in the elderly, making correct diagnosis important, as shunt placement to drain the excess cerebrospinal fluid (CSF) has been demonstrated to be a safe and effective treatment [4]. Meanwhile, misdiagnosis and missed treatment can lead to decreased quality of life and cognitive deterioration. NPH presents as enlargement of the lateral ventricles in the brain, while maintaining normal CSF pressure levels. CSF normally flows through the subarachnoid space and ventricles, but in patients with NPH, there is an abnormality of CSF flow and absorption that results in CSF accumulation in

the ventricular system, causing the ventricles to swell, as illustrated in Figure 2.1. This swelling causes a multitude of neurological symptoms if left untreated.

NPH afflicts mainly elderly patients, often with co-morbid factors such as Parkinson’s or Alzheimer’s disease, and is commonly accompanied by symptoms of dementia. More specifically, cognitive dysfunction, changes in gait, and urinary incontinence are major symptoms indicating the presence of NPH [4]. Current diagnostic methods for NPH involve a mixture of clinical and imaging approaches [4]. Most commonly, a memory test is conducted, along with observation of gait and inquiry of urinary continence status. In addition, a Computed Tomography (CT) scan is often acquired to visually determine lateral ventricle size.

Sometimes, the **Evans’ Index**, a 2D manually computed ratio illustrated in Figure 2.2, is computed as a proxy for the size of the lateral ventricles in comparison to the size of the brain as a whole. The Evan’s Index is the ratio of the transverse diameter of the anterior horns of the lateral ventricles to the greatest internal diameter of the skull in a single slice of a 3D volume CT or Magnetic Resonance Imaging (MRI) Scan.

Current guidelines define ventricular enlargement an Evan’s index of greater than 0.3 [2]. However, determining the Evans’ Index is time-intensive, manual, and prone to error due to varying imaging conditions and subjectivity of measurement. It has been shown that the Evan’s ratio in fact varies greatly depending on the level (slice location) of the brain CT scan image at which the frontal horns and maximal inner skull diameters are measured [2]. More importantly, the Evans’ index measurement only takes into account a proxy for the ventricle and skull sizes, as shown in Figure 2.2 and does not account for other important volumes such as the subarachnoid space—the smaller folded channels also containing CSF located near the surface of the brain—and gray-white matter volumes. These volumes are important due to the presence of increased subarachnoid space and decreased gray-white matter in subjects with dementia, which is a major confounding




Figure 2.1: Example CT image slice of area with widest frontal horns in (a) a normal subject and (b) a subject diagnosed with NPH. (c) and (d) are masks overlaid on the original image indicating the ventricles (red), gray-white matter (green), and subarachnoid space (blue). In this example, the width of the frontal horns appear similar in the normal subject and the NPH subject, but the total volume of the ventricles differs. This is one example of a case where 3D volumetric measurements may be helpful in differentiating potential patients with NPH from those without NPH. These annotations are carefully created from the 3D data and are available to the public.




Figure 2.2: Demonstration of the measurements involved in obtaining the Evans' Index in a normal (left) and NPH (right) patient. A is the distance between the widest part of the frontal horns. B is the width of the widest part of the skull. The Evans' Index is computed as A/B .

factor in diagnosing NPH.

A 3-dimensional multi-class volumetric method of measuring the relevant regions of the brain could help to mitigate these challenges and holds promise for improving NPH differential diagnosis [5]. Manual 3D segmentation of CT scans is a time-consuming process, with one detailed volumetric segmentation of the ventricles, subarachnoid space, and gray-white matter taking up to twenty hours to complete. Moreover, delineating these regions in detail requires domain expertise, limiting the possibility of such segmentation being done on a regular basis. An automated method to segment CT scans would allow for widespread usage of volumetric analysis as an aid in NPH diagnosis, as well as the monitoring of patients over time to determine the effects of treatment.

An automated method of computing the Evan's ratio from CT is presented in [6], but this method loses the volumetric advantage of directly computing the volumes from CT scans. For segmentation of the lateral ventricles, [7] uses expert priors to aid in patch based segmentation of the lateral ventricles in MRI. [8] implements a method for automated ventricular volume measurement in MRI using the strong force algorithm

to find the best statistical features and a support vector machine to compute the final classification. The paper claims feasibility in CT, but does not discuss NPH. Another method of lateral ventricle segmentation in MRI is presented in [9] using a fuzzy representative line. The authors in [10] explore challenges in ventricle segmentation from MRI using neural networks, finding difficulties transferring segmentation models trained normal patients to work with NPH patients. [11] performs highly detailed segmentation of CT scans for the purpose of NPH analysis, using 4-6 slices per scan to capture the ventricular region of the scan.

To our knowledge, our proposed method is the first to predict NPH by first fully segmenting the subarachnoid space, lateral ventricles, and gray/white matter of the brain from a CT scan. An automated method of segmenting all of the relevant regions in the brain from a CT scan would greatly aid in the discovery and treatment of patients with NPH. Moreover, changes in brain connectivity due to the enlarged lateral ventricles may be modeled from this segmentation and its effects used to aid in NPH prediction.

2.2.1 The Connectome

The connectome is a map of the white tissue structures connecting different regions in the brain, which consist of axonal bundles that neurons use to communicate with one another. These white tissue structures can be seen using diffusion MRI, an imaging method which allows us to see the diffusion properties of water through the matter being imaged. since water travels in a directed pattern through axonal bundles, they are visible through this method. Diffusion orientation computation methods and fiber tracking [12, 13] reveal network properties of the brain that are unique to each connectome, in a process called diffusion tractography.

2.3 Contributions

We hypothesise that the enlargement of the lateral ventricles in patients with NPH cause changes to the connectome, especially in the regions closest to the enlarged ventricular space (i.e., the frontal horns). Since patient specific diffusion imaging is not gathered in a typical NPH clinical workflow, we propose an inventive method to approximate changes in the connectome using the segmented scan and a diffusion MRI atlas. The network metric derived from this approximated connectome is used as additional information during the NPH prediction process. To our knowledge, this is the first time connectome information has been combined with volumetric segmentation features to predict the presence of a disorder such as NPH.

The major contributions of chapter are as follows.

1. We have developed a robust, data-driven CT segmentation method which is derived from data from NPH and age-matched normal patients which can aid in NPH diagnosis. Our algorithm is the first of its kind to take into account the spatial distribution of the regions of interest. The robustness of our algorithm allows for the segmentation of CT scans from both normal and NPH subjects, which has historically been a challenge even with deep learning [10].
2. Additionally, we combine diffusion tractography information with the segmentation in a novel way to improve NPH prediction.
3. The creation of a novel and valuable data set of detailed and full manual segmentations of CT brain scans. This data will be made publicly available and can further benefit the medical research community. The code is available on GitHub at https://github.com/UCSB-VRL/NPH_Prediction/tree/connectome.

2.4 Data

2.4.1 Data collection

To study the morphological differences between NPH and non-NPH patients and create a pipeline for predicting NPH, CT brain scans are collected from patients with and without NPH. To incorporate additional white matter structural (diffusion tractography) information into the NPH prediction process, diffusion MRI of normal subjects between the ages of 75 and 85 are collected from the Alzheimer’s Disease Neuroimaging Initiative (ADNI).

The CT data comes from two sources: the University of California Irvine Medical Center (UCI) and the Santa Barbara Cottage Hospital. This is a retrospective study, with all images de-identified as specified by the IRB agreement between each medical center and the University of California, Santa Barbara.

There is no protocol determining the number of slices, orientation, or other imaging parameters for the data used in this study. CT scans of $n = 65$ subjects are included in the study, with 42 subjects having a diagnosis of normal and 23 subjects having a diagnosis of NPH. Axial scans are acquired as part of the treatment process, and the number of slices varied from 25 to 207. For the subjects from UCI, the average subject age is 75 ± 15 years. For the subjects from Cottage Hospital, the average age of the subjects is 72 ± 14 years.

2.4.2 Data Processing and Annotation

40 manual segmentations are performed by members of the research team under direct supervision and validation by a neurological surgeon. Each annotation took approximately 20 hours to complete, and includes the subarachnoid space, ventricles, gray-white

matter, and cerebellum. The annotation is first initialized using semi-supervised snake segmentation through the ITK-SNAP tool[14], then further refined manually. There are a total of approximately 2,000 completed manually segmented slices used in this study, taking approximately 800 hours to annotate. The rest of the scans are not used for segmentation training or validation, but are added to the dataset for diagnosis training and validation.

For annotation consistency using ITK-SNAP, the viewing window is adjusted to have a viewing intensity minimum of -980 and maximum of 80 . After window adjustment, curve-based contrast adjustment is done with 3 control points. The middle control point is set at $x = -10$ and $y = 0.020$.

The refined annotation is compared to the initial semi-supervised segmentation. The Evan’s index, measured under direct supervision of a neurological surgeon, is calculated for all subjects.

As noted earlier, diffusion tractography information is incorporated in a novel way into the NPH prediction process to improve the prediction results. In addition to the CT data, Diffusion MRI of 11 normal subjects between the ages of 75 and 85 from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) are used to create an average connectome with which to conduct connectivity studies comparing NPH and non-NPH subjects.

2.5 Approach

2.5.1 Algorithm Overview

First, the CT image is segmented into three regions of interest (ROI) - the white/gray matter, subarachnoid space, and lateral ventricles - using a modified trained 3D UNet.

Algorithm 1: NPH Prediction. CT Scan Segmentation

Input : CT Scan of the subject, MNI152 Atlas, Population Average diffusion data for diffusion tractography computation in the MNI152 space**Output:** 3D Segmentation of CT Scan into 3 regions, NPH Score

- 1 Compute the Affine Transform Matrix H that aligns the MNI152 atlas to the subject space
 - 2 Compute probability maps using ground truth annotations and transform it into the subject space using H
 - 3 The CT Scans are segmented using modified 3D-UNet and affine transformed Probability Maps. The resulting segmentation is mapped back to the MNI152 space using H^{-1} for the tractography computations below
 - 4 Compute ROA(region of avoidance) Tractographs using the ventricles as the ROA
 - 5 Compute Patient Specific connectivity matrix and the derived network properties
 - 6 Using the segmentation volumes and network properties, predict whether the subject has NPH.
-

These ROIs are important to NPH prediction due to the enlargement of the lateral ventricles in relation to the rest of the brain which occurs during NPH. The UNet uses both the original CT scan as well as a probability map created from the ground truth annotations to provide contextual data to the up-sampling layers. The segmented volumes are used to create a subject-specific connectome from average older subject diffusion MRI tractography. The segmented volumetric and diffusion tractography information are used as input features to a fully connected layer to perform feature fusion. The fused features are then used to make the final prediction of NPH vs. non-NPH. Figure 2.3 illustrates the main components of the prediction process. Figure 2.4 shows the layers in the modified 3D UNet. Algorithm 1 summarizes the various steps.

2.5.2 Segmentation Using a Modified 3D UNet

We modified a 3D UNet architecture to perform accurate segmentations of the lateral ventricles, subarachnoid space, and gray/white matter in Computed Tomography (CT) scans of the brain. The relative values in each CT scan are preserved to maintain density

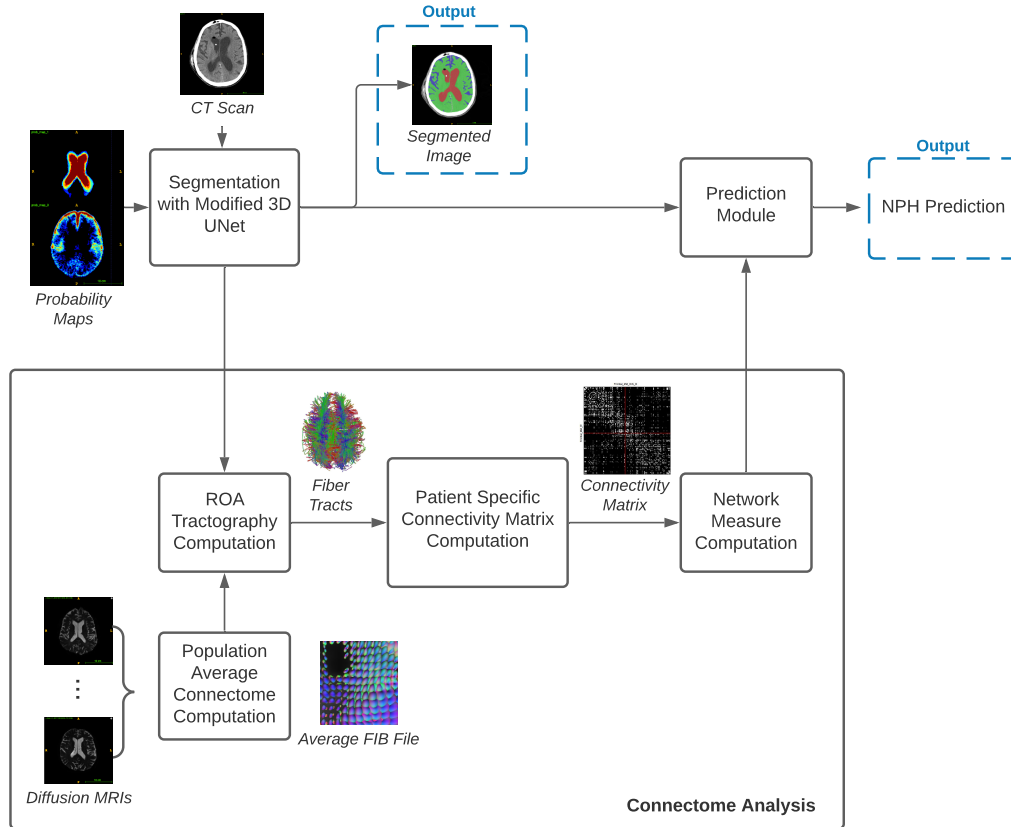


Figure 2.3: Flowchart for CT segmentation and NPH prediction. Segmentation of the CT scan is achieved with a modified 3D UNet. The segmentation result is used to create a Region of Avoidance (ROA) during the computation of the average connectome of normal elderly people. From the computed tractography results, a patient-specific connectivity map is created and the associated network properties are used in combination with the features from the modified 3D UNet (prediction module) to reach a final prediction of probable NPH.

information. In preparation for the next steps, an affine transform matrix is computed by matching the features of the CT scan to an MRI MNI-152 atlas using FSL FLIRT [15]. The original image is then passed to a modified 3D UNet, along with a probability map derived from the transformed volume which is further described in section 2.5.2.

3-dimensional convolutional neural network models are currently the state of the art for medical image segmentation. Because many of the most successful models are

variants of 3D U-Net [16, 17], we have adopted it as the basic model from which we create a segmentation pipeline. The full model is shown in Figure 2.4. The input size is $256 \times 256 \times 128$. There are 4 downsampling blocks each containing two $3 \times 3 \times 3$ convolutional layers and one rectified linear unit (RELU) [18]. These blocks each reduce the dimensions of the previous input by half while doubling the number of feature layers. The last downsampling block is preceded by max pooling. Following the downsampling layers, there are 4 upsampling layers. Each upsampling layer consists of 2 iterations of a $3 \times 3 \times 3$ convolutional layer followed by a RELU, a $2 \times 2 \times 2$ upsampling layer, a $1 \times 1 \times 1$ convolutional layer, Group Normalization [19], and another RELU. Each upsampling block increases the dimensions by 2 in all 3 dimensions, while halving the number of feature layers. The output of each upsampling block is concatenated with the output of the corresponding downsampling block. For the entire network, each $3 \times 3 \times 3$ convolutional layer is followed by group normalization. After the first RELU of the 2nd and 3rd upsampling block, the output at that layer is passed through a $1 \times 1 \times 1$ convolutional layer, followed by summation and a $2 \times 2 \times 2$ upsampling layer. The output of these extra blocks are concatenated with the next layer as shown in Figure 2.4. The output of the final upsampling block is also concatenated with the output of the probability map layers (described in the next paragraph), and passed through one more $3 \times 3 \times 3$ convolutional layer, a RELU and one $1 \times 1 \times 1$ layer. Finally, the end result is passed through a softmax layer and a $1 \times 1 \times 1$ convolution layer followed by summation. The final output is the 3D segmentation of the input CT scan, at a size of $256 \times 256 \times 128$. Our main modification to the 3D UNet is the probability map layers, which are described in the next paragraph.

