2016

Automatic machine learning-guided methods for 3D synapse quantification in confocal neuron images

Jonathan William Sanders

Abstract

Automatic Machine Learning-Guided Methods for 3D Synapse Quantification in Confocal Neuron Images

Jonathan Sanders, M.S.
Department of Computer Science
Northern Illinois University, 2016
Jie Zhou, Ph.D., Director

This thesis explores computational methods for automatically detecting and quantifying synapses in complex 3D neuronal images. The resulting approach is a novel combination of traditional image processing, machine learning algorithms, and multi-channel comparison methods designed to overcome the unique challenges posed by these images. The methods investigated combine the strengths of each of these components in order to produce an overall method that is capable of fully detecting the synapses in large 3D confocal neuron images with minimal interaction. Human annotation of 3D neuron images remains prohibitively difficult and subjective, and computational analysis tools are highly desirable in the expanding field of 3D neuronal imaging. Validation techniques were also designed and implemented in order to test these methods for this thesis, including construction of a gold standard set of manually annotated synapse images. These are unique in their own right as there are currently no other data sets available for comparison. These methods were tested on multiple partial dendritic tree 3D images and a complete 3D dendrite with good outcomes. Quantitative validation was performed using the gold standard set to check the accuracy of synapse quantification, also with favorable results.

NORTHERN ILLINOIS UNIVERSITY
DE KALB, ILLINOIS

AUGUST 2016

AUTOMATIC MACHINE LEARNING-GUIDED METHODS FOR 3D SYNAPSE QUANTIFICATION IN CONFOCAL NEURON IMAGES

BY

JONATHAN WILLIAM SANDERS

A THESIS SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREEE

MASTER OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

Thesis Director:

Jie Zhou, Ph.D.

Acknowledgements

I would like to thank Dr. Bing Ye and his Lab members at the University of Michigan Life Sciences Institute, especially Gabriella Sterne, for providing the confocal neuronal images for this thesis as well as for their support and feedback. Their willingness to collaborate and provide these high quality images was critical to the pursuit of this thesis.

I would also like to thank Dr. Hanchaun Peng and his lab at the Allen Institute for Brain Science for the development of the open source, multi-dimensional data visualization and analysis tool Vaa3D that was used to provide 3D visualizations and construct testing data for this thesis. I would like to thank Santosh Lamichhane, Edward Hottendorf and the rest of the NIU ILAAL members and alumni for the development of BIOCAT, synapse density analysis tools, neural reconstruction tools, and other related works and for their feedback and assistance in testing.

Additionally, I would like to thank my committee members Dr. Reva Freedman and Dr. Minmei Hou for their advice, constructive criticism, and continued support. I would like to especially thank my thesis advisor Dr. Jie Zhou for the initial idea for this thesis and for her support and guidance throughout its development.

Dedication

This thesis is dedicated to my wonderful wife and daughter.

Table of Contents

Chapter Page

Chapter Page

List of Figures

Page

List of Tables

Page

Chapter 1

Introduction

Within the brain there are many distinct cell types, all fulfilling their unique role. One cell in particular is the subject of intense scrutiny and still has yet to be fully understood by modern science and medicine. This cell is the neuron, a highly specialized cell that when networked with other neurons comprises the underlying framework for all thoughts and actions. The neuron is the fundamental unit of the nervous system and brain that processes information and relays it to various parts of the body. These specialized cells form the basis for the circuitry of the brain in a similar way to the transistors used to construct modern computer hardware. There exist small gaps called synapses that serve as the junctions between neurons. Understanding synapse structure and their distribution among neurons is important to understanding neural functioning as a whole as well as neurological diseases and development (Fiala, Spacek, & Harris, 2002). Similarly, the distribution of synapses within a neuron is important to understanding the assembly, function, and plasticity of the nervous system as a whole, which has implications in many biological fields in addition to medical research (Kerschensteiner, Morgan, Parker, Lewis, & Wong, 2009; Liu, 2004; Morgan, Schubert, & Wong, 2008; Soto et al., 2011)

Capturing images of neurons is challenging, and most conventional imaging techniques have two major pitfalls: they are in only two dimensions, and they are cluttered with overlapping neurons. These images are large and complex, and attempts to analyze those using manual methods prove to be both difficult and subjective. The ability to image the distribution of synapses within a single neuron via a laser scanning confocal microscope has recently been made possible by advancements in staining techniques, but even at single-neuron resolution, manual annotation is still extremely labor intensive (Scott, Raabe, & Luo, 2002; Shrestha & Grueber, 2011). Due to natural variances in the intensity of the stain and the limitations of the imaging system, there is often residual noise and contrast variation in the neuron images. Noise in this case refers to defects in the image composed of unwanted structures and aberrations resulting from the image capture process. These factors complicate analysis and require the use of more robust detection algorithms than simple segmentation in order to extract useful information about the subcellular distribution of synapses.

Current annotation methods for analyzing these images are slow and rely heavily on manual input. This increases subjectivity and is very inefficient for a large-scale study of these images. Instead, computational tools need to be used to examine the data for the sake of efficiency and reliability. There are computational challenges associated with analyzing this data, but overcoming them is worth the gain in speed and effectiveness. Neuron structure and function is of keen interest to medical research initiatives. It is well documented that the subcellular distribution of synapses is important to the pathology of many neurological

disorders (Fiala, Spacek, & Harris, 2002; Kerschensteiner, Morgan, Parker, Lewis, & Wong, 2009; Kim et al., 2012).

The purpose of this thesis is specifically to design and implement algorithms for the accurate quantification of synapses from 3D laser confocal microscope images. These algorithms employ machine learning, multi-channel colocalization and adaptive thresholding techniques. Additionally, it was necessary to create methods and a gold standard data set for validating these algorithms, which was a complex research undertaking on its own. There are currently no other widely available gold standard data sets available for synapses in 3D confocal images.

The general organization of this thesis is as follows: first, the biological background for synapse quantification within the context of neuroscience is explored in Chapter 2. This is followed by an explanation of the imaging technology used to capture the images used for this thesis in Chapter 3. Chapter 4 is dedicated to defining other terminology necessary to understand the synapse detection methods. Chapter 5 discusses the related literature. Chapter 6 describes the methods used to detect synapses and the process of constructing a validation set to verify the success of these methods. Finally, Chapter 7 presents the qualitative and quantitative results and a discussion of these results, and Chapter 8 highlights the conclusions.

Chapter 2

Background

Neurons function by transmitting electro-chemical impulses from cell to cell via specialized gaps known as synapses. This process, called neurotransmission, involves many cellular structures and complex chemical interactions. One of the core themes in the biological sciences is the relationship between structure and function. To better understand the function of something, get a more detailed look at the structure. In this way, the structure of neurons is strongly correlated to their function in the body. The following description of neurons is still a generalization, and like all living systems there can be a great deal of variation and exceptions. Neurons are branching, thin, and are divided into three main components: the axons, dendrites, and the soma. The axon and dendrites usually branch away from the soma and are the information sending and receiving ends respectively. Signals enter the neuron on the dendrite side of the cell and are transmitted across the soma and down the axon where they are then passed on to other neurons (Lodish et al., 2000).

The signal being relayed by a neuron is carried in the form of an electrical potential across the outer membrane. This action potential is created by the controlled flow of charged ions from special channels and pumps on the surface of the neuron. In the resting state, positive sodium ions are pumped out of the cell and positive potassium ions are pumped into

the cell. The ion channels on the surface of the cell are normally closed off but are sensitive to electrical charge. Once a signal has started and the voltage changes in one region of the neuron, nearby ion channels open to allow positively charged sodium to flood back into the cell. This creates a sharp increase in voltage across the membrane. That exponentially triggers more ion channels to open. Shortly after that, the potassium channels open, allowing the positive potassium ions to flood out of the cell, equalizing the voltage once again. The signal then cascades down the surface of the neuron in this way, and the ion pumps work to move the sodium and potassium ions back to the ready state (Huxley, 2002).

The areas where neurons pass along their signal to each other are highly specialized structures known as synapses. The axons and dendrites of neighboring cells don't usually make physical contact with one another. Instead they are separated by a small gap. The electrical signal propagated down the outside of the neuron isn't normally strong enough to pass over these gaps. Instead, complex chemical machinery takes over in order to convert the electrical signal into a purely chemical one. The small areas that sit opposed to one another across this synaptic gap are densely populated with receptors and specialized channels. The arrival of the action potential at a synapse induces these channels to release organic molecules called neurotransmitters into the synapse space. These are capable of diffusing across the space and reaching the receptors on the other side. The side that releases neurotransmitters is typically referred to as the pre-synaptic neuron, and the receptor side is the post-synaptic neuron. This describes a spatial relationship between the synaptic components that also implies a directionality to the overall flow of information from neuron to neuron (Turbes & Schneider,

1989). The presence of these channels and receptors as well as the terminology referring to the pre- and post-synaptic neurons will be relevant in later chapters. There are several kinds of neurotransmitter molecules, each associated with different types of neurons and signals. Often, a receptor will set off a new action potential in the post-synaptic neuron, allowing the signal to flow continuously from one neuron to the next. These are called excitatory receptors and their adjacent synapses are called excitatory synapses. Some neurotransmitter/receptor combinations even function to decrease the action potential in the post-synaptic neuron, reducing the chance that it will transmit a signal. These are inhibitory synapses. Though they are not directly comparable, it is easy to picture neurons arranged in a similar way to common computer hardware logic gates in order to control the flow of signals through the nervous system.

Despite wide-ranging research in physiology and neurology, the subcellular distribution and density of neuronal components, particularly synapses, is still poorly understood. It is not known for certain what correlations lie between the locations and density of synapses and the structure of the associated neurons. The exact structural relationships between synapses and neurons is likely very important to the function of the nervous system, and a great deal can be learned from analyzing synapse distribution (Defilipe, Alonso-Nanclares, & Arrellano, 2002; Menon et al., 2013).