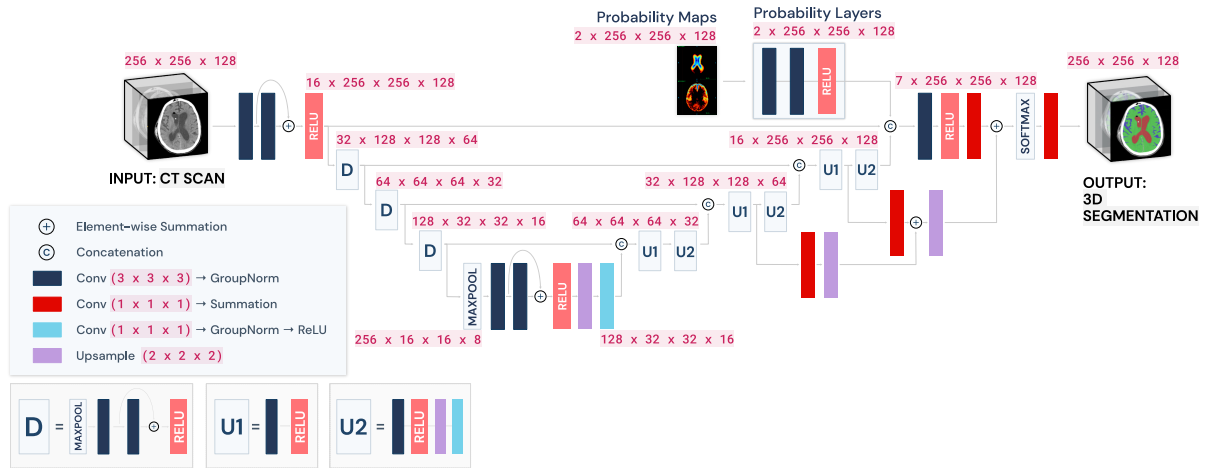


Figure 2.4: Modified 3D UNet. The 3D CT Scan is fed into a UNet, with the standard downsampling and upsampling layers. The layer with aggregated probability maps improves the performance of the network for the ventricle and especially for the subarachnoid class.

Probability Maps

While the base 3D U-Net model provided good results for brain CT segmentation, the trained network still failed to differentiate between subarachnoid space and ventricle on some border regions. In addition to this, subarachnoid space would often be misclassified as white/gray matter if the subarachnoid spaces are thin, as the intensity values are lighter for thin subarachnoid spaces. Since the occurrence of subarachnoid space and ventricles are highly similar between subjects, a location-aware layer is added to the network to boost performance in these areas. Regional probability maps are generated from the ground truth annotations and used in the up-sampling layer of the 3D UNet. The probability maps are created by averaging all of the ground truth annotations once they have been registered to a common space (MNI152). The same probability map is used for each patient, but it is transformed into that patient’s space using a reverse affine transformation that is calculated through the registration process. Two convolutional layers and a ReLU are used to learn the weights with which to apply each probability

map, and the output of the ReLU is concatenated to the output features of the 3D UNet at the last upsampling layer. The dimensions of the input and output to the probability maps layers is $2 \times 256 \times 256 \times 128$, with the 2 being the number of probability maps we are using (corresponding to the ventricle and subarachnoid space). To see the probability map layers, see Figure 2.4. These probability maps are then transformed into the patient space with an affine transformation. We use the ground truth annotations to create spatial probability maps for each of the 2 significant regions - ventricles and subarachnoid spaces. Thus, each voxel location in the probability map represents the likelihood of that voxel belonging to one of these regions.

2.5.3 Diffusion Tractography Analysis

Population Average Connectome Computation

Since the swelling of the ventricles that occur during NPH may result in damaged white matter tracts that affect brain connectivity as a whole, we analyzed the effect of using white matter disturbance information in the NPH prediction process. However, diffusion MRI are not commonly collected from patients with suspected NPH in a clinical setting. Therefore, we used a proxy for the possible white matter disturbances by modifying the average diffusion tractography from normal (control) older subjects in the ADNI dataset with the automated segmentation outputs of each subject. Deterministic fiber tractography using QSDR reconstruction [12] is conducted of the average diffusion tractography from 11 subjects in the ADNI dataset with a seed count of one million seeds per sample. This results in approximately 50,000 tracts per subject.

ROA Tractography Computation

For each subject in the NPH dataset, the predicted ventricle segmentation is used as a Region of Avoidance (ROA) during fiber tractography. Affine transformations are computed using FSL [20, 21] by first thresholding the skull and computing a rigid transformation from each CT scan to an MNI152 MRI atlas, then an affine transform using the soft tissue inside of the brain. The affine matrix for transforming each CT to MNI152 space is then used on the ventricle segmentation for each subject scan to project the ROA into MNI152 space for comparison.

Patient Specific Connectivity Matrix and Network Metrics Computation

A connectivity matrix computed using the Automated Anatomical Labeling 2 (AAL2) atlas [22] in MNI-152 space is constructed for each set of fibers, and the network properties of the matrix are used as additional features to a shallow, fully connected network which also uses the features from the last downsampling layer of the UNet to differentiate NPH and non-NPH subjects. These network properties are briefly described in Table 2.1 and are explained in detail in [23].

2.5.4 NPH Prediction Module - Feature Fusion

The volumes per class from the modified 3D UNet segmentation results are concatenated with the patient-specific network properties defined in Table 2.1. These features are then used to train a linear SVM with the L2 regularization. The output of the SVM is a binary prediction of NPH (1) or not NPH (0).

Network Coefficients	Brief Description
density	fraction of present connections to possible connections
clustering coeff average	fraction of triangles (node’s neighbors that are also neighbors of each other) around a node (vertex connected to other vertices)
transitivity	ratio of triangles to triplets (three nodes that are connected by either two or three ties) in the network
network characteristic path length	average shortest path length in the network
small-worldness	average path length of the network divided by the average path length of a random network with the same node and edge (connection between two nodes) count as the network being analyzed
global efficiency	average efficiency (1/distance) between all sets of nodes
diameter of graph	maximum eccentricity (maximal shortest path length between a node and any other node)
radius of graph	minimum eccentricity
assortativity coefficient	correlation of the degree (number of overall connections) of connected nodes.
rich k club, $k = (5, 10, 15, 20)$	fraction of edges that connect nodes of degree k or higher out of the maximum number of edges that such nodes might share

Table 2.1: Network coefficients used for NPH vs Non-NPH comparison. All coefficients included a weighted and a binary version, except for the network density.

2.6 Experiments and Results

2.6.1 Segmentation Results

The Modified 3D UNet with the probability map modification is trained on 30 segmented CT scans using a learning rate of 0.001, Adam optimizer, weight decay of 0.0001, dropout rate of 0.5, and hard per image cross entropy for 300 epochs. The segmentation masks have 6 classes—the ventricle, subarachnoid space, and gray/white matter, and the same tissue classes in the cerebellum. The cerebellum is differentiated from the rest of the brain so that its volume can be observed and monitored separately for future medical

studies. 5-fold cross-validation is used to verify the segmentation accuracy of the model.

Some basic machine learning methods are implemented to compare with the proposed method. They include random forest (RF) classification, 3D morphological geodesic active contours (MGAC) [24], and 3D morphological Chan-Vese (MCV) [25].

The implementations of these alternatives methods of ventricle segmentation use thresholding to find the skull region and remove any labels outside of this region. Each implementation first computes and applies the affine transformations into MNI152 space, then computes and applies the inverse transformation after completing segmentation. For the MCV and MGAC methods, the volumes are seeded in anatomically informed locations pertaining to the gray-white matter. The regions are then grown according to their perspective algorithms.

For the scores in Table 2.2, the Dice Score,

$$\frac{2|X \cap Y|}{|X| + |Y|} = \frac{2TP}{2TP + FP + FN} \quad (2.1)$$

where X and Y are two classes (positive and negative for each class) and TP = True Positives, FP = False Positives, and FN = False Negatives, is used for each class. The average Dice Score is taken over all of the subjects. The dice score calculates the union over intersection of a given class, and is especially useful for determining the effectiveness of a segmentation algorithm when the class labels vary in size.

The results of Table 2.2 show that our proposed method outperforms the baseline methods at segmenting the regions of interest. Our U-Net without probability map enhancement also performs well in the three categories, but the use of probability maps reduces the variance in the dice scores.

The proposed method is unique in that it allows for better separation of ventricle space and subarachnoid space. The subarachnoid space and the ventricles are both

Method	Ventricle	Gray-White Matter	Subarachnoid
3D U-Net + Prob. Maps	85.0 ± 0.1 %	93.5 ± 0.5 %	71.6 ± 5 %
3D U-Net	85.1 ± 7.2 %	93.3 ± 0.8 %	68.8 ± 13.2 %
RF + MCV	84.2 ± 3.8 %	86.8 ± 1.8 %	35.1 ± 9.5 %
Random Forest	64.6 ± 11.5 %	86.5 ± 2.0 %	N/A
3D MGAC	24.8 ± 17.4 %	81.3 ± 2.1 %	N/A
3D MCV	13.3 ± 14.1 %	79.9 ± 2.0 %	N/A

Table 2.2: Comparison of Dice Scores for various ventricle and gray-white mass segmentation algorithms for CT scans. The scores are reported as mean \pm standard deviation. All methods are our implementations created for the purpose of comparison. The marked improvement in subarachnoid segmentation performance is critical for probability maps are generated from the training set for each fold.

composed of cerebrospinal fluid, so they show up with similar intensities on a CT scan. Our method successfully separates these similar-looking regions with high performance.

The following plots further analyze the results of the segmentations for NPH and normal subjects. Figures 2.10–2.9 show the mean and standard deviation of volumes of each region by slice in the axial direction. Figure 2.11 shows the distribution of subjects by diagnosis based on their connectivity network properties.

2.6.2 Diagnosis Prediction and Comparison with Evan’s Ratio

For comparison, each scan is labeled with the Evan’s ratio as measured under direct supervision of a neurological surgeon. NPH prediction on the labeled subset using only the Evan’s ratio are first computed by the current guidelines, with subjects having an Evan’s ratio greater than or equal to 0.3 classified as NPH, and the remaining subjects classified as non-NPH.

Precision and Recall are calculated for 3 methods of NPH prediction - thresholding of the manually annotated Evans’ Index, using fully connected layers with the modified 3D UNet features, and using fully connected layers with the modified 3D UNet features along with the network properties.

Precision is defined as the number of true predicted positives over the number of all




Figure 2.5: Example of segmentation results: (a) shows the original scan, (b) shows the ground truth annotations, (c) shows the results of the segmentation network without probability maps, and (d) shows the results of the segmentation network with probability maps. It can be seen that the biggest difference exists in the subarachnoid space, which takes up a small portion of the brain, but is important for clinical analysis. The slices shown are taken at every third index in the axial direction.

	Precision (train / test)	Recall (train / test)
Evan's Index, Thresholding	86	70
Volumetric Features (model 1)	$86 \pm 3 / 86 \pm 13$	$80 \pm 7 / 76 \pm 17$
Network Properties (model 2)	$80 \pm 5 / 78 \pm 17$	$85 \pm 7 / 75 \pm 21$
Volumetric Features + Network Properties (model 3)	$96 \pm 3 / 93 \pm 12$	$89 \pm 4 / 89 \pm 13$

Table 2.3: NPH prediction scores using various methods and features. All rows except the first row (Evan's Index) used a linear support vector machine for training and testing. The predictive models are trained and tested for 100 iterations using 5-fold cross validation with randomized selection at each fold using scikit-learn, as explained in [26].

predicted positives. Recall is defined as the number of true predicted positives over the total number of actual positives. Essentially, precision is a proxy for how many selected elements are relevant, while recall is a proxy for how many relevant elements are selected.

As seen in Table 2.3, Using the volumetric information and network properties for prediction of NPH outperformed Evans' Index thresholding in both Precision and Recall. We did not implement automated Evans' Index calculations from [6] for comparison purposes, because the paper claims equivalence to manual Evans' Index calculations as

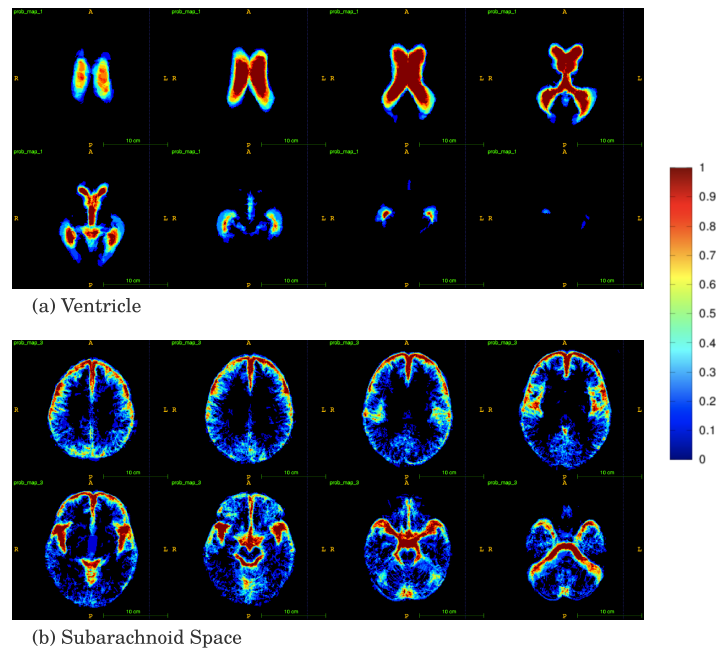


Figure 2.6: Probability maps for the ventricle (a) and subarachnoid (b) class. Each image shows one slice in the z dimension. The progression of the images (column-wise, then row-wise) goes from the top of the brain to the bottom of the brain. Cooler colors represent lower values while warmer colors represent higher values. Black represents a value of zero. The legend bar has been scaled to 1.

the best case scenario. One-sided t-tests yielded $p < 0.001$ for precision and recall between model 3 and model 2, and $p < 0.1$ for precision and recall between model 3 and model 1. Performance did not vary when we used the 3D UNet model without probability maps. This is expected, for the purpose of the probability maps is to achieve higher accuracy in the finer details of the segmentation for clinical analysis purposes.

We plot the AUC and feature importance of the SVM in Figures 2.12 and 2.13.

Feature importance are calculated by using 5-fold cross validation to generate class membership probabilities using logistic regression on the SVM's scores [27]. After obtaining the class membership probabilities, Individual Conditional Expectation plots [28] are generated using the ml-insights software. From Figure 2.13, it can be seen that the

most important features when predicting NPH with volume information only are the total volume, subarachnoid space, and cerebral mass. This is interesting because the expected features of importance are total volume and ventricle volume. However, given the top 3 features, the ventricle volume can be derived, so due to the correlated nature of the features, the feature importance plot still makes sense. For plots (b) and (c), it can be seen that the diameter and radius of the graph are important network metrics for predicting NPH. When all of the features are used for NPH prediction, a mixture of volumetric and network features are deemed important. All of the volumetric features are in the top 30% of the most important features for NPH prediction.

2.7 Summary

The paper presents a fully automated, volumetric method of lateral ventricles, subarachnoid space and gray-white mass segmentation in CT scans. Additionally, we propose a fully automated method to predict NPH diagnosis, which in conjunction with the clinical symptomatology, can facilitate the diagnosis of NPH and rule out subjects who do not meet the radiographic criteria of an NPH diagnosis. This technological system outperforms the thresholding method using Evan's ratio and can be used as a screening tool to identify or stratify possible NPH cases in a clinical setting. Our segmentation model can be used to segment and study the volumes of CT brain scans in cases other than NPH.

Furthermore, we contribute a novel method of using network properties to aid in NPH diagnosis prediction, and explores potential brain regions most affected by NPH by studying the connectivity changes that may occur during ventricle dilation.

2.7.1 Study Limitations and Future Work

The work presented in this chapter is intended as a proof of concept, with representative samples of CT scans from subjects in each category. The network is sensitive to differences in CT appearance due to age, so a fusion of networks may be used in the future to accommodate patients of all ages, not only those in the NPH risk range. Additionally, it would be beneficial to provide more examples of CT scans with high resolution in the sagittal and coronal planes, to make the model robust to scanning direction. One intrinsic limitation of this work is the idiopathic nature of NPH, meaning it is diagnosed when other conditions have been ruled out. This could lead to a biased dataset. We would like to include more confounding factors in future studies, such as the presence of Alzheimer's, Parkinson's, or non-NPH hydrocephalus.

Currently, the runtime for diffusion tractography and network analysis takes approximately 1.5 minutes per subject on a standard computer with 64 GB of RAM. It is possible to explore alternatives to the tractography process using methods such as probabilistic tractography, as described in [29] or tractography using deep learning methods, as presented in [30, 31, 32].

To conclude, auto segmentations of CT scans are helpful because CT scans are more common, readily available, and accessible compared to MRI scans. Since auto segmentations of CT scans have not been done before in such detail, this machine-learning based algorithm can be applied to study many other neurological conditions. In particular, changes over time with serial CT scans can be examined both retrospectively and prospectively. For example, CT scans of patients with Alzheimer's disease, TBI, or NPH can be analyzed to compare the ratios of ventricles to subarachnoid to cerebrum. Changes in these ratios can hold valuable information about the progression/resolution of these disease entities. Furthermore, with respect to NPH, this ability to do this type of

detailed volumetric analysis space offers tremendous potential for determining whether or not there are changes reflective of proper shunt functioning.




Figure 2.7: Example segmentation generated by the proposed algorithm. The first (a) set of images consists of cross sections of a normal subject and the second (b) set of images are cross sections of a subject diagnosed with NPH. The images in first row of each set are cross sections of the original scan, the images in the second row are the ground truth segmentation for the scan, the images in the third row are the results of the 3D UNet without the probability map modification, and the images in the fourth row are the results of the proposed algorithm with the probability map layers. The cross sections are in descending order from top to bottom of the brain.

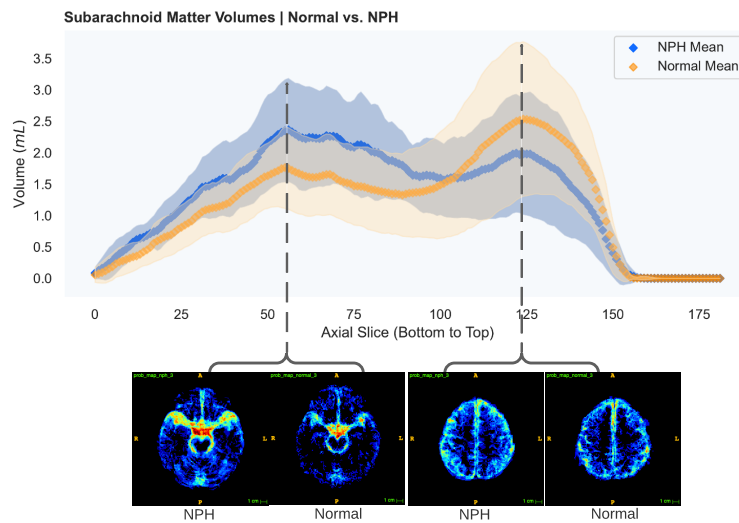


Figure 2.8: Average volumes in mL/slice of the subarachnoid in NPH (blue) and normal (orange) subjects after registration to MNI-152 space. The standard deviation is also plotted in light blue(NPH) and light orange (normal). Probability maps at inflection points are shown underneath the plot. It can be seen that the average subarachnoid space volume of subjects with NPH is greater towards the bottom of the head near the spine, and smaller towards the top of the head, compared with normal subjects.

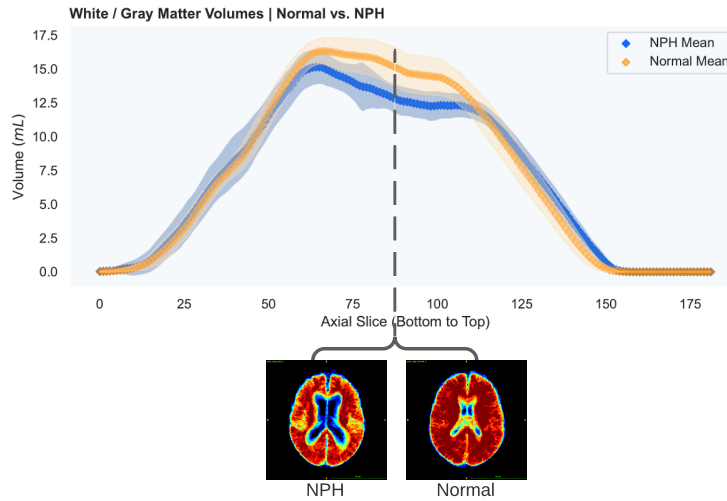


Figure 2.9: Average volumes in mL per slice of the white/gray matter in NPH (blue) and normal (orange) subjects after registration to MNI-152 space. The standard deviation is also plotted in light blue(NPH) and light orange (normal). Probability maps at the location with the largest difference between the means is shown underneath the plot.

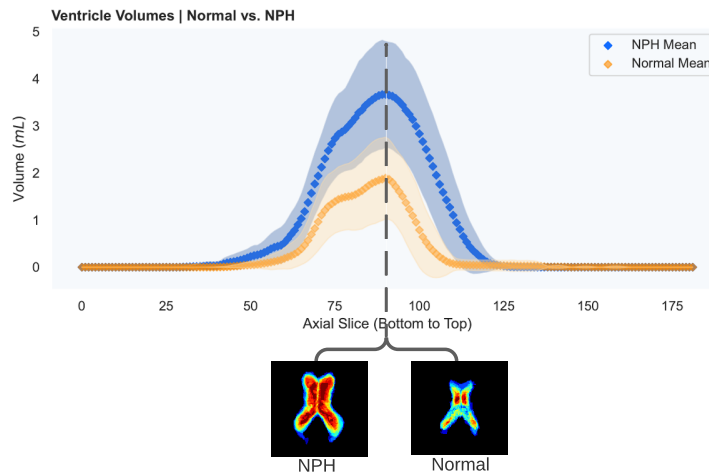


Figure 2.10: Average volumes in mL per slice of the ventricles in NPH (blue) and normal (orange) subjects after registration to MNI-152 space. The standard deviation is also plotted in light blue(NPH) and light orange (normal). Probability maps at the location with the largest difference between the means is shown underneath the plot. As expected, the volumes are greater for subjects with NPH than those without NPH.




Figure 2.11: Average location of network metric after PCA reduction from 26 connectivity features to 2 principal components. For the most part, a separation can be seen between normal and NPH subjects, but there is some mixing of the classes as well. The trend towards separation shows why the network properties are useful features for NPH prediction.

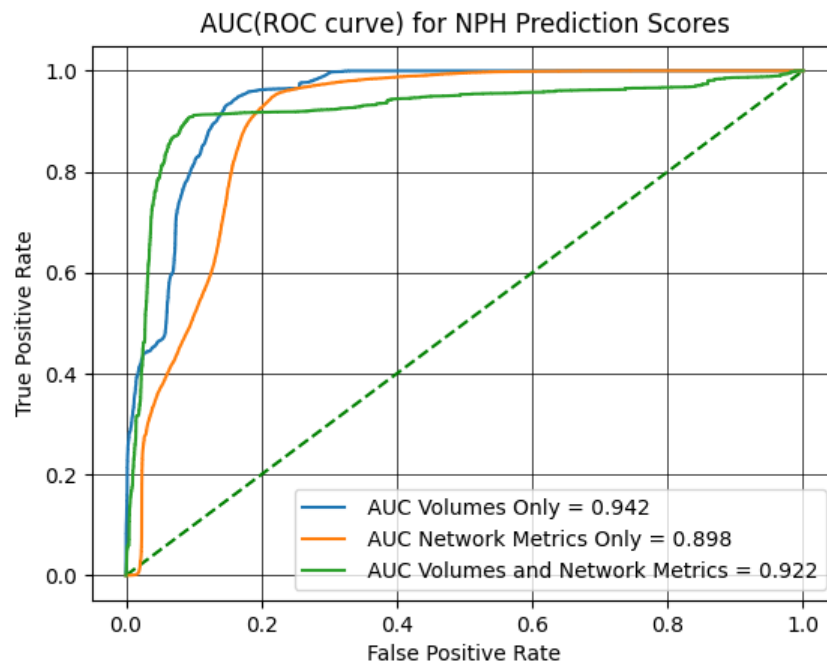


Figure 2.12: Test AUC for the linear SVM trained on different features: ventricle, subarachnoid, white/gray matter, and overall volumes (4 features); network properties (26 features); and volume and network properties (30 features). While the AUC for volumes and network metrics is lower than for volumes only, the inflection point of the model associated with volumes and network metrics has consistently higher performance over 100 iterations using 5-fold validation. Each iteration uses a new generator for the 5 folds.

Chapter 3

Neuron Level Analysis: Ciona Connectome

This chapter introduces the *Ciona intestinalis* larva as a suitable proxy of the more complex human brain with which to study individual neuron-to-neuron interactions, and analysis methods for connecting fluorescent microscopy data with electron microscopy data. Specifically, we consider establishing a correspondence between fluorescent imaging data with annotations that are derived manually from EM data, towards identifying neurotransmitter expressions of individual neurons. In Chapter 4, we extend this to include analysis on the EM imaging data, for detecting synapses and classifying these synapses based on excitatory or inhibitory neurotransmitter expressions. Much of this Chapter contents is derived from our paper [33]. The in-situ hybridization experiments and data collection are carried out by our collaborators in the Smith Lab at the University of California, Santa Barbara.

First, the motivation for analyzing the brain on a neuronal scale using Ciona is explained, followed by a section on the background of research on the central nervous system of Ciona larva. The datasets are introduced next, with two main parts - confo-

cal microscopy on in-situ hybridization results, and electron microscopy for connectome analysis. The annotation and data extraction process are detailed, and existing work on point cloud registration is explained.

After the introduction of the dataset and background, we present the use of a point-cloud registration technique to match nuclei centroid locations between the EM dataset and the in-situ hybridization results, in order to identify which cells exhibit which neurotransmitter expressions. The results and analysis of the algorithm are given and summarized.

3.1 Significance of Ciona Connectome

While tissue level analysis is commonplace in the field of brain imaging, the human brain is far too large and complex to study neuron-to-neuron interactions at a cellular and sub-cellular level with current available imaging and computational technology. A common central nervous system (CNS) architecture is observed in all chordates, from vertebrates such as humans to basal chordates like the ascidian *Ciona*. Ascidians, including those of the widely studied ascidian genus *Ciona*, are members of the chordate sub-phylum Tunicata, and comprise the closest extant relatives of the vertebrates [34]. While much smaller and containing many fewer cells than those of vertebrates, the ascidian tadpole larvae resembles vertebrate larvae in having a prominent notochord running the length of the muscular tail and a dorsal CNS with regional anterior-to-posterior homology to vertebrate CNSs [35]. Despite these homologies with vertebrates, the *Ciona* CNS contains only 177 neurons [1]. This makes it a prime candidate for connectomic analysis, as each additional neuron introduces exponential complexity to the system.

An essential aspect of constructing a connectome, a comprehensive map of neurons and their connections in the brain, is to identify the synapses which form chemical and

electrical communication links between neurons [17]. Traditionally, synapse detection of synapses is completed manually through expert analysis of thousands of EM images. This is a time consuming process that can take thousands of hours for one specimen. For example, the first Ciona connectome was constructed manually over a period of 5 years from serial section EM images [36].

Another crucial but missing component of understanding the connectome is the classification of the function and properties of each synapse. With some exceptions, Ciona neurons express only one neurotransmitter per cell, but different cells in the same may express different neurotransmitters from each other. While some experimental methods such as in situ hybridization can identify cells, or clusters of cells, expressing transcripts indicating neurotransmitter use [33], the resolving power of light microscopy cannot reveal the connectivity of these neurons. Additionally, in cases with intermingled neurotransmitter expressions and high variability in cell location across specimens, it was not possible to determine the neurotransmitter expression of the cell through in situ hybridization alone.

3.1.1 Background

Anatomically, the Ciona larval central nervous system is comprised anteriorly of the brain vesicle (BV; also known as the sensory vesicle), a region homologous to the vertebrate forebrain and midbrain, followed by the neck region, a homolog of the vertebrate midbrain/hindbrain junction as seen in Figure 3.1. Immediately posterior to the neck is the motor ganglion (MG; also known as the visceral ganglion). The MG is thought to be homologous to the vertebrate hindbrain and/or spinal cord, and contains ten motor neurons (MN) as well as a number of interneurons – including the two descending decussating neurons (ddN). An illustration of the neuronal composition of the Ciona larva is




Figure 3.1: From [33]: the minimal visuomotor circuit is shown with circles representing classes of neurons with the number of cells of each class indicated in the parentheses of the key. Abbreviations: dor., dorsal; vent., ventral; ant., anterior; post., posterior; PR-II, photoreceptor group II; PR-I, photoreceptor group I; pr-AMG RN, photoreceptor ascending motor ganglion relay neuron; prRN, photoreceptor relay neuron; MGIN, motor ganglion interneuron; MN, motor neuron. L, left; R, right. Cell types are color coded according to [1].

shown in Figure 3.1.

Ascidians, including members of the widely-studied *Ciona* genus, have a biphasic life cycle. At the start of their life, most ascidians spend their first few days as free-swimming tadpole larvae. It is at the larval stage that ascidians display unmistakable chordate traits, including a prominent notochord running the length of a muscular tail and a dorsal central nervous system (CNS). *Ciona* stands apart among chordates in having a complete larval connectome [1]. A comprehensive understanding of nervous system function requires detailed descriptions of circuits in which all contributing neurons and their synaptic connections are known and characterized – in other words, a connectome. Both the small scale of synapses ($< 0.5\mu\text{m}$), which necessitates serial section electron microscopy (ssEM) or other similar methods, and the inherent difficulty in visualizing and tracing synaptic connections has so far limited the number connectomic descriptions of complex nervous systems to discrete CNS regions [37, 38]. Until the recent description of the larval *Ciona* connectome [1], whole-nervous system connectomics was limited to the nematode *C. elegans*, although extensive connectomes are now available for *Platynereis* [39] and *Drosophila* [40, 41] larvae. Even for the simplest vertebrate models, such as

larval zebrafish with 100,000 neurons [42], a full connectome is likely many years away. Given the challenges of vertebrate connectomics, the ascidian larva provides a unique model for bridging connectomic-level understanding of simple invertebrate models and vertebrates [1]. Understanding not only the connections between these neurons, but also the types of signals being sent between one neuron to another will allow us to fully understand the neural pathways of the *Ciona* larva. We can then use the information we have gained about neuronal connectivity and activity to better understand the mechanics behind *Ciona* behavior to controlled environmental stimuli, as observed in experiments such as those in [33].

3.2 Dataset

Ciona intestinalis-type B (Marine Biological Laboratories, Woods Hole) are used for in situ studies to match the animals used in the connectome study [1]. *Ciona intestinalis* presents with certain neurotransmitters, which require transporters that can be detected with in situ staining. These neurotransmitters include the inhibitory neurotransmitters gamma-Aminobutyric acid (GABA) and glycine, and the excitatory neurotransmitters acetylcholine and dopamine.

3.2.1 Cell Types

From [1], 21 cell types are identified through electron microscopy imaging by their location and function. For our work, we are primarily interested in the photoreceptors (PR-I and PR-II) and the relay neurons (prRN, pr-AMG RN, AntRN, PBRN, PCRN, and PNRN). The photoreceptors react to changes in light levels in the environment, while the relay neurons convey messages from the photoreceptors to the rest of the nervous system.

3.2.2 In-Situ Hybridization of Ciona Neurons

The neurotransmitter expression of a cell can be detected through detection of its messenger Ribonucleic Acid (mRNA). Fluorescent markers attached to opposing strands of the signature mRNA indicating each neurotransmitter expression can be used to detect the mRNA present in a cell using light microscopy. This process is called in situ staining.