Chapter 3

Imaging and Microscopy

Microscopy is a broad term for using tools to view objects normally smaller than the resolution of the eye. Conventional light field microscopy relies on an external or ambient light source to illuminate the specimen and generate the image. This technique is fast, easy to perform, and is still in common use in almost every biological lab. Most scientists' first exposure to microscopic imaging will be in the form of a traditional wide-field light microscope. The two other aspects of microscopy that are very important in the context of this thesis are fluorescence staining and confocal microscopy. The key feature of fluorescence imaging is the use of fluorescent particles to generate the light used to illuminate the sample as opposed to relying on an external light source. Confocal microscopy is a technique for increasing the optical resolution and lowering the impact of noise in the form of ambient light while allowing for precision focusing and resolution that is superior to general wide-field microscopy. Most microscope technologies are not mutually exclusive in practical application, and almost any way to combine them has some value to the modern scientist. It is important to understand these techniques and their associated challenges in order to see the true application of these technologies in the study of neuronal structure.

Light

It is also important to have a basic understanding of light physics in order to simplify the following discussion of imaging technology. Light is a form of electromagnetic radiation that can be described as both a wave and a particle, but the wave nature of light is more relevant in this context. The most readily observable type of electromagnetic radiation is the visible light that human eyes use to see. Light waves all travel at the same speed in a given medium, but they can appear in different colors, including those above and below the range of human vision, based on their wavelength. Shorter wavelengths confer higher energy, with violet being the shortest wavelength of visible light and red being the longest. The ability to discern between different wavelengths of light and control the optical properties of light waves is key to modern microscopy. It should be noted that light waves can overlap and interfere with one another just like other types of mechanical waves.

Confocal Imaging

One key piece of imaging technology is laser confocal microscopy. One of the major drawbacks of conventional microscopes is that they can only resolve objects down to a minimum size. This size is related to the wavelength of light used to perform the imaging and the numerical aperture of the lenses used. Numerical aperture is a rating of how well a lens captures light from a sample, and higher rated lenses are more effective for imaging. Shorter wavelengths of light yield higher resolution. Even with the best lenses and optimal lighting, the

lower limit of resolution can still be larger than the structure being imaged. The main reason resolution degrades at this limit is due to the wave properties of light. When many individual beams of light are passed through a sample and focused into the lenses of a microscope, they interfere with one another causing overlapping patterns called Airy disks around objects in the image. Near maximum resolution, these interference patterns result in near total degradation of the image.

Confocal microscopy is a technique for reducing the impact of this kind of interference by using a very narrow and highly consistent light beam focused by a pinhole aperture to image the sample. Instead of generating a complete image of the sample by bathing it in light, the precisely focused beam of light is scanned across the image, one point at a time, and the result is a composite image of much higher resolution than could be achieved with wide-field microscopy. The precision focus of a confocal microscope also allows it to image at varying depth of a sample allowing for the construction of 3D images (Webb, 1996). It should be noted that the Z-directional scanning in 3D confocal images is based on the mechanical movement of the microscope stage, not the laser apparatus. Due to this, there is a discrepancy in Z-directional resolution compared to the X and Y directions. The Z direction is of much lower resolution comparatively, and considerations must be made with this in mind when analyzing these images.

Fluorescence

Fluorescence is a property of certain chemicals to emit light after absorbing light or other electromagnetic radiation. Unlike simply reflecting light, where light bounces off of the object, fluorescent particles absorb the excitatory photon and re-emit a new photon of a different wavelength. This property can be exploited to great effect in microscopy by applying fluorescent stains to samples and then exposing them to a light source of the correct wavelength to excite the fluorescent stain. Fluorescent stains may even absorb light above the visible spectrum such as ultra-violet while emitting light in the visible spectrum, allowing the florescent material to glow without the interference of another visible light source.

Many fluorescent materials do not glow indefinitely while exposed to their excitation wavelength. This is due to an effect known as photobleaching. While exposed to the light source, the absorption and emission of light can chemically damage the fluorescent particles, temporarily or permanently lowering their light emitting potential (Ghauharali & Brakenhoff, 2000). When used to illuminate microscope samples, this implies that care has to be taken not to over-expose samples and ruin a potential imaging opportunity.

In order to target fluorescent stains so that they will attach to desired structures in the sample, a technique known as immunofluorescence is used. Many fluorescent materials are biological compounds and can be chemically bonded to immune system proteins. The immune system uses a vast array of chemical antibodies that recognize specific chemical antigens. The relationship of these two molecules is much like two puzzle pieces fitting together. It is possible to bond a fluorescent particle or fluorophore to an antibody that targets a specific antigen in

order to illuminate features in a sample (Maity, Sheff, & Fisher, 2013). For example, if it was desirable to view the DNA in a sample cell, a fluorophore attached to an antibody that only attaches to DNA could be used. Either by filtering out the wavelength of light used to illuminate the fluorophore or using an ultraviolet source, a very clean image of the DNA could be resolved.

Immunostaining is subject to a key problem that can be difficult to work around. Antibodies are usually specific to their antigen, but it is always possible and even inevitable that some of the immunostain ends up in undesirable regions of the sample. This is due to clumping of the stain, failure to bond to the antigen, or any number of other biological factors. Immunostain that is not in the desired position can be the source of lots of noise when imaging.

Laser Scanning Confocal Images of Neurons

These techniques of confocal imaging and immunostaining are well suited to being combined. Laser confocal fluorescence microscopy is widely used in many scientific fields. By performing immunostaining and then imaging the sample with a confocal microscope with a wavelength that matches the fluorophore used, very high-resolution three-dimensional images can be created. The narrow beam of the confocal microscope and the act of scanning it across the sample minimizes the possibility of photobleaching and degrading the fluorophore. Photobleaching can still occur if the sample is scanned too rapidly or too slowly, and proper imaging takes a great deal of skill and good judgment from the microscope operator. Neurons

can be stained with immunofluorescent techniques, and their complex three-dimensional geometry makes them a good fit for laser scanning confocal imaging.