In situ staining to 5 larvae using probes for the vesicular GABA transporter (VGAT), vesicular glutamate transporter (VGLUT), tyrosine hydroxylase, and vesicular acetylcholine transporter (VACHT) was done as described in [43]. Larvae are also stained with 4',6-diamidino-2-phenylindole (DAPI), a fluorescent stain that binds strongly to adenine–thymine-rich regions in DNA in order to view the nuclei. The resulting larvae are imaged by confocal microscopy. Nuclei are segmented using the IMARIS software package, Bitplane, and manually pruned to find nuclei locations within regions displaying fluorescence from in-situ hybridization. An illustration of the results of this process is shown in Figure 3.2, and an example of the nuclei locations for the photoreceptors is shown in Figure 3.3.




Figure 3.2: From [33]: example result of in-situ hybridization and fluorescent microscopy result on the photoreceptors in Ciona [33]. The cell nuclei are visible due to nuclei staining, but the synapses and other subcellular structures which are smaller than the wavelength of light are not visible. This means we do not know the cell identities of each neuron based solely on light microscopy.

From Figure 3.3, the heterogeneity of photo receptor configurations can be seen.

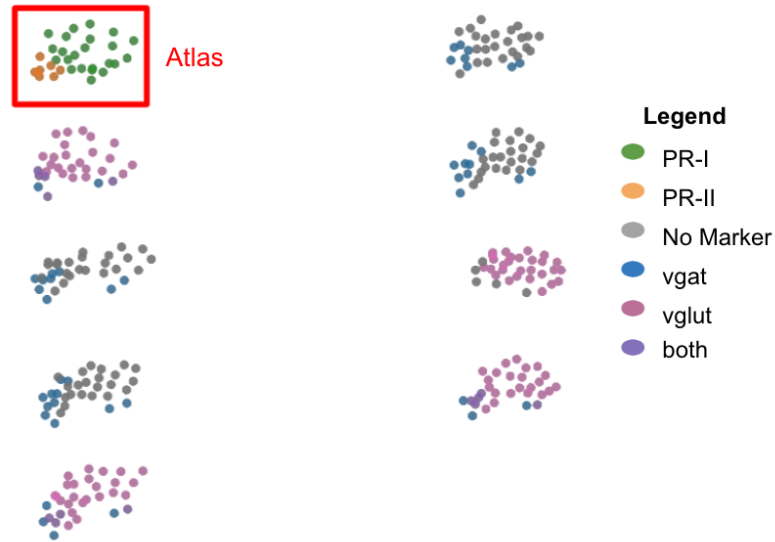


Figure 3.3: Plots of the nuclei centroids of 8 *Ciona* specimens after manually notating the centroids. The color corresponds to the in-situ hybridization appearance under confocal light microscopy. Some samples, with both pink and blue, have been hybridized with both VGAT and VGLUT, while some have only been hybridized with one of the two markers. The gray centroids have no visible marker for that particular specimen.

Computational methods can help us to match each neuron in the in-situ hybridization results to its identity, and therefore, its connectomic qualities, derived from electron microscopy imaging data.

3.2.3 Electron Microscopy of *Ciona* Neuronal Structures

For point set registration, a summary of the EM annotation results are used to extract nuclei centroid locations for each cell. An example of the aggregated information for one cell is shown in Figure 3.4.

A continuous series of cross sections are cut from a *Ciona* larva after fixation for electron microscopy (EM). A total of 3375 60 nm sections from the anterior brain vesicle and motor ganglion are collected and imaged at 3.85 nm per pixel [1]. The data collected surpasses 1 terabyte. The original *Ciona* EM serial section image data set was collected,


Cell ID	Cell Type	Ciliated	Soma location				Morphology	Cell body volume	Number of pre-synaptic sites	Number of post-synaptic sites	Reconstruction Scale bars: 10 μ m (thin); 1 μ m (thick)	
			Side	Brain region	Z	X						Y
1	cor-ass BVIN	ciliated to bp	L	anterior BV	0.83	9.5	30.6	Single simple axon with slight expansion at terminal in PBV	94.7	2	21	

Figure 3.4: Example row of aggregate and derived data from EM annotations. The relevant fields for point cloud registration are the X, Y, and Z locations. While this example shows soma location, for our work we used nuclei location.




Figure 3.5: Example of manual annotations of synapses, indicated by the cyan arrows. The zoomed in image shows a patch containing a synapse, with the vesicles encircled with bright green and the cell boundaries marked with dotted lines. The direction of the arrow does not indicate synaptic direction.

processed and annotated using the program RECONSTRUCT [44]. The annotations we focused on are the synapse annotations, which are stored as points in 3 dimensions with a naming system to indicate the pre- and post-synaptic neuron. In order to facilitate analysis, python scripts are written to interface with the program and stored data and extract image patches corresponding to annotated regions. A series of geometric transforms was used to co-register the annotation coordinates with the aligned 3D image stack coordinates. The RECONSTRUCT images used in this project are available upon request. The image processing and extraction steps are described in section 3.2.3.

Data Curation and Annotation

EM cell IDs for each neuron was stored in an excel spreadsheet, along with each cell's location in 3D coordinates. This data was extracted and normalized for matching with the collected fluorescent microscopy data.

To extract synapse images from EM data, each synapse object which was labelled using RECONSTRUCT was parsed based on pre- and post- synaptic cell ID. Then, the location of the object was found through a series of inverse transformations and magnifications as follows.

The EM images are scaled and aligned in the z-dimension while annotating. Image scaling is specified by a pixel size (magnification) parameter while alignment is represented by a non-linear transformation associated with the image. Each transformation maps trace points or image pixels into the section using a combination of basis functions representing an elementary motion such as translation, orientation, scaling, and deformation. We extract the underlying annotations (actual section coordinates in microns) from each section by combining these movement components in different proportions. Each image is associated with an independent transformation which determines the size and location of the element on the section. Applying the inverse transform on the contour point (x,y) , we obtain the points (x', y') , on which applying the forward image transformation brings the points to the original image domain. For our study, synapse is represented by seven points forming an arrow as shown in Figure 3. After getting the coordinates of these points on the original image domain, we determine the centroid of these points, and extract a 500x500 dimensional patch around this centroid point. We perform this operation for all the synapse annotations to obtain approximately 25,000 total number of patches.

Visual inspection of the synapses in Ciona do not reveal any distinguishable differ-




Figure 3.6: Illustration of Data Preparation Method. The first set of images represent the image patches which are obtained through electron microscopy. After image stitching, the individual image patches are registered and stitched together on one plane. After 3D registration, the image planes are registered lengthwise from the head to the tail of the organism. Finally, manual annotation of the registered image planes result in the labelled EM data that we use to extract synapses with which to train a deep network. The purple colored object indicates the boundaries of a cell which has been selected in RECONSTRUCT.

ences between known excitatory and inhibitory neurons. Some examples are included in Figure 3.7. Ciona synapses typically include the vesicle cluster on the transmitting end, which varies in count and size, and the post-synaptic density on the receiving end, which varies in size and density. No apparent pattern in these features is visible upon manual inspection.

From the datasets described in this chapter, it is evident that no one modality contains all of the information needed to untangle the workings of the Ciona connectome. Cross-modality analysis between data at the cellular level (i.e. confocal light microscopy) and data at the sub-cellular level (i.e. electron microscopy) must be conducted in order to gain further understanding of each neuron and its place in the brain.

3.2.4 Prior Work

To our knowledge, this is the first time a semi-automated method of matching fluorescent microscopy data with electron microscopy data has been developed.

For 3D point cloud matching, there are several techniques which are currently used to match sets of points in 3 dimensions. One such method is the iterative point cloud




Figure 3.7: Examples of (a) inhibitory and (b) excitatory synapses. The synaptic region is circled in green. The vesicles are the small circles in the center of the green bordered regions. From the image examples, it can be seen that there are varying vesicle counts and sizes between inhibitory and excitatory synapses, as well as little visible post-synaptic density, which is a darkened region on the other side of the cell boundary from the vesicles.

matching algorithm by the Insight Toolkit [14]. This method relies on the Iterative Closest Point algorithm [45], which is an iterative method to solve the least squares problem of finding the appropriate rigid or affine transformation matrix to minimize the difference between two point clouds. We initially tried using this method, but found that the algorithm often ends up in poor local minima which did not find the optimal solution. The sensitivity of the ITK Iterative Closest Point method made it less suitable for our purposes.

Another 3D point cloud matching algorithm is Robust Point Matching, which uses a soft correspondence to find the optimal affine transformation between two point clouds [46]. While this method was more robust than Iterative Closest Point, it still got easily stuck in local minima, making it a poor fit for our dataset.

3.3 Methods

3.3.1 Point Cloud Registration

For both the EM and confocal microscopy datasets, nuclei centroid locations are manually annotated for each neuron. The nuclei centroids are split by cell types, and each type is registered on its own. In this chapter, we focus on the ocellus, or the simple eye, and the relay neurons, which relay information from the ocellus to the motor system. Point cloud registration is used on the 3D coordinates of each in-situ dataset to match each cell in the dataset to a cell in the EM dataset. First, a rotation matrix is calculated based on the 3-dimensional vectors between the ddN and/or anterior cells and the center of the cluster of cells in the target set of cell locations and the source set of cell locations. The source set is then rotated to an approximate orientation to the target set. Next, the Coherent Point Drift Algorithm is used to calculate an affine transformation matrix between the source set and the target set of cells [47], and confusion matrices are constructed with all of the data available in order to analyze the consistency of the registration for each cell type. Figure 3.8 shows a summary of the steps involved in registration and consistency verification.

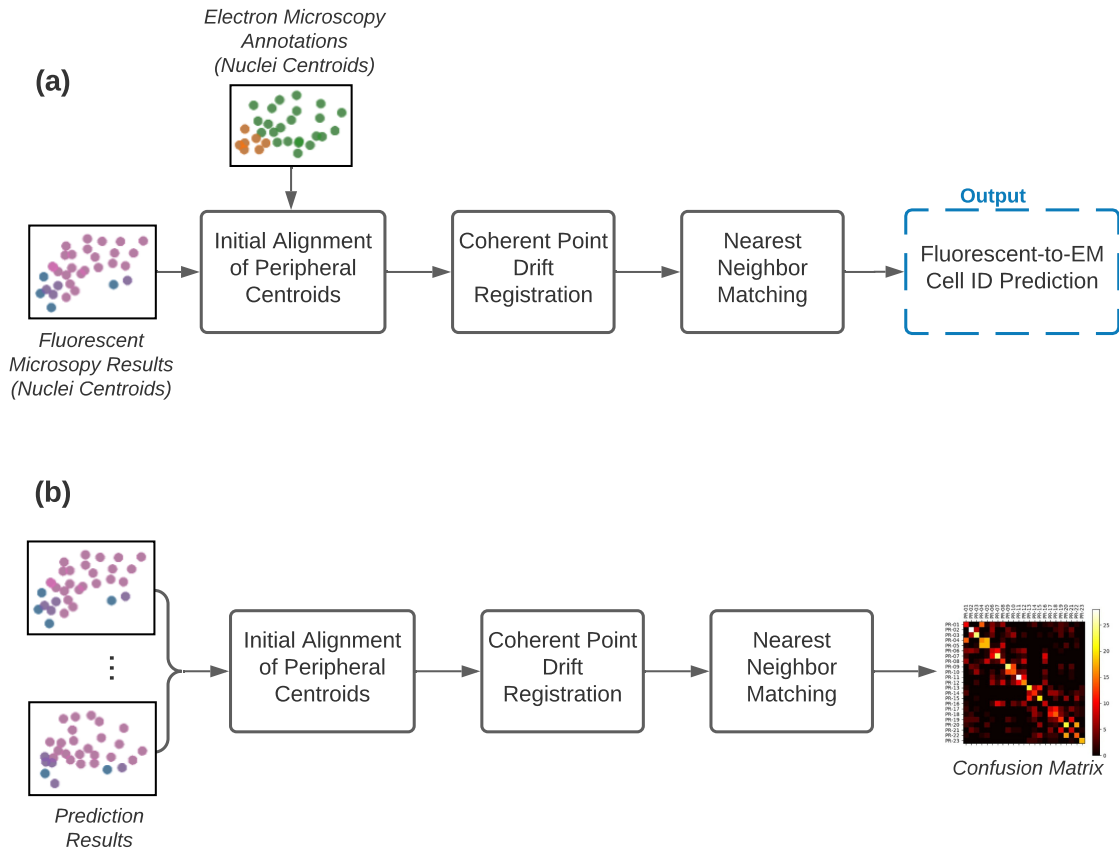


Figure 3.8: (a) The registration process for predicting cell IDs of fluorescent imaging results for each type of neurons, and (b) The process of generating a confusion matrix for the registration results of multiple fluorescent light microscopy datasets to one EM dataset.

3.3.2 Coherent Point Drift Algorithm

The Coherent Point Drift algorithm models the source set as a Gaussian Mixture Model (GMM) [48], and the target set is treated as observations from the GMM. A GMM is constructed from the moving point cloud (confocal microscopy data) and the fixed point cloud is treated as observations from the moving GMM. Due to the probabilistic nature of the algorithm, this method is more robust than the other iterative

Algorithm 2: Coherent Point Drift and Nearest Neighbors Matching

- Input** : 3D locations of nuclei centroids from fluorescent confocal microscopy
and 3D locations of nuclei centroids from electron microscopy
- Output:** 1-1 matching of confocal microscopy cell to electron microscopy cell
- 1 Initialization of affine transformation matrix to the Identity matrix.
 - 2 Repeat steps 3 and 4 until the error is under a threshold:
 - 3 **Expectation** step: find the Gaussian from which each point in the moving set was sampled from
 - 4 **Maximization** step: maximize likelihood that observed points are sampled from the GMM (solve for the affine transformation matrix)
 - 5 Apply calculated transform to moving point set
 - 6 For each point in the moving set, find the nearest neighbor point in the static set
-

methods described in this chapter. The expectation maximization algorithm [49] is used to optimize the cost function as follows. First, we find the Gaussian from which the observed point cloud was sampled from (the expectation step). Next, we maximize the negative log-likelihood that the observed points are sampled from the GMM with respect to transformation parameters (the maximization step). The transformation matrix is calculated to maximize the Maximum A Posteriori estimation that the observed point cloud is drawn from the GMM. The GMM centroids are forced to move coherently, preserving the topological structure of the set of points [47]. Following the registration algorithm, A nearest neighbor mapping based on Euclidean distance is then used to find the closest corresponding point in the fixed cell set for each cell in the transformed cell set. The algorithm is described in Algorithm 2. An illustration of the progression of the iterative Coherent Point Drift algorithm is shown in Figure 3.9.

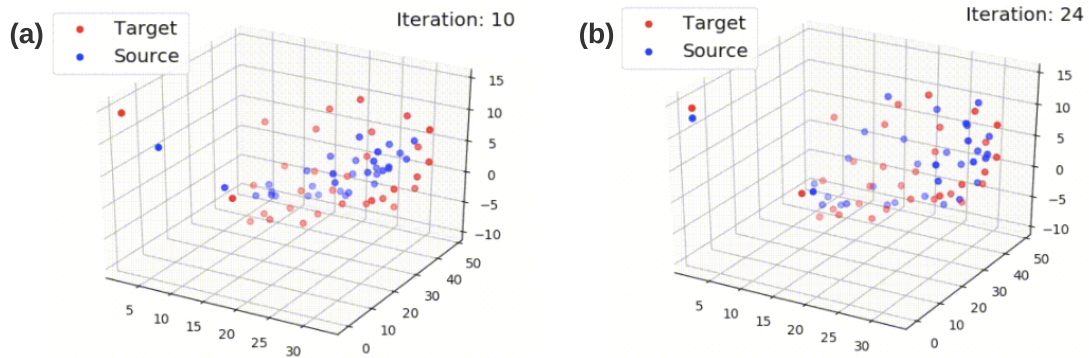


Figure 3.9: For one sample of the photo-receptor neurons: (a) An earlier iteration in the Coherent Point Drift process, and (b) Convergence of the Coherent Point Drift Registration Process. The EM-derived point locations are in red (Target) and the transformed fluorescent microscopy points are in blue (Source).