Automatic Detection

Though it is now possible to image neurons using modern technology in increasingly higher quality and quantity, the technology to analyze and interpret them has lagged behind. Great advancements have been made in morphological tracing of the neuron structure, including some here at the NIU Image Learning and Analytics Lab, but automatic synapse analysis in images of complex neurons is not as well explored (Myatt, Hadlington, Ascoli, & Nasuto, 2012; Xiao & Peng, 2013). Current biological research about the subcellular distribution of synapses is mostly limited to painstaking manual analysis of small subsections of the neuron. Additionally, operating in a three-dimensional space is very difficult for a human to do and can introduce a large degree of subjectivity into any manual analysis (Zhou & Peng, 2011). In many other fields of biological study, especially in genetic data and image analysis, semi- and fully automatic analysis techniques are very desirable (Libbrecht & Noble, 2015). Reducing the impact of subjectivity and human error is a compelling reason to explore automatic analysis.

Chapter 4

Terminology

Digital Images

The neuronal images analyzed in this thesis are three-dimensional multi-channel digital images produced by the previously discussed microscope techniques, namely immunofluorescence and laser confocal microscopy. Even though modern digital images are very sophisticated and are of high quality, there are certain important technical considerations that need to be taken into account when analyzing them. Unlike their real-world counterparts, digital images are abstractions that only partially represent an object. They cannot fully reproduce every aspect of the object, even in three-dimensional images. If you imagine a real object sitting in a Cartesian coordinate system, there are an infinite number of points that could be used to represent its volume and an infinite number of colors that could be used to describe it. A digital storage medium has only a finite amount of space to store information about an image, so it must be simplified and abstracted in order to create a digital representation of a real object or image.

Fundamental to the nature of digital images is the process of choosing a finite way of representing the object. In computer graphics, the unit chosen to represent a discrete point in a digital image is called a pixel in two dimensions and a voxel in three dimensions.

In this research, since three-dimensional images are the primary concern, the term *voxel* will be used in most cases. Imaging algorithms must perform a process called sampling where a finite set of pixels or voxels are used to represent the whole image. This necessitates the sacrifice of some details of the object but allows for the image to be represented in a format that can be interpreted by a computer and operated on mathematically. Each sampled point in the object corresponds to an exact location in a Cartesian coordinate space, and the image is essentially converted into a matrix. When an image is represented this way, each pixel or voxel may have an associated set of values that describe how that region appeared in the original object. These values correspond to the value of the matrix element at the x and y locations in the image matrix.

The simplest type of digital image is a binary image. In this case, each voxel is represented by its x, y, and z coordinates and a corresponding bit value indicating if the voxel is part of the object or not. A binary image is therefore said to have a bit depth of one. If a larger value is used to describe the voxel, usually 0-255, the image can now represent intensity as a monochromatic gradient. Such images would have a bit depth of eight, though higher bit depths can also be used to represent a finer gradient. Any color can be broken up into red, green, and blue intensity values, a color model that is a good fit for digital images due to the display specifications of computer screens. Based on this assertion, a color image can be constructed by storing three separate values for each voxel representing the intensity of its red, green, and blue color components.

The images used in this research are in the form of three value voxels, each corresponding to an associated color channel, with a slight but useful modification. The signal wavelength emitted by fluorophores in a confocal fluorescent microscope image is rarely a pure red, green, or blue wavelength. Instead of recording the exact color breakdown of the observed signal as a full-color image, the signal wavelength can be interpreted in a gray-scale based only on intensity. Each type of fluorescent stain can then be arbitrarily assigned a color channel and a false color image can be constructed from three different fluorescent signals. In the neuron images, one stain is usually used to image the neuron morphology, leaving other channels available to image the synapses. This is very beneficial for analysis efforts because each channel discretely represents information about a particular facet of the neuron. It is possible to consider one channel at a time when processing these images, and that greatly simplifies the task of separating what is a synapse from what is neuron morphology.

Segmentation

Separating synapses from the rest of a neuron image is an example of an image segmentation task. Segmentation refers to determining which voxels in the image are in the foreground (objects of interest) and which are in the background. The end result of image segmentation is usually a binary image describing the foreground and background with the bit values for each voxel. There is no single guiding algorithm by which segmentation is accomplished. Instead, many different approaches can be used to accomplish the desired

result. Segmentation can be as simple as thresholding and image processing based on a particular value or may involve many complex steps. A threshold refers to a value used to discretize the voxels in an image into binary values. It works by simple comparison to the value of a voxel. If the value is higher than the threshold, the corresponding voxel in the binarized segmentation result is conventionally set to the value 1. If it is below the threshold value, it is set to 0 instead. Furthermore, the threshold for an image need not be a global value. It may be the case that image intensity varies from region to region inside the image, and a single global value is not capable of representing all of the foreground. Instead, there can be many thresholds corresponding to different regions in the image. In this case, it is called local thresholding, though the end result is still a binarized image. In order to determine a threshold for an image, it is especially useful to construct an image histogram representing the values of all voxels in the image. An image histogram plots the values of all voxels in the image by frequency. This allows for the determination of threshold values by simple analysis of the histogram and is very important in many segmentation approaches. In the case of identifying synapses inside of three-dimensional neuron images, simple segmentation procedures, even local thresholding algorithms, are insufficient to completely determine what is a synapse and what is not a synapse. Indeed, the exact identification of synapses in an image is a contentious process even within the neuroscience community (Burette, Collman, Micheva, Smith, & Weinberg, 2015).

Robust Automatic Threshold Selection

There is one very specialized type of binary segmentation involved in the synapse detection process that necessitates more comprehensive explanation. This is robust automatic threshold selection, or RATS. The RATS algorithm is a local thresholding approach that attempts to limit the impact of contrast variance among regions in the image and instead calculates many local thresholds specific to sub-regions in the image (Wilkinson, 1998). RATS functions by combining two key steps in order to binarize the whole image. The first is the use of a recursive sectioning of the original image into smaller regions by use of a quad-tree. This works by dividing the image into four equal-sized regions and then dividing each of those regions into four more regions until they are a minimum size. Then each of the sub-regions is subjected to analysis of a discrete differentiation operator called a Sobel operator to approximate the image gradient around each pixel in each sub-region recursively all the way to the leaf regions of a specified minimum size. Then RATS uses the gradient -sum of each pixel in the sub-region to calculate the local threshold for that region. In order to limit the effects of noise and reduce stark variances between sub-regions, if a sub-region's calculated threshold is below a specified noise value, it inherits the threshold of the parent region in the tree structure. After each sub-region has been calculated, the final thresholds are interpolated for the whole image and used to assign the final binary value for the segmentation.

Machine Learning and Computer Vision

Separating the important components from images is a vast field of research in its own right and is often related to the field of computer vision. Human visual systems perform this kind of analysis extremely rapidly and in highly diverse ways. A person is capable of extracting relevant objects from a scene in almost any conditions where their vision is otherwise not impaired. In order to do this, the visual system uses an assemblage of different types of distinguishing features to identify objects. It is a great challenge to enable a computer system to be able to interpret the matrix representations of images in a meaningful way. This expansive concept delves into many specialized fields and is the subject of a great deal of interest in modern computer science. The work in this thesis was accomplished using relatively well-established computer vision and machine learning concepts that were applied in a new way and in combination with other technologies in such a way as to be unique. It is necessary to have a general understanding of the concepts of machine learning to appreciate their application in this context.

Machine learning is used most often to solve a classification problem, where it is desirable to sort data into several different classes by making a decision based on the attributes of the data. The key difference between a purely statistical classification and a machine learning algorithm is the capacity for the machine learning algorithm to learn from a training set of data and develop an internal model for prediction. The general flow for machine learning and classification involves several steps shown in Figure 1. First, the classification problem

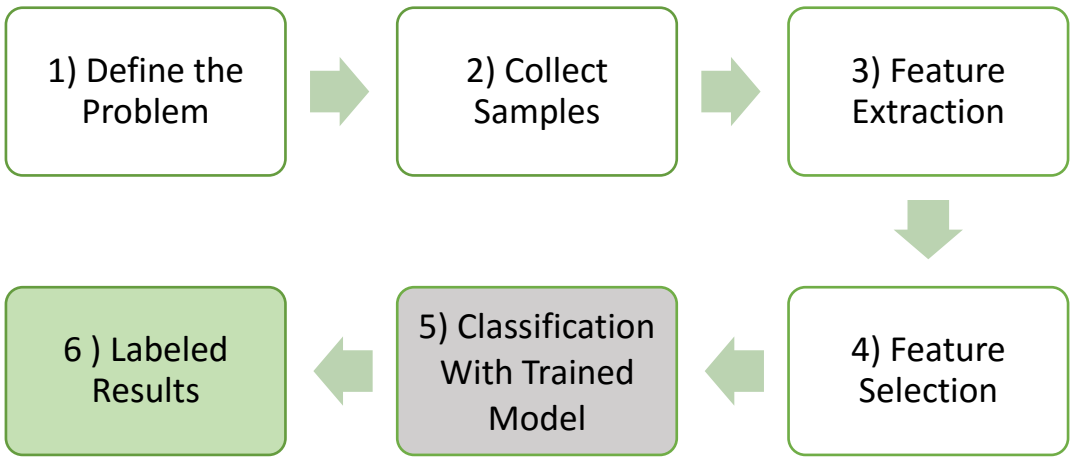

Figure 1: Typical machine learning based classification flow.

must be defined in appropriate terms (1). Then raw samples are collected of the items to be classified (2). This includes sound information, pixel values, recorded values from sensors, or any other form of digitally representable data. Next, the samples are refined by feature extraction (3), though it is possible to directly operate on the raw sample data. Any valuable descriptor of a sample can be a feature, and there exists a wide range of feature extraction algorithms that can calculate or choose valuable features from a sample. Feature extraction (3) is a way to highlight more important aspects of the sample data. In supervised machine learning applications, some of the samples are manually labeled according to their class, and

these samples are used to train the machine learning algorithm. Different machine learning algorithms have different methods for training themselves, but the overall goal is to produce a model that has minimal error when classifying the training data. Once the model has been trained and features are extracted, the set of features can again be refined by statistical feature selection (4). Finally the samples are classified (5) by the machine learning model which assigns them a label (6).