3.3.3 Registration Consistency Analysis

Each dataset containing neurotransmitter information is registered to every other dataset of the same type using the algorithm detailed above, completing all possible combinations of pairwise registrations. The EM-registration based cell assignments of each cell in both sets is then compared to each other to see if they agree. A confusion matrix [50] is generated from the results to view the consistency of each registration by cell. The confusion matrix shows the number of times a cell assignment in one dataset corresponds with each other cell assignment in another dataset.

3.4 Results

Confusion matrix

As seen in Figure 3.10, good consistency was achieved in the registration process for the photoreceptor(PR) neurons. We are able to discern that for the PR-I cells, PR-9

expressed VGAT and PR-10 expressed both VGAT and VGLUT, while the rest of the PR-I neurons expressed only VGLUT. For the PR-II cells, we are able to discern that they all expressed VGAT, with some of them also expressing VGLUT. However, the confusion matrix results showed some ambiguity in the identities of the PR-II neurons, so we are not able to pin down exactly which cells expressed VGAT only and which cells had dual expression. Based on our results, we tentatively conjecture that PR-b and PR-e may be the cells expressing solely VGAT.

Our results are not as clear cut for the relay neurons (RNs). The confusion matrix in Figure 3.11 shows the inconsistency of cross-registration results for these cells.

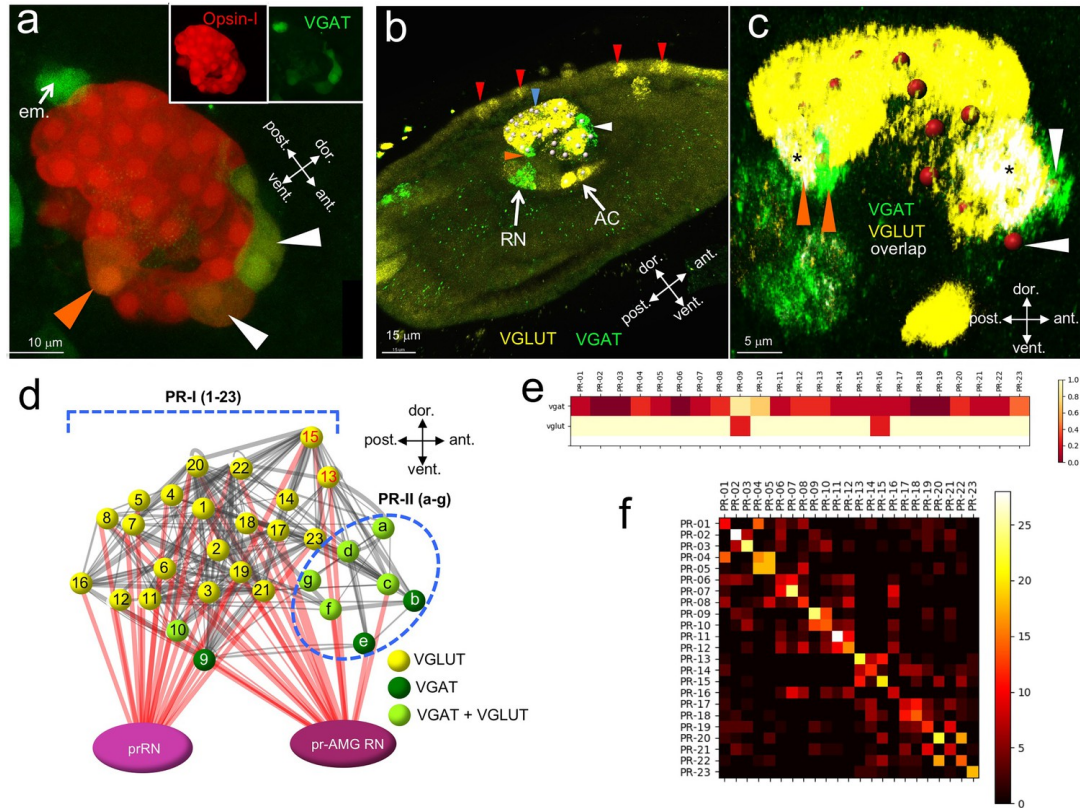


Figure 3.10: From [33]. (a) Coexpression of opsin and VGAT reporter constructs in the ocellus (white and orange arrowheads). Insets show expression of Opsin-1 and VGAT individually. (b) Expression of VGLUT and VGAT in the brain vesicle and epidermis by in situ hybridization. VGAT was observed in an anterior (white arrowhead) and posterior (orange arrowhead) domain of the ocellus. Blue arrowhead indicates VGLUT expression in the ocellus, and red arrowheads indicate VGLUT-expressing epidermal sensory neurons. (c) Posterior VGAT-expression in the ocellus consists of two cells (orange arrowheads), one exclusively expressing VGAT, and one coexpressing VGAT and VGLUT. Two cells in the anterior exclusively express VGAT (white arrowheads). Nuclei are shown as red spheres. Asterix indicate overlap of VGAT and VGLUT. (d) Neurotransmitter predictions color-coded on a schematic diagram of the ocellus photoreceptors. Lines between photoreceptors indicate chemical synaptic connections taken from [51], with red lines indicating projections to the relay neurons. (e) Heat map of neurotransmitter predictions from registration for photoreceptor group I (cells 01–23). Scale assigns color to proportion of iterations predicting VGAT or VGLUT within a particular cell. (f) Confusion matrix of registration of photoreceptor group I cells (cells 01–23). High values (light colors) in the diagonal indicate higher confidence. Abbreviations: dor., dorsal; vent., ventral; ant., anterior; post., posterior; em., eminent cell; RN, relay neuron; AC, antenna cells; pr-AMG RN, photoreceptor ascending motor ganglion relay neuron; prRN, photoreceptor relay neuron; VGAT, vesicular GABA transporter; VGLUT, vesicular glutamate transporter; PR-I, photoreceptor group I (01–23).

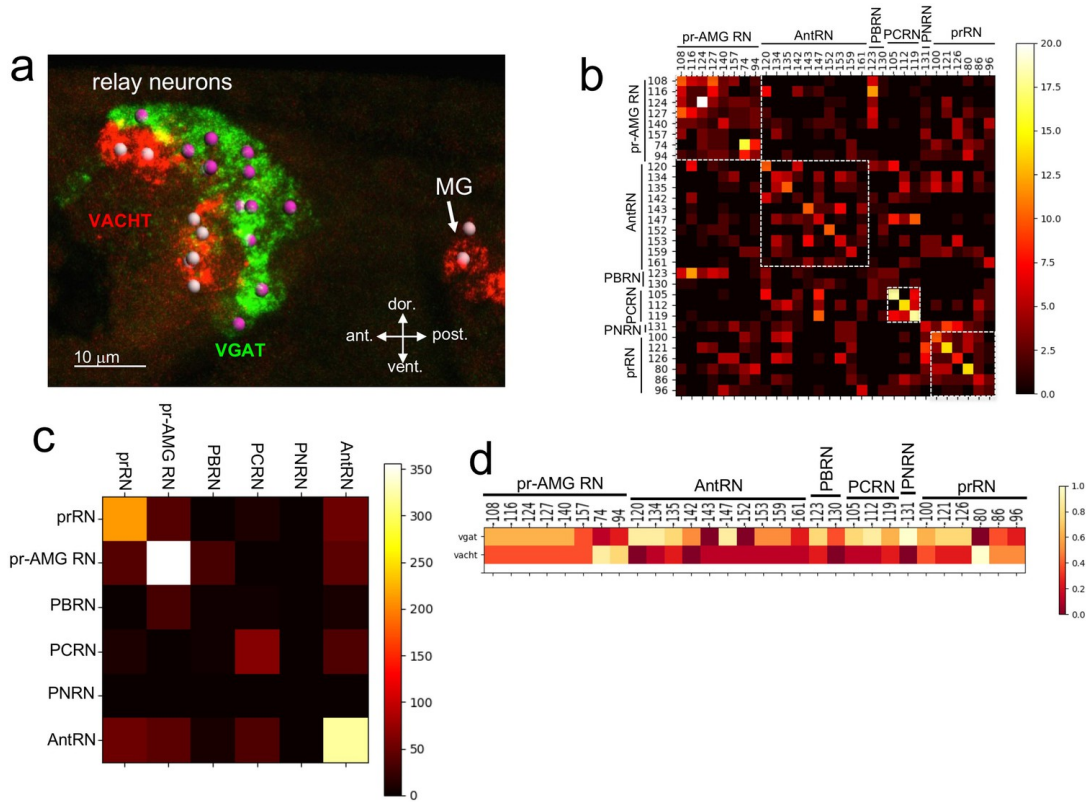


Figure 3.11: Neurotransmitter use in the relay neurons. From (a), VGAT expression seems to be higher posteriorly while VACHT expression seems to be more anterior-ventral. However, from (b) it is evident that point-cloud registration was not able to consistently match cells from in-situ hybridization with EM cell centroids. (c) shows that while point cloud matching was not able to resolve individual cell identities, it was reasonably able to stay consistent within cell types. If we assume within-type consistency, we can predict from (d) that AntRN are primarily VGAT, as well as PCRN and PNRN.

3.5 Summary

On a cellular and sub-cellular level of granularity, multiple modes of microscopy are used to image cells for different types of information. Confocal light microscopy on in-situ hybridized specimens allows us to view the types of neurotransmitters expressed in discrete regions of the brain, while electron microscopy allows us to view the connections between each neuron via sub-cellular synaptic structures.

The neuronal model we are able to build from the results of computational analysis on in-situ hybridization of *Ciona Intestinalis* support a model for two parallel visuomotor pathways, one mediated by the PR-Is and sensitive to the direction of light, and the other mediated by the PR-IIs and sensitive to changes in ambient light.

From in situ hybridization we observed that some PR-IIs exclusively express VGAT, while other co-express VGAT and VGLUT. The significance of VGAT/VGLUT coexpression in the *Ciona* visuomotor pathway is not yet clear, although similar coexpression is widely observed in mammalian brains (Fattorini et al., 2015; Zander et al., 2010) and invertebrates (Fabian-Fine et al., 2015). It is speculated that co-release of GABA and GLUT may serve to tune excitatory/inhibitory balance. While the connectome shows that not all of the PR-IIs project to the RNs, with a subset instead forming extensive connections to other PR-Is and PR-IIs, the connectome indicates that several of the VGAT-exclusive PR-IIs do project to the pr-AMG RNs (Figure 2e and Ryan et al., 2016), consistent with our hypothesis that the PR-II output to the pr-AMG RNs is predominantly inhibitory.

In order to determine with greater certainty the neurotransmitter expression of the relay neurons, more analysis needed to be done on the multimodal dataset. The next chapter will describe an automated synapse detection and classification method based on the information we have gained from the current chapter.

Chapter 4

Synapse Level Analysis: Ciona Connetome

The point cloud registration presented in Chapter 3 allowed us to determine the neurotransmitter expressions of a subset of neurons in the Ciona nervous system. However, a large portion of neurons remain unclassified due to variations in configuration and experimental limitations. In order to draw further conclusions about neurons in the Ciona brain, the synapses visible in electron microscopy (EM) are to be further examined. More specifically, we would like to associate the synaptic structure with the neurotransmitter valence (excitatory or inhibitory). This is the first attempt, to the best of our knowledge, to connect Ciona synaptic structure with their valence. Towards this, we present a convolutional neural network to predict the neurotransmitter valence of a synapse based on its appearance in electron microscopy images. The ability to determine neurotransmitter expression directly from EM images allows us to confirm the results of experiments on behavior, as well as in-situ hybridization results. Because each step introduces potential errors, directly predicting neurotransmitter expression from EM images can reduce misidentification of individual cells. With more training data and validation, we may be

able to bypass the time and resource-consuming experiments.

This chapter presents the algorithms and computational methods used to analyze and predict the function and form of sub-cellular EM structures based on the multi-modal data introduced in the previous chapter. A deep learning framework for image detection and classification is detailed. Statistical and computational methods are used to verify algorithm performance and consistency, and attention maps are used to discover important features to the deep learning model. The results and analysis of the deep networks are presented. Finally, a summary of the approach and results are given, along with commentary on the results and future directions.

4.1 Prior Work

Previous work have reported successful application of computer vision methods for automatically detecting synapses in EM images of *Drosophila*, mouse and rabbit neurons [52, 53, 54]. However, the synapses of these organisms contain unique features which the detection systems rely on heavily. The algorithm for *Drosophila* synapse detection [52] uses primarily the t-shaped feature and postsynaptic density to detect the synapse with a 3D UNet [55], while the algorithms for mouse synapse detection [53], [54] was created for cryogenic EM. One approach applies pixel-level classification and graph cut segmentation on EM images to identify potential synapses, then filters the potential synapses with a random forest classifier [53]. The classifiers are trained on manually annotated samples. A second approach finds handcrafted synaptic cleft features for the presynaptic and postsynaptic regions, and uses LogitBoost to perform the final synapse detection [54]. A third publication describes the uses of a fusion of ribbon, cleft, and vesicle features of a rabbit retina synapse to detect retinal synapses through kernel learning [56]. Ribbons are not ubiquitous amongst synapses and clefts are not always visible for different types

of synapses, so this method would not work on all types of synapses.

Studies have qualitatively shown differences between excitatory and inhibitory synapses - namely that the vesicle shape and post-synaptic density appearance varied between the two functions [57, 58, 59]. However, these papers are mostly qualitative, and do not provide predictive functionality. Furthermore, these studies do not analyze a large number of synapses, and are not directly translatable to Ciona synapses due to studying different species. More recently, synapse classification using a 3D UNet was done on EM images of *Drosophila* neurons [60]. This body of work is most similar to the one presented in this chapter, but differs in granularity, dealing with large groups of neurons approximately ten times greater in number than Ciona. In contrast, we seek to resolve the neurotransmitter class of individual neurons and verify predictions with prior knowledge derived from a variety of experiments in other modalities. Our previous study [33] combined in situ hybridization with the existing connectome derived from EM to determine the neurotransmitter expression of neuron types, such as the photoreceptors. Point cloud matching was done to match relative cell locations in 3 dimensions between fluorescent microscopy and electron microscopy results. While we are successful in determining the neurotransmitter expression of individual cells belonging to the photoreceptors, we are unable to determine with certainty the neurotransmitter expression of individual cells belonging to the relay neurons and other types. The present study aims to take steps towards resolving the neurotransmitter assignments in these ambiguous regions while applying neural network approaches to this problem.

4.2 Synapse Detection and Classification

We propose a deep learning convolutional neural network to detect Ciona *intestinalis* synapses from serial section transmission electron microscopy (EM) images, as well

as a multi-modal method of predicting neurotransmitter expression based on several modalities of data obtained in different ways - EM imaging, light microscopy of in situ hybridization, and behavioral observation experiments.

Our contributions to the synapse classification problem are as follows:

- An automated method to detect and localize synapses from EM data with high accuracy.
- A method to predict the neurotransmitter class of neurons in Ciona based on its synapse structure. Creation and analysis of class activation maps (CAM) from the neural network to derive the synaptic features which are identified as important to neurotransmitter prediction.
- Model-based predictions of neurotransmitter class by cell which are previously unknown, and can be used in further experiments to help determine the true neurotransmitter expression of said cells.

4.2.1 Algorithm Details

The major steps involved in Ciona synapse detection and classification are shown in Figure 4.1. The synapse detection workflow includes image patch extraction and training of a ResNeXT network. The neurotransmitter prediction workflow is similar, but sorts the images based on the assigned IDs [1] of the presynaptic cell before training the deep learning network. Post-prediction class activation maps of the convolutional neural network are computed for better understanding of important imaging features for classification. Feature maps are reduced to 2 dimensional space and plotted to visualize the feature-space distance between synapses of various neurotransmitter expressions.