Haar Wavelet Feature Extractor

The feature extractor chosen for use in this thesis is the Haar wavelet transform. This extractor applies the discrete Haar wavelet matrix to the data passed to it, producing an output of features equal in dimension to the original data. This transform has the effect of highlighting sharp contrast differences in the horizontal, vertical, and diagonal directions. This transformation is scalable to any number of dimensions. However, full extension of these features from 2D to 3D can lead to a big increase in the number of features that is cubic to the side length of the region. So we used an anisotropic extension that performs full transforms along the x-y direction and then obtains the weighted combination of the coefficients along the z direction (Zhou, Lamichhane, Sterne, Ye, & Peng, 2013). Such features can be particularly suitable for 3D confocal microscopic images since the transformation adapts to the anisotropic nature of confocal imaging where z resolution is typically less than x-y resolution.

Support Vector Machine Classifier

A support vector machine is a type of supervised learning model that uses training points to construct a feature space and separate that space into two areas. Once this space has been divided, testing data can be mapped into this space, and the class of each point is predicted by where it falls relative to the dividing plane of the SVM. The SVM works by optimizing the maximum separation of the training points such that the separating plane is as far as possible from the nearest training sample of each of the classes (Chapelle, Haffner, & Vapnik, 1999). For cases where the training classes are not clearly separable by a plane, the feature space can be transformed over a kernel space using a kernel function to make the problem separable.

BIOCAT – Modular Bio-Image Analysis

BIOCAT stands for BIOlogical Image Classification and Analysis Tool. It is a bio-image analysis and modular machine learning tool developed by the NIU Image Learning and Analytics Lab (ILAAL) to provide a GUI-driven interface for the application of feature extraction and machine learning algorithms for biological image data (Zhou et al., 2013). Many ILAAL members have worked on BIOCAT over time, and it is a key resource both within the lab and for public use. The BIOCAT GUI was used to explore possible machine learning algorithms for use in synapse detection, and the BIOCAT API was used to invoke the machine learning algorithms within the actual synapse detection methods.

Chapter 5

Literature Review

While machine learning has been used in analysis of two-dimensional biological images, there are few examples of machine learning-guided synapse quantification, and even fewer in the three-dimensional setting. There are specific challenges associated with analyzing confocal 3D neuron images. They are anisotropic on the Z direction, meaning that the resolution is much lower between the Z axis and the X/Y axes due to the scanning mechanisms of the microscope. Immunostained images often have staining artifacts present in them, and they can vary in contrast across the image.

Most biological image analysis is performed on 2D images. Operating in 2D greatly simplifies the problem by directly reducing the computational complexity of analysis. Existing general methods for quantifying 2D biological images are largely not applicable to 3D images or are prohibitively inefficient in 3D (Carpenter, Kamentsky, & Eliceiri, 2012; Shamir, Delaney, Orlov, Eckley, & Goldberg, 2010; Wählby, Lindblad, Vondrus, Bengtsson, & Björkesten, 2002). The majority of these methods were developed for general image processing and result in low robustness when applied to complex 3D images. Existing object or threshold-based detection algorithms are either constrained to 2D analysis or are also insufficient to handle the complex

3D images (Kim et al., 2012; Meseke et al., 2009). In general, there are very few automated 3D biological image analysis tools available.

There have been advancements in 3D biological object detection for structures such as cells and nuclei (Bjornsson et al., 2008; Chen, Velliste, Weinstein, Jarvik, & Murphy, 2003; Zhou & Peng, 2011), but as far as we are aware, there are no prior automatic learning-guided algorithms specifically for the extraction of synaptic information from 3D confocal microscope images. There has recently been research into machine learning-guided analysis of electron microscope (EM) neuron images (Navlakha, Suhan, Barth, & Bar-Joseph, 2013). This study is the most directly comparable in terms of operating on 3D images and specifically targeting synapses, as well as the application of machine learning algorithms for synapse detection, but it is fundamentally different due to the use of EM images. These types of images are of much higher resolution but are more difficult to produce and may interfere with the natural arrangement of subcellular structures.

Chapter 6

Methods

Overview

Isolating synapses and extracting information about them from a three-dimensional neuron image is a complex task. The greatest barrier to manual annotation is the large size of these images and their three-dimensional geometry. It is difficult for the human eye to navigate three-dimensional representations, as we can only see a 2D view on a computer screen. The process of synapse detection and quantification involves several key steps and algorithms. The main goal of this thesis was to explore approaches to identifying synapses in these images using largely automatic means in order to reduce the impact of human subjectivity in the detection of synapses and to allow for much faster and more accurate processing of three-dimensional images compared to other techniques. To accomplish this, a general algorithm was implemented as a small set of plugins for the ImageJ platform (Abramoff, Magalhaes, & Ram, 2004) and standalone Java applications that were capable of detecting synapses in the test images when operated in a pipeline. The overall theme of the algorithm was to start with a large set of candidate voxels and reduce them at every step until only the synaptic centers remained.

The overall flow for synapse quantification is depicted below in Figure 2. The images to be processed were first split into their individual color channels corresponding to the cell morphology and synapse channels present in the image and preprocessed using several classical image enhancement techniques (1). After this, the core synapse detection algorithm was used to extract the candidates from the main synapse channel (2). In order to use a machine learning module in the detection, the correct model needed to be selected. BIOCAT was used to determine what model was to be used for synapse detection. This was accomplished by loading regions of interest representing positively identified synapses and negatively identified regions into BIOCAT's interactive model selection feature. This was used to compare different combinations of feature extractors and classifiers directly. Once the model was chosen based on this comparison, and trained on a set of data from the images to be analyzed, the learning-guided synapse detection algorithm was applied to the synaptic channel. The positively identified synaptic markers resulting from this step were then compared based on proximity to the structures in the other image channels. If there was another synaptic channel present, that comparison was completed (3) before comparison to the morphology (4). The intention of this proximity comparison was to reduce the impact of false positives reported due to aberrant stain in the synapse channel and increase the overall fidelity of the final synaptic count. Finally, the synapses were quantified based on their distribution in the dendritic arbor (5). The details of each of these steps will be described in more detail in the next subsections.

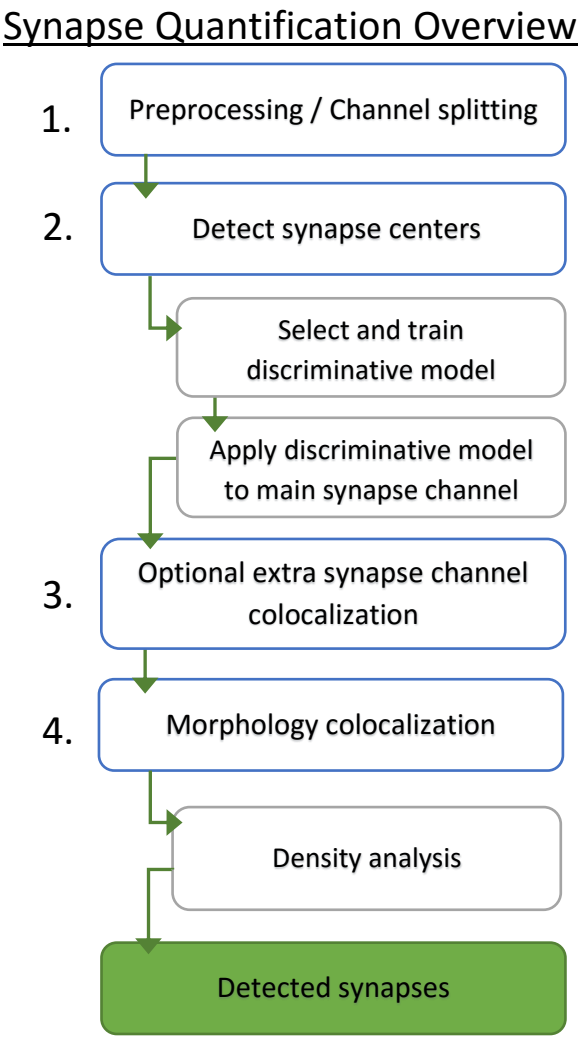

Figure 2: Main flow for synapse detection.

Preprocessing

Before being analyzed using the learning-guided algorithm, the original multi-channel image had to be split and some preprocessing performed. Splitting the image was accomplished using the ImageJ platform to separate the color channels of each image into independent gray-scale images. Figure 3 below shows an example of an image split into component color channels. From here each channel, synapse and morphology, was processed independently until they are recombined to formulate the final synapse detection result. All channels were subject to a rolling-ball background subtraction of radius 5 in order to reduce low-intensity background noise. Rolling-ball subtraction is the process of correcting for uneven background intensity in an image by averaging the intensity of a region under a scanning ball and subtracting the averaged value from the image. By passing the ball over the complete image, areas of low intensity are reduced to a more even intensity while foreground objects are largely unaffected. On images where intracellular noise was a problem, the morphology channel was subtracted from the post-synaptic marker channel in a pixel-to-pixel operation in an effort to remove the noise. The effects of this preprocessing are demonstrated in Figure 4.

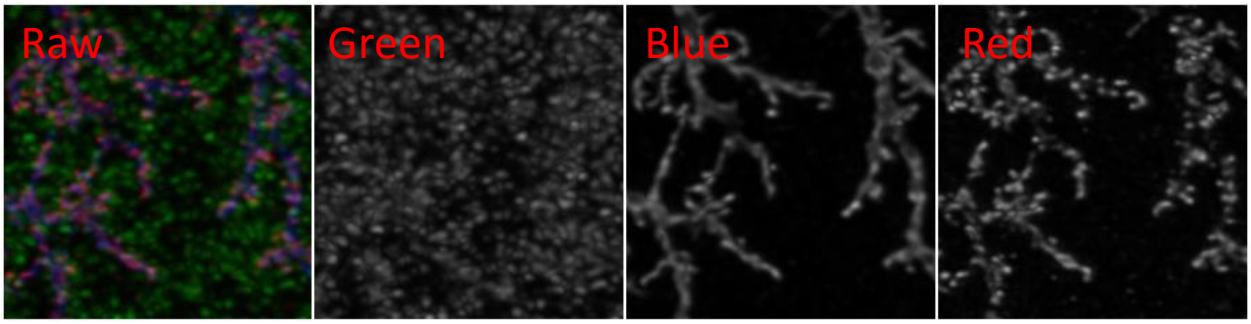

Figure 3: Splitting a neuron image into individual color channels. In this image, the green channel is pre-synaptic markers, the blue is morphology, and the red is the post-synaptic marker.