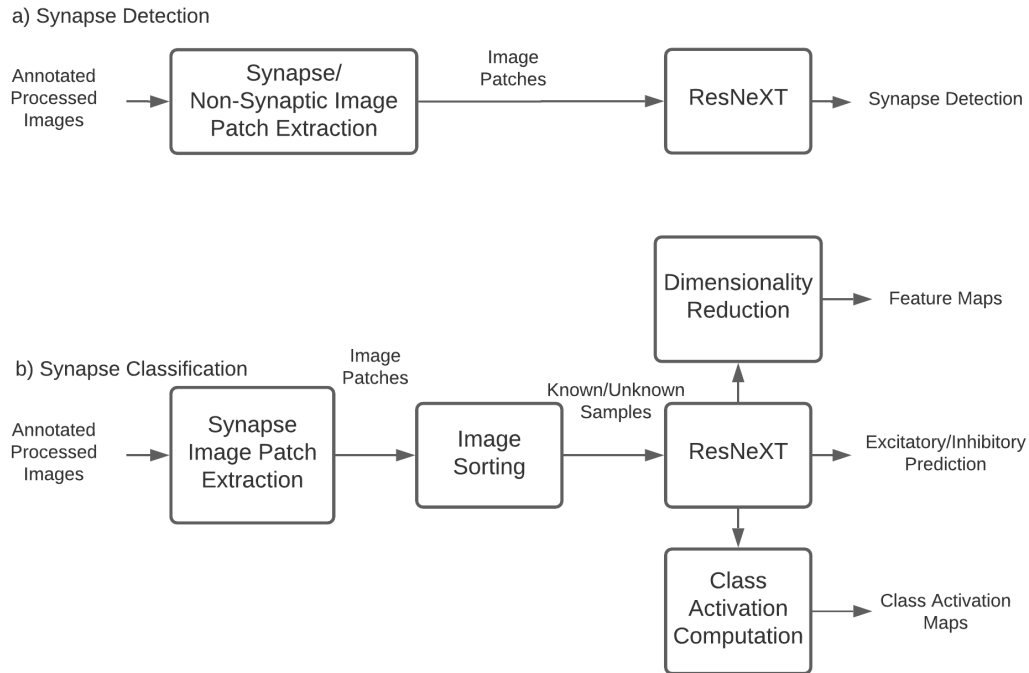


Figure 4.1: Synapse Detection / Classification flowchart.

For both synapse detection and classification, the training process is used as follows. A ResNeXt-50 network architecture [61] pretrained on the ImageNet [62] dataset was re-trained on the extracted image patches. First, the last fully connected layer was replaced with a randomly initialized fully connected layer with an input of 2048 and an output of 2. The output has 2 possible classes, with 0 being inhibitory and 1 being excitatory. The training was done in two stages. First, all layers of the ResNeXt are frozen, with no gradient, except for the last fully connected layer. The model is trained for 100 epochs, and the best model is used for the second round of training. On the second round of training, the entire network is unfrozen, every layer of the network is tuned with retraining for 200 epochs. The architecture of the deep network is shown in Table 4.1.

stage	layers		output
conv1	7x7, 64	stride 2	112x112
conv2 (x3)	3x3 max pool	stride 2	56x56
	1x1, 128	C = 32	
	3x3, 128		
	1x1, 256		
conv3 (x4)	1x1, 256	C = 32	28x28
	3x3, 256		
	1x1, 512		
conv4 (x6)	1x1, 512	C = 32	14x14
	3x3, 512		
	1x1, 1024		
conv5 (x3)	1x1, 1024	C = 32	7x7
	3x3, 1024		
	1x1, 2048		
	global average pool 1000-d fc, softmax		prediction (0 or 1)

Table 4.1: ResNeXt-50 architecture.

Synapse Detection

Image patches containing annotated synapses and image patches containing non-synaptic structures are used for the training (80%) and testing (20%) sets. 1,400 500x500 image patches containing synapses, and 1,392 500x500 image patches containing non-synaptic structures, are used for training a synapse detection network. Non-synaptic structures used are annotated variously as botryosomes, coated vesicles, basal bodies, and autophagosomes¹. Annotated structures which are avoided due to the possibility of synapse presence included terminals, vesicles, dense core vesicles, and gap junctions. Misclassified patches are analyzed visually for confounding factors.

¹Botryosomes are compartment-like organelles resembling a cluster of vesicles [63]. Coated vesicles are vesicles which transport lipids and protein and are coated with proteins on a membrane [64]. Autophagosomes are spherical structures with double membranes which are used to degrade damaged or expired organelles and proteins [65].

Synapse Neurotransmitter Prediction

Image patches are grouped by presynaptic and postsynaptic cell ID [1] before splitting into training, validation, and testing sets. A similarity computation was applied to each patch to ensure that no duplicates or extremely overlapping patches are used. Due to the limited number of synapses and known neuron types available, we combined the two inhibitory neurotransmitter expressions (glycine, gamma-Aminobutyric acid (GABA)) and the two excitatory neurotransmitter expressions (glutamate, acetylcholine). 470 excitatory synapses and 338 inhibitory synapses are used for training a synapse detection network. 1246 excitatory synapses and 396 inhibitory synapses are used for testing. Each synapse was composed of 3-10 image patches in the z dimension.

9 neuron types with 4 known neurotransmitters, as determined by in situ hybridization [33], are used in training and testing. These types and their neurotransmitter expressions are shown in table 4.5. Each presynaptic neuron had 1-41 associated synapses. Neurotransmitter class was predicted using the trained network on an additional 13 neuron types with previously unknown neurotransmitter expressions. For each presynaptic neuron, a majority vote was made from the predictions of each synapse belonging to the neuron. Based on the strength of consensus, network confidence, and prior knowledge from in situ hybridization experiments [33], predictions are made on the neurotransmitter expression of neurons which surpassed a confidence interval. For each neuron, we tallied the number of predictions for each neurotransmitter valence. This tally is referred to as a vote. The valence with the most votes is chosen as the raw prediction of that presynaptic neuron. If the votes for each valence are close (e.g., inhibitory vote is not more than $e^{1.02}$ times of the excitatory votes) or that the total number of votes is less than 3, we determine the prediction to be inconclusive.

Class Activation Maps

512x512 image patches with labels are passed through the neurotransmitter prediction neural network. The final layer of our network uses global average pooling to go from a size of 16x16x2048 to 2048 features, then compute probabilities for each of the N classes using a fully connected layer without bias. For each pixel in the 16x16 feature map, we then compute the amount that this feature contributes to the output class. We apply the fully-connected weights to the features at each pixel, cutting out the average pooling step. The basic idea for deriving class activation maps is described in [66]. After extracting the activation maps from the model, we use connected component analysis to find seed points, then employ a watershed algorithm to segment the activation map into disjoint regions to further analyze spatial and intensity information about the activations. Next, we removed regions that are smaller than a set threshold, which was set empirically to 5000 pixels. Following this step, for each remaining connected region, we then calculated the x and y coordinates of the centroid for the region and the average activation intensity inside the region of interest.

Feature Maps

To better understand the network's decision boundaries, the high-dimensional features from the output layer of the prediction network are reduced to 20 dimensions using Principal Component Analysis. The 20 dimensional feature is then further reduced to 2 dimensions using t-distributed Stochastic Neighbor Embedding [67], which better visualizes the clustering characteristics of the features. The features are grouped in various ways to gain more insight into how they are clustered.

4.3 Results

4.3.1 Model Performance

For synapse detection, a training accuracy of 0.99 and a testing accuracy of 0.98 was achieved. For training, 2780 samples are used, with half of the samples containing synapses and half without synapses. For testing, 692 samples are used, also with a 50/50 split of synapse vs. non-synapse sample. There are no false negatives, but there are 13 false positives which are detected by the model. Examples of some false positives are shown in Figure 4.2.

Performance of the network on classification is shown in Table 4.2. Precision is the number of True Positives (TP) divided by the sum of TP and True Negatives (TN), and recall is TP divided by the total number of positive samples. Precision determines how often selected items are relevant, and recall determines how often relevant items are selected.

Of the failed detections (all false positives), 10 cases are image patches annotated as coated vesicles, 2 cases are annotated as botryosomes, and 1 case annotated as an autophagosome. Some representative patches are included in Figure 4.2. From the failed cases, it can be seen that they tend to contain cell boundaries and vesicles, which are features associated with synapses. The lack of false negatives is reassuring, as the goal of the detection network is to detect likely synapses for screening by experts, so false positives are better tolerated than false negatives.




Figure 4.2: Examples of false positives for the synapse detection network. a), b), and c) contain coated vesicles, d) and e) contain botrysome, and f) contains an autophagosome. It can be seen that either cell boundaries (blue arrows) or groups of vesicles (red arrows) are visible in many of the cases.

	Precision (per synapse)	Recall (per synapse)	Precision (per cell)	Recall (per cell)
inhibitory	0.98 train / 0.98 test	0.93 train / 0.91 test	1.0	0.95
excitatory	0.96 train / 0.93 test	0.99 train / 0.99 test	0.98	1.0

Table 4.2: Automated synapse classification performance - per synapse (train and test) and per cell.

	Number of Cells	Number of Synapses	Accuracy (per cell)	Accuracy (per synapse)
gly	3	22	0.67	0.71 train / 0.6 test
gaba	19	457	1.0	0.92 train / 0.95 test
ach	19	292	1.0	0.97 train / 1.0 test
glut	23	342	1.0	1.0 train / 0.98 test
overall	64	1113	0.98	0.95 train / 0.97 test

Table 4.3: Automated synapse classification performance breakdown by neurotransmitter expression. As seen in the table, the performance for glycine was the worst, most likely due to the low number of cells and synapses available to train the model. Performance for GABA, acetylcholine, and glutamate are similar to each other.


4.3.2 Model Analysis

To get a better idea of the features which are important to the neurotransmitter predictor network, we visualized the activation maps of the network on different classes, as described in section 4.2.1. The results of the visualization show that cell boundaries, vesicles, and postsynaptic density are the main focus of the majority of the attention for the trained network. Some examples of activation maps are shown in Figure 4.3

The feature maps derived from the model outputs are shown in Figures 4.4 - 4.6.

One potential failure mode of the deep learning model is the possibility of picking up on features which are specific to cell type (as defined in section 3.2.1), rather than neurotransmitter expression. From the feature maps, it is evident that the synapses are placed in clusters which mostly correspond to their neurotransmitter class, inhibitory or excitatory. From Figure 4.6 it appears that the differences between different excitatory neurotransmitters (acetylcholine and glutamate) are also captured by the model, even though this information was not explicitly included during training. From Figure 4.4, it can be seen that the features of synapses tend to be spread out throughout the feature space, and co-mingle amongst cell types. This is promising, because this indicates that the model is picking up on differences between synapses which are more indicative of neurotransmitter class, rather than type-specific differences.

a) Activation Maps for Excitatory Neurons



b) Activation Maps for Inhibitory Neurons




Figure 4.3: Processed Class Activation Maps for Excitatory and Inhibitory Synapse Image Patches. The network seems to pay more attention to the vesicles when predicting inhibitory neurons, and the cell boundary when predicting excitatory neurons. The blue arrows indicate cell boundaries while the red arrows indicate vesicles. The green outline shows the main region of interest of the network.

Table 4.5 shows the performance of the model on synapses belonging to cells of known neurotransmitter expression. Some of these synapses are used for training, and others are used for testing. We tallied the number of predictions by presynaptic neuron type and valence (inhibitory and excitatory) in the training set. The tally is repeated for the testing set. The tally for each valence is referred to as a vote. The valence with the most votes is chosen as the raw prediction of that presynaptic neuron. If the votes for each valence are close (e.g., inhibitory vote is not more than $e^{1.02}$ times of the excitatory votes) or that the total number of votes is less than 3, we determine the prediction to be




Figure 4.4: Visualization of features for 30 neuron types after reduction to 2 dimensions using principal component analysis and t-distributed stochastic neighbor embedding. Each data point on the plot is the computed feature of a synapse, with synapses that span multiple frames averaged across all frames. There is quite a bit of intermingling of the features between cell types, which is an encouraging sign that the model is picking up on differences not unique to each cell type.

inconclusive. Those are the values in the Predicted Valence section. We compared the valence of the raw prediction with the neurotransmitter expression determined by in situ hybridization to calculate precision and recall.

Table 4.4 shows the predictions of the model on relay neurons of unknown neurotransmitter expression. We noticed that the prediction model tends towards predicting more excitatory synapses in the relay neuron type, since the observed average number of excitatory relay neurons from in situ hybridization was 11, while the number of predicted excitatory relay neurons was 14 [33]. All by one of the pr-AMG relay neurons was predicted to be excitatory, with varying degrees of likelihood. This is different from our predictions in [10], which indicated an inclination towards inhibitory pr-AMG relay neurons, but with low confidence. However, the predicted number of excitatory and in-




Figure 4.5: Feature visualization for synapses belonging to cells with known neurotransmitter expression, grouped by valence. Two distinct groups can be seen, one primarily composed of excitatory synapses and one primarily composed of inhibitory synapses.

inhibitory neurons in the relay neuron type is closer to the observed number in [33] than the predicted results from 3D point cloud matching, which suggests a promising direction for resolving the neurotransmitter expressions in that region. Further experiments and analysis is needed to determine with certainty the neurotransmitter expression of each relay neuron.

The full prediction table is included in Table 4.5. Summaries of the available data are shown in Figures 4.7 and 4.8.




Figure 4.6: Feature visualization for synapses with known neurotransmitter expression, grouped by neurotransmitter. Even though the prediction model was trained only to differentiate between excitatory and inhibitory synapses, it can be seen that the features tend towards separation by neurotransmitter expression, with the exception of glycine, likely do to the lack of training samples.




Figure 4.7: Number of unique cells per cell type of selected groups of interest.




Figure 4.8: Frequency of cells with various number of synapses. It can be seen that the majority of cells have between 10-20 synapses.

4.3.3 Comparison with Manual Overlay

The crux of our analysis lies in our hypothesis that the same cells appear in similar locations across Ciona specimens. To provide an additional point of comparison, we manually overlay the fluorescent imaging results from in-situ hybridization with the cell centroids provided by the annotated EM data, as seen in Figure 4.9. Using the Unity software [68], the centroid and volume of each neuron in the four selected cell types as given in [1] is rendered in 3D. Each neuron is approximated with a sphere of its corresponding volume. As seen on the leftmost image in Figure 9, image stacks of in situ from [33] which contain expression of both vesicular GABA transporter (VGAT, inhibitory) and vesicular Acetylcholine transporter (VACHT, excitatory) in the posterior Brain Vesicle are also rendered into the software with real-world dimensions. Since certain parts of the VGAT structure are well known and consistent, such as in the photoreceptors

and at the posterior end of the relay neurons, this was used to manually align the in situ with the connectome. The matching criteria is as follows: the posterior border doesn't pass the most posterior antRN, the dorsal cap marks the Eminens cells, and the two patches, a smaller posterior one and larger anterior one, on the right mark the two photoreceptor types, PR-I (only pr9 and pr10) and PR-II, respectively. After this alignment is done, the smaller VACHT-labelled regions are brought into view for analysis. 7 in situs from [33] are aligned using the mentioned structures and similarity across the in situs as guides. Once they are aligned, a collision detector is used to compute the number of voxels in contact with each neuron. An illustration of the matching process is shown in Figure 4.9, and the comparative results are shown in Table 4.4.




Figure 4.9: Illustration of manual overlay of in-situ hybridization results for VGAT (left image) and EM-derived cell centroids (middle image) in 3D. Of the four cell types shown, The areas with VGAT are encapsulated by the cell models, as seen in the rightmost image. See text for more details.

Name	cell type	Excitatory Voxel	Inhibitory Voxel	Excitatory Score	Inhibitory Score	Network Prediction	Overlay Estimation
90	Bipolar prIN	6234	532	1.10E-02	1.40E-03	excitatory	excitatory
92	Bipolar prIN	6767	259	1.20E-02	6.79E-04	excitatory	excitatory
93	non-sensory RN	308	905	5.44E-04	2.37E-03	inconclusive	inconclusive
103	non-sensory RN	5	867	8.84E-06	2.27E-03	excitatory	inconclusive
106	non-sensory RN	51	1924	9.01E-05	5.05E-03	inconclusive	inhibitory
122	non-sensory RN	25	2113	4.42E-05	5.54E-03	excitatory	inhibitory
125	non-sensory RN	83	12696	1.47E-04	3.33E-02	inhibitory	inhibitory
4	PNIN	0	0	0	0	excitatory	inconclusive
6	PNIN	0	0	0	0	excitatory	inconclusive
20	PNIN	0	0	0	0	excitatory	inconclusive
29	PNIN	0	0	0	0	excitatory	inconclusive
30	PNIN	0	0	0	0	excitatory	inconclusive
85	PNIN	6909	484	1.22E-02	1.27E-03	excitatory	excitatory
61	PNIN	1673	0	2.96E-03	0	excitatory	excitatory
65	PNIN	8273	146	1.46E-02	3.83E-04	excitatory	excitatory
88	PNIN	13979	236	2.47E-02	6.19E-04	excitatory	excitatory
131	PNRN	1	6828	1.77E-06	1.79E-02	excitatory	inhibitory
74	pr-AMG RN	1665	0	2.94E-03	0	excitatory	excitatory
94	pr-AMG RN	8900	9602	1.57E-02	2.52E-02	excitatory	inconclusive
108	pr-AMG RN	5903	1558	1.04E-02	4.09E-03	excitatory	inconclusive
116	pr-AMG RN	7398	423	1.31E-02	1.11E-03	inconclusive	excitatory
124	pr-AMG RN	6215	6184	1.10E-02	1.62E-02	excitatory	inconclusive
127	pr-AMG RN	871	6454	1.54E-03	1.69E-02	excitatory	inhibitory
140	pr-AMG RN	1485	7463	2.62E-03	1.96E-02	excitatory	inconclusive
157	pr-AMG RN	941	3979	1.66E-03	1.04E-02	inconclusive	inconclusive
123	pr-BTN RN	6077	2991	1.07E-02	7.84E-03	inconclusive	inconclusive
130	pr-BTN RN	312	84	5.51E-04	2.20E-04	inhibitory	inconclusive
105	pr-cor RN	2953	760	5.22E-03	1.99E-03	inconclusive	excitatory
112	pr-cor RN	197	6530	3.48E-04	1.71E-02	inconclusive	inhibitory
119	pr-cor RN	236	3234	4.17E-04	8.48E-03	inconclusive	inhibitory
80	prRN	2301	1352	4.07E-03	3.55E-03	inconclusive	inconclusive
86	prRN	4822	1696	8.52E-03	4.45E-03	inconclusive	inconclusive
96	prRN	3660	4863	6.47E-03	1.28E-02	excitatory	inconclusive
100	prRN	174	11273	3.08E-04	2.96E-02	inconclusive	inhibitory
121	prRN	751	7371	1.33E-03	1.93E-02	inconclusive	inhibitory
126	prRN	1096	4588	1.94E-03	1.20E-02	inconclusive	inhibitory

Table 4.4: Predicted neurotransmitter valence for the relay neurons (RNs). RNs used for training with known neurotransmitter expression are omitted from this table and included in the appendix. The column with our model predictions is ‘Network Prediction’, which is ‘inconclusive’ if there are fewer than 3 synapses for a given cell, or if the number of excitatory and inhibitory predictions are similar, that is, one is not more than 1.5 times the other.

cell id	in situ	llr test	llr train	excitatory test	inhibitory test	excitatory train	inhibitory train	prediction test	prediction train
ACIN1L	gly	0.24	-1.73	1	0	0	6	excitatory	inhibitory
ACIN2L	gly	-0.68	0.12	1	2	4	3	inhibitory	excitatory
ACIN2R	gly	-0.46	-0.54	0	1	1	3	inhibitory	inhibitory
AMG1	gaba	-1.08	-1.66	0	5	1	9	inhibitory	inhibitory
AMG2	gaba	-0.67	-3.09	1	4	1	14	inhibitory	inhibitory
AMG3	gaba	-0.21	-3.59	0	1	0	14	inhibitory	inhibitory
AMG4	gaba	-0.47	-2.42	0	2	2	12	inhibitory	inhibitory
AMG5	ach	2.49	2.87	8	0	14	2	excitatory	excitatory
AMG6	gaba	-0.49	-3.55	1	2	1	14	inhibitory	inhibitory
AMG7	gaba	-0.6	-3.1	0	3	1	16	inhibitory	inhibitory
Ant1	glut	0.21	3.89	3	1	14	0	excitatory	excitatory
Ant2	glut	2.46	7.25	9	0	25	1	excitatory	excitatory
120	gaba	-1.56	-11.75	0	4	1	29	inhibitory	inhibitory
134	gaba	-1.58	-6.74	0	5	1	20	inhibitory	inhibitory
135	gaba	-1.15	-5.12	0	3	1	19	inhibitory	inhibitory
142	gaba	-1.54	-7.14	0	5	2	23	inhibitory	inhibitory
143	gaba	-1.75	-6.55	0	5	3	20	inhibitory	inhibitory
147	gaba	-2.94	-7.12	0	7	0	20	inhibitory	inhibitory
152	gaba	-2.58	-5.79	0	8	3	21	inhibitory	inhibitory
153	gaba	-1.52	-3.36	0	5	1	10	inhibitory	inhibitory
159	gaba	-1.79	-5.6	0	5	2	16	inhibitory	inhibitory
161	gaba	-3.02	-5.22	0	6	1	16	inhibitory	inhibitory
ddNL	ach	1.39	4.69	4	0	16	0	excitatory	excitatory
ddNR	ach	1.37	3.79	4	0	12	0	excitatory	excitatory
Em1	gaba	-1.11	-7.92	1	6	3	31	inhibitory	inhibitory
Em2	gaba	-1.6	-6.26	1	6	3	24	inhibitory	inhibitory
MGIN1L	ach	0.87	6.27	5	0	25	1	excitatory	excitatory
MGIN1R	ach	0.43	6	3	0	22	0	excitatory	excitatory
MGIN2L	ach	0.82	4.85	3	0	14	0	excitatory	excitatory
MGIN2R	ach	2.01	3.88	6	0	14	1	excitatory	excitatory
MGIN3L	ach	0.33	2.91	1	0	7	0	excitatory	excitatory
MGIN3R	ach	0.68	2	2	0	7	1	excitatory	excitatory
MN1L	ach	1.13	5.77	3	0	18	0	excitatory	excitatory
MN1R	ach	2.9	3.85	8	0	16	1	excitatory	excitatory
MN2L	ach	1.84	3.46	4	0	9	0	excitatory	excitatory
MN2R	ach	1.56	4.81	4	0	11	0	excitatory	excitatory
MN3L	ach	0.43	1.91	1	0	5	0	excitatory	excitatory
MN3R	ach	0.85	2.77	3	0	7	0	excitatory	excitatory
MN4L	ach	1.91	1.02	6	0	4	0	excitatory	excitatory
MN4R	ach	0.76	2.92	2	0	7	0	excitatory	excitatory
MN5L	ach	0.4	1.34	1	0	3	0	excitatory	excitatory
MN5R	ach	0.37	2.54	1	0	6	0	excitatory	excitatory
pr1	glut	1.5	2.05	4	0	9	0	excitatory	excitatory
pr11	glut	0.67	2.41	3	0	10	0	excitatory	excitatory
pr12	glut	0.85	4.72	3	0	16	0	excitatory	excitatory
pr13	glut	1.1	2.96	3	0	11	0	excitatory	excitatory
pr14	glut	0.16	-	1	0	0	0	excitatory	excitatory
pr15	glut	1.98	7.33	7	0	27	0	excitatory	excitatory
pr16	glut	0.89	1.83	3	0	7	0	excitatory	excitatory
pr17	glut	0.44	3.35	2	0	13	0	excitatory	excitatory
pr18	glut	0.14	0.58	1	0	2	0	excitatory	excitatory
pr19	glut	1.38	1.48	5	0	7	0	excitatory	excitatory
pr2	glut	0.39	1.3	2	0	6	0	excitatory	excitatory
pr20	glut	2.92	5.94	9	0	22	0	excitatory	excitatory
pr21	glut	0.43	2.18	2	0	9	0	excitatory	excitatory
pr22	glut	0.58	2.4	3	0	10	0	excitatory	excitatory
pr23	glut	1.06	4.16	4	0	14	0	excitatory	excitatory
pr3	glut	0.54	3.06	2	0	12	0	excitatory	excitatory
pr4	glut	0.66	-	3	1	0	0	excitatory	excitatory
pr5	glut	0.85	2.41	3	0	8	0	excitatory	excitatory
pr6	glut	0.75	2.14	2	0	9	0	excitatory	excitatory
pr7	glut	1.18	4.82	4	0	16	0	excitatory	excitatory
pr8	glut	0.66	2.7	3	0	11	0	excitatory	excitatory
pr9	gaba	-0.36	-1.3	0	1	4	11	inhibitory	inhibitory

Table 4.5: Predictions on synapses with known ground truth.

4.4 Summary

The proposed synapse detection model has good performance on the detection of synapses in electron microscopy for Ciona. While there are some false positives from the model prediction (Table 4.5), this is desirable compared to false negatives, because an expert can then screen the predictions to determine true synapses. The ratio of synapses to non-synaptic structures in a typical EM of Ciona is on the order of 1:1000. Instead of scanning an entire image for possible synapses, the model can drastically reduce the annotation time needed for synapse annotation. The synapse prediction model has helped to identify possible neurotransmitter expression for cells from certain neuronal types, which are previously unknown. While we cannot be absolutely certain that the model has predicted correctly for synapses belonging to cells with previously unknown neurotransmitter expressions, the output of the model seems reasonable given the feature analysis we have done. Comparison with previous in situ hybridization results have also shown that the prediction of the model is likely correct. For the relay neurons, in [33] we had previously found an average of 16 VGAT-positive neurons and 11 VACHT-positive neurons. The prediction of the network matches these numbers well, and much better than the point-cloud registration method used in [33]. We hope that the model predictions will work in conjunction with both prior and future analysis to help resolve the neurotransmitter expression of individual neurons in regions which have been difficult to resolve using in-situ hybridization and other experimental methods.

4.4.1 Future Work

More work can be done on the analysis of cellular and subcellular features in neurons with undetermined neurotransmitter expression. Overall connectivity and cell shape can be obtained from the annotated EM data and may be useful tools in better understanding

the relationship between structural features and neurotransmitter expression expression. The connection between and synapse structure can also be further explored, and may help with a more robust neurotransmitter prediction method.

Chapter 5

Discussion

The primary goal of this dissertation is to introduce and explore the concept of relating brain structure to function. The structure of the brain is deduced from the results of multiple imaging modalities spanning a wide range of scales, while the function refers to an observed behavior in response to controlled and known stimuli, or a patient diagnosis or prognosis.

Chapter 2 introduced the problem of diagnosing Normal Pressure Hydrocephalus (NPH) and the clinical workflow for diagnosing and assessing NPH, which currently consists of a CT scan of the head and observation of symptoms, such as a shuffling gait, confusion, and urinary incontinence. A robust segmentation algorithm is presented, along with the use of network metrics to predict NPH diagnosis. The network metrics are computed from a common atlas which is modified based on the segmentation results to fit each patient.

In Chapter 3, the *Ciona* organism is introduced as a simpler chordate closely related to humans on the evolutionary tree. The structure of the *Ciona* neuronal system is described, and datasets from different modalities from the cellular level (fluorescent confocal microscopy) to the sub-cellular level (electron microscopy) are presented. The

annotation process for the data is also described.

Finally, Chapter 4 details the computational methods we employed to collate and match the neurotransmitter expression data from confocal microscopy with the cell connection maps from electron microscopy and gain new insights into the neuronal processes of Ciona. Reports on statistical and feature analysis is provided to determine the validity and consistency of the computational results.

5.1 Broader Impacts

The exploration of tractographic features for to predict patient outcome has many applications in the analysis of the brain. [17] uses tractographic features computed from brain tumor segmentations to improve patient outcome prediction. Exploration of connectivity matrix generation methods from tractography led to the implementation of a Reeb graph approach to tractography bundle formation [69, 70] and a demonstrative application to the ADNI dataset to differentiate dementia and normal subjects. There remains much more to be explored on the topic of creating features from structural information derived from brain imaging, and connecting that information with brain function, whether it be a diagnosis or a prediction of behavior.

We are in the process of implementing our brain CT segmentation algorithm on BisQue. Currently, a version of the algorithm is available on the cloud for use by clinical researchers. A graphical user interface allows the user to upload a CT scan from their own database, process it in the cloud using servers in our laboratory, and obtain and store the segmentation results, along with volumetric estimates and a diagnosis prediction of NPH or no NPH. Figure 5.1 is a snapshot of the user interface.

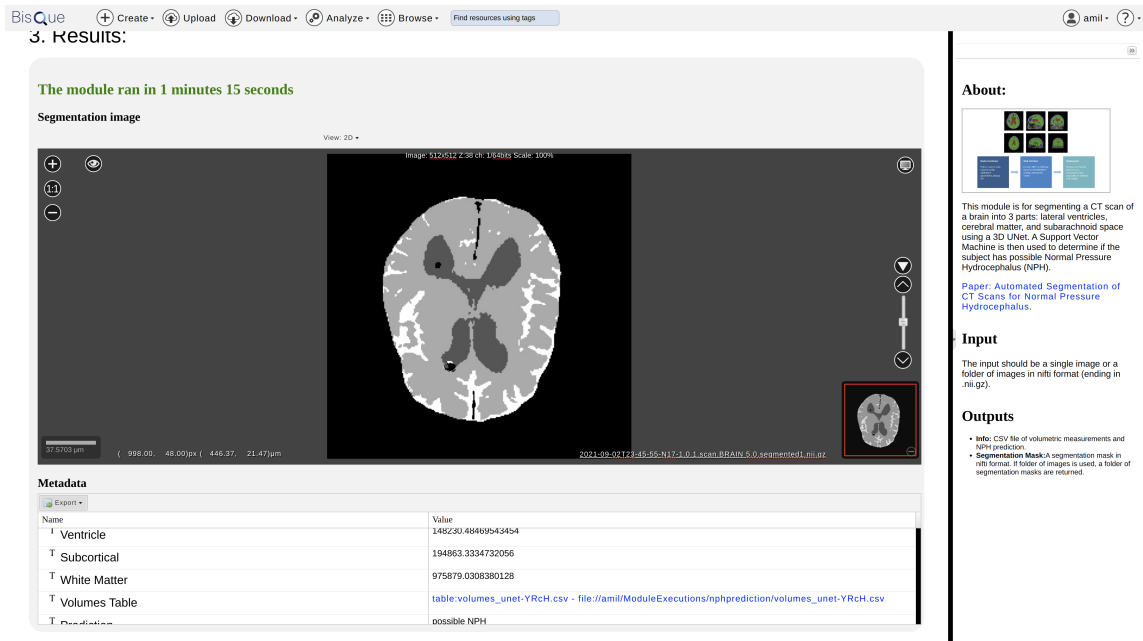


Figure 5.1: Example result of the brain CT segmentation algorithm from Chapter 1 on BisQue. The BisQue implementation of the algorithm is currently used as an initial segmentation method for CT scans in the Chen Lab in neurosurgery at the University of California, Irvine.

5.2 Future Directions

The work presented in this thesis proposed multiple methods to combine multi-modal imaging data with computational methods to predict observed behavior and function. While we developed our algorithms to accomplish specific tasks, such as predicting Normal Pressure Hydrocephalus from CT scans and determining neurotransmitter expression of Ciona neurons, the concepts and approaches used in this work can be generalized to other areas of study related to brain imaging.

On the tissue level, one promising direction is in the study of the effects of hormones on brain structure and function. The work done by [71] show that there are significant cognitive differences based on age and sex. It would be interesting to explore the structural differences in the brain correlated with hormonal changes and age through computational methods. Likewise, image processing and computer vision techniques could be applied

to analyze brain images of patients with other diagnoses, such as dementia or depression.

On the cellular level, more work can be done to verify the neurotransmitter expressions of neurons in regions where model confidence was low, or disagreed with previous results. While we had the complete electron microscopy data and annotations for one Ciona sample, it would be tremendously helpful to have a second fully annotated electron microscopy dataset of a second sample for comparison and further analysis. The annotation of a second dataset is currently in the works at Dalhousie University in Halifax, Canada. The full annotations include information about cell boundary locations for the main cell body, axons, and dendrites. Shape features can be extracted from this information and used to provide additional context for neurotransmitter expression or cell behavior information. Likewise, network metrics for each cell could be incorporated into the neurotransmitter prediction process.