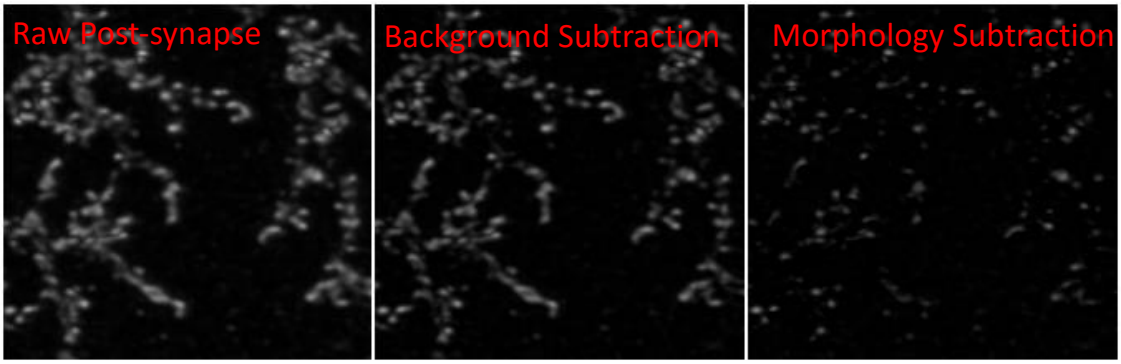

Figure 4: Processing the synapse channel to remove noise.

Model Selection and Training

Many machine learning models were considered for use for synaptic detection. Model selection was accomplished using the BIOCAT user interface and its comparison functionality. First, training data had to be selected for each family of neuron images so that the machine learning models could be constructed. The machine learning problem was formulated as a voxel neighborhood-based binary predictor. That is, for each voxel to be classified, the region around it should be considered and used to come to a decision about that voxel. For each neuron image, sets of positive and negative 9x9x3 voxel regions of interest (ROIs) were selected randomly, representing synapse centers and areas that were not synapse centers, respectively. Examples of these types of training ROIs can be seen below in Figure 5. These regions were loaded into BIOCAT using its "model selection and training" mode. BIOCAT allows for the combination of feature extractor algorithms and feature classifiers into algorithm chains that can then be run, saved, and compared directly. Candidate chains were created from commonly useful modules and were tested on the training data sets.

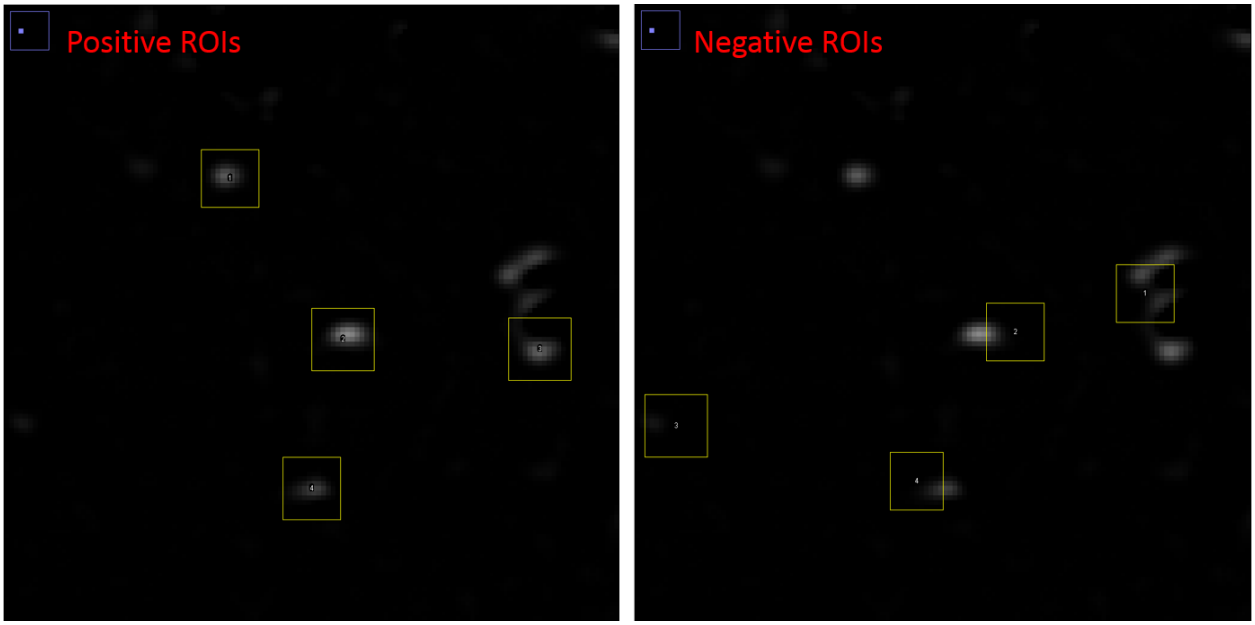

Figure 5: Example