Bibliography

- [1] K. Ryan, Z. Lu, and I. A. Meinertzhagen, *The CNS connectome of a tadpole larva of *Ciona intestinalis* (L.) highlights sidedness in the brain of a chordate sibling*, *Elife* **5** (Dec., 2016) e16962.
- [2] A. K. Toma, E. Holl, N. D. Kitchen, and L. D. Watkins, *Evans index revisited: The need for an alternative in normal pressure hydrocephalus*, *Neurosurgery* **68** (Apr., 2011) 939–944.
- [3] D. Jaraj, K. Rabiei, T. Marlow, C. Jensen, I. Skoog, and C. Wikkelso, *Prevalence of idiopathic normal-pressure hydrocephalus*, *Neurology* **82** (Mar., 2014) 1449–1454.
- [4] D. Shprecher, J. Schwalb, and R. Kurlan, *Normal pressure hydrocephalus: Diagnosis and treatment*, *Current Neurology and Neuroscience Reports* **8** (Sept., 2008) 371–376.
- [5] D. W. Moore, I. Kovanlikaya, L. A. Heier, A. Raj, C. Huang, K.-W. Chu, and N. R. Relkin, *A pilot study of quantitative MRI measurements of ventricular volume and cortical atrophy for the differential diagnosis of normal pressure hydrocephalus*, *Neurology Research International* **2012** (2012) 1–6.
- [6] N. Takahashi, T. Kinoshita, T. Ohmura, E. Matsuyama, and H. Toyoshima, *Automated method to compute evans index for diagnosis of idiopathic normal pressure hydrocephalus on brain CT images*, in *Medical Imaging 2017: Computer-Aided Diagnosis* (S. G. Armato and N. A. Petrick, eds.), SPIE, Mar., 2017.
- [7] P. Coupé, J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, and D. L. Collins, *Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation*, *NeuroImage* **54** (Jan., 2011) 940–954.
- [8] F. Yepes-Calderon, M. D. Nelson, and J. G. McComb, *Automatically measuring brain ventricular volume within PACS using artificial intelligence*, *PLOS ONE* **13** (Mar., 2018) e0193152.

- [9] S. Kobashi, K. Kondo, and Y. Hata, *Fully automated segmentation of cerebral ventricles from 3-d SPGR MR images using fuzzy representative line*, *Soft Computing* **10** (Jan., 2006) 1181–1191.
- [10] M. Shao, S. Han, A. Carass, X. Li, A. M. Blizt, J. L. Prince, and L. M. Ellingsen, *Shortcomings of ventricle segmentation using deep convolutional networks*, in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*, pp. 79–86. Springer International Publishing, 2018.
- [11] J. C. Cai, Z. Akkus, K. A. Philbrick, A. Boonrod, S. Hoodeshenas, A. D. Weston, P. Rouzrokh, G. M. Conte, A. Zeinoddini, D. C. Vogelsang, Q. Huang, and B. J. Erickson, *Fully automated segmentation of head CT neuroanatomy using deep learning*, *Radiology: Artificial Intelligence* **2** (Sept., 2020) e190183.
- [12] F.-C. Yeh and W.-Y. I. Tseng, *NTU-90: A high angular resolution brain atlas constructed by q-space diffeomorphic reconstruction*, *NeuroImage* **58** (Sept., 2011) 91–99.
- [13] F.-C. Yeh, D. Badre, and T. Verstynen, *Connectometry: A statistical approach harnessing the analytical potential of the local connectome*, *NeuroImage* **125** (Jan., 2016) 162–171.
- [14] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig, *User-guided 3d active contour segmentation of anatomical structures: Significantly improved efficiency and reliability*, *NeuroImage* **31** (July, 2006) 1116–1128.
- [15] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith, *FSL*, *NeuroImage* **62** (Aug., 2012) 782–790.
- [16] T. D. Bui, J. Shin, and T. Moon, *3d densely convolutional networks for volumetric segmentation*, 1709.0319.
- [17] P.-Y. Kao, T. Ngo, A. Zhang, J. Chen, and B. Manjunath, *Brain tumor segmentation and tractographic feature extraction from structural mr images for overall survival prediction*, in *Proceedings 2018 Conference on Medical Image Computing and Computer Assisted Intervention*, 11, 2018.
- [18] A. F. Agarap, *Deep learning using rectified linear units (relu)*, 1803.0837.
- [19] Y. Wu and K. He, *Group normalization*, 1803.0849.
- [20] M. Jenkinson and S. Smith, *A global optimisation method for robust affine registration of brain images*, *Medical Image Analysis* **5** (June, 2001) 143–156.

- [21] M. Jenkinson, P. Bannister, M. Brady, and S. Smith, *Improved optimization for the robust and accurate linear registration and motion correction of brain images*, *NeuroImage* **17** (Oct., 2002) 825–841.
- [22] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot, *Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain*, *NeuroImage* **15** (Jan., 2002) 273–289.
- [23] E. Bullmore and O. Sporns, *Complex brain networks: graph theoretical analysis of structural and functional systems*, *Nature Reviews Neuroscience* **10** (Feb., 2009) 186–198.
- [24] V. Caselles, R. Kimmel, and G. Sapiro, *Geodesic active contours*, *International Journal of Computer Vision* **22** (1997), no. 1 61–79.
- [25] T. F. Chan and L. A. Vese, *Active contours without edges*, *IEEE Transactions on Image Processing* **10** (2001), no. 2 266–277.
- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Édouard Duchesnay, *Scikit-learn: Machine learning in python*, *Journal of Machine Learning Research* **12** (2011), no. 85 2825–2830.
- [27] J. Platt and N. Karampatziakis, *Probabilistic outputs for svms and comparisons to regularized likelihood methods*, 2007.
- [28] A. Goldstein, A. Kapelner, J. Bleich, and E. Pitkin, *Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation*, 2013.
- [29] T. Sarwar, K. Ramamohanarao, and A. Zalesky, *Mapping connectomes with diffusion MRI: deterministic or probabilistic tractography?*, *Magnetic Resonance in Medicine* **81** (Oct., 2018) 1368–1384.
- [30] P. Poulin, M.-A. Côté, J.-C. Houde, L. Petit, P. F. Neher, K. H. Maier-Hein, H. Larochelle, and M. Descoteaux, *Learn to track: Deep learning for tractography*, in *Medical Image Computing and Computer Assisted Intervention MICCAI 2017*, 2017.
- [31] Q. Tian, B. Bilgic, Q. Fan, C. Liao, C. Ngamsombat, Y. Hu, T. Witzel, K. Setsompop, J. R. Polimeni, and S. Y. Huang, *DeepDTI: High-fidelity six-direction diffusion tensor imaging using deep learning*, *NeuroImage* **219** (Oct., 2020) 117017.

- [32] I. Benou and T. R. Raviv, *DeepTract: A probabilistic deep learning framework for white matter fiber tractography*, in *Lecture Notes in Computer Science*, pp. 626–635. Springer International Publishing, 2019.
- [33] M. J. Kourakis, C. Borba, A. Zhang, E. Newman-Smith, P. Salas, B. Manjunath, and W. C. Smith, *Parallel visual circuitry in a basal chordate*, *eLife* **8** (Apr., 2019) e44753. Publisher: eLife Sciences Publications, Ltd.
- [34] F. Delsuc, H. Brinkmann, D. Chourrout, and H. Philippe, *Tunicates and not cephalochordates are the closest living relatives of vertebrates*, *Nature* **439** (Feb., 2006) 965–8.
- [35] C. Hudson, *The central nervous system of ascidian larvae*, *Wiley Interdiscip Rev Dev Biol* (June, 2016).
- [36] K. Ryan, *The Connectome of the Larval Brain of Ciona intestinalis (L.)*. PhD thesis, 2015.
- [37] M. Helmstaedter, K. L. Briggman, S. C. Turaga, V. Jain, H. S. Seung, and W. Denk, *Connectomic reconstruction of the inner plexiform layer in the mouse retina*, *Nature* **500** (Aug., 2013) 168–74.
- [38] S. Y. Takemura, A. Bharioke, Z. Lu, A. Nern, S. Vitaladevuni, P. K. Rivlin, W. T. Katz, D. J. Olbris, S. M. Plaza, P. Winston, T. Zhao, J. A. Horne, R. D. Fetter, S. Takemura, K. Blazek, L. A. Chang, O. Ogundeyi, M. A. Saunders, V. Shapiro, C. Sigmund, G. M. Rubin, L. K. Scheffer, I. A. Meinertzhagen, and D. B. Chklovskii, *A visual motion detection circuit suggested by Drosophila connectomics*, *Nature* **500** (Aug., 2013) 175–81.
- [39] N. Randel, A. Asadulina, L. A. Bezares-Calderon, C. Veraszto, E. A. Williams, M. Conzelmann, R. Shahidi, and G. Jekely, *Neuronal connectome of a sensory-motor circuit for visual navigation*, *Elife* **3** (2014).
- [40] K. Eichler, F. Li, A. Litwin-Kumar, Y. Park, I. Andrade, C. M. Schneider-Mizell, T. Saumweber, A. Huser, C. Eschbach, B. Gerber, R. D. Fetter, J. W. Truman, C. E. Priebe, L. F. Abbott, A. S. Thum, M. Zlatic, and A. Cardona, *The complete connectome of a learning and memory centre in an insect brain*, *Nature* **548** (Aug., 2017) 175–182.
- [41] I. Larderet, P. M. J. Fritsch, N. Gendre, G. L. Neagu-Maier, R. D. Fetter, C. M. Schneider-Mizell, J. W. Truman, M. Zlatic, A. Cardona, and S. G. Sprecher, *Organization of the Drosophila larval visual circuit*, *Elife* **6** (Aug., 2017).
- [42] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller, *Whole-brain functional imaging at cellular resolution using light-sheet microscopy*, *Nat Methods* **10** (May, 2013) 413–20.

- [43] H. M. T. Choi, M. Schwarzkopf, M. E. Fornace, A. Acharya, G. Artavanis, J. Stegmaier, A. Cunha, and N. A. Pierce, *Third-generation in situ hybridization chain reaction: multiplexed, quantitative, sensitive, versatile, robust*, *Development* **145** (June, 2018).
- [44] J. C. Fiala, *Reconstruct: a free editor for serial section microscopy*, *Journal of Microscopy* **218** (2005), no. 1 52–61. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2818.2005.01466.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2818.2005.01466.x).
- [45] K. S. Arun, T. S. Huang, and S. D. Blostein, *Least-squares fitting of two 3-d point sets*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-9** (1987), no. 5 698–700.
- [46] S. Gold, A. Rangarajan, C.-P. Lu, S. Pappu, and E. Mjolsness, *New algorithms for 2d and 3d point matching: pose estimation and correspondence*, *Pattern Recognition* **31** (1998), no. 8 1019–1031.
- [47] A. Myronenko and X. Song, *Point set registration: coherent point drift*, *IEEE Trans Pattern Anal Mach Intell* **32** (Dec., 2010) 2262–75.
- [48] C. K. Williams and C. E. Rasmussen, *Gaussian processes for regression*, .
- [49] T. K. Moon, *The expectation-maximization algorithm*, *IEEE Signal processing magazine* **13** (1996), no. 6 47–60.
- [50] S. V. Stehman, *Selecting and interpreting measures of thematic classification accuracy*, *Remote Sensing of Environment* **62** (1997), no. 1 77–89.
- [51] N. Satoh, *Developmental biology of ascidians*. No. 29 in Developmental and cell biology series. Cambridge University Press, Cambridge England ; New York, 1994.
- [52] G. B. Huang, L. K. Scheffer, and S. M. Plaza, *Fully-Automatic Synapse Prediction and Validation on a Large Data Set*, *Frontiers in Neural Circuits* **12** (2018). Publisher: Frontiers.
- [53] A. Kreshuk, U. Koethe, E. Pax, D. D. Bock, and F. A. Hamprecht, *Automated Detection of Synapses in Serial Section Transmission Electron Microscopy Image Stacks*, *PLOS ONE* **9** (Feb., 2014) e87351. Publisher: Public Library of Science.
- [54] B. Staffler, M. Berning, K. M. Boergens, A. Gour, P. v. d. Smagt, and M. Helmstaedter, *SynEM, automated synapse detection for connectomics*, *eLife* **6** (July, 2017) e26414. Publisher: eLife Sciences Publications, Ltd.
- [55] O. Ronneberger, P. Fischer, and T. Brox, *U-net: Convolutional networks for biomedical image segmentation*, in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, vol. 9351, pp. 234–241. Springer

International Publishing, Cham, 2015. Series Title: Lecture Notes in Computer Science.

- [56] V. Jagadeesh, J. Anderson, B. Jones, R. Marc, S. Fisher, and B. S. Manjunath, *Synapse classification and localization in Electron Micrographs*, *Pattern Recognition Letters* **43** (July, 2014) 17–24.
- [57] H. L. Atwood, F. Lang, and W. A. Morin, *Synaptic Vesicles: Selective Depletion in Crayfish Excitatory and Inhibitory Axons*, *Science* **176** (June, 1972) 1353–1355. Publisher: American Association for the Advancement of Science Section: Reports.
- [58] C.-L. Tao, Y.-T. Liu, R. Sun, B. Zhang, L. Qi, S. Shivakoti, C.-L. Tian, P. Zhang, P.-M. Lau, Z. H. Zhou, and G.-Q. Bi, *Differentiation and Characterization of Excitatory and Inhibitory Synapses by Cryo-electron Tomography and Correlative Microscopy*, *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **38** (2018), no. 6 1493–1510.
- [59] K. Uchizono, *Characteristics of Excitatory and Inhibitory Synapses in the Central Nervous System of the Cat*, *Nature* **207** (Aug., 1965) 642–643. Number: 4997 Publisher: Nature Publishing Group.
- [60] N. Eckstein, A. S. Bates, M. Du, V. Hartenstein, G. S. Jefferis, and J. Funke, *Neurotransmitter Classification from Electron Microscopy Images at Synaptic Sites in Drosophila*, preprint, Neuroscience, June, 2020.
- [61] S. Xie, R. Girshick, P. Dollar, Z. Tu, and K. He, *Aggregated Residual Transformations for Deep Neural Networks*, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (Honolulu, HI), pp. 5987–5995, IEEE, July, 2017.
- [62] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei, *ImageNet: A large-scale hierarchical image database*, in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, June, 2009. ISSN: 1063-6919.
- [63] A. J. Hannan, P. C. Kind, and C. Blakemore, *Phospholipase c-beta1 expression correlates with neuronal differentiation and synaptic plasticity in rat somatosensory cortex*, *Neuropharmacology* **37** (Apr., 1998) 593–605.
- [64] M. Faini, R. Beck, F. T. Wieland, and J. A. Briggs, *Vesicle coats: structure, function, and general principles of assembly*, *Trends in Cell Biology* **23** (June, 2013) 279–288.
- [65] Z. Xie and D. J. Klionsky, *Autophagosome formation: core machinery and adaptations*, *Nature Cell Biology* **9** (Oct., 2007) 1102–1109.

- [66] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, *Learning Deep Features for Discriminative Localization*, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (Las Vegas, NV, USA), pp. 2921–2929, IEEE, June, 2016.
- [67] L. v. d. Maaten and G. Hinton, *Visualizing Data using t-SNE*, *Journal of Machine Learning Research* **9** (2008), no. 86 2579–2605.
- [68] J. K. Haas, *A history of the unity game engine*, .
- [69] J. Sun, M. Cieslak, S. Grafton, and S. Suri, *A reeb graph approach to tractography*, in *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*, ACM, Nov., 2015.
- [70] B. M. S. Shailja, A. Zhang, *A computational geometry approach for modeling neuronal fiber pathways*, .
- [71] S. Yu, A. P. Boone, C. He, R. C. Davis, M. Hegarty, E. R. Chrastil, and E. G. Jacobs, *Age-related changes in spatial navigation are evident by midlife and differ by sex*, *Psychological Science* **32** (Apr., 2021) 692–704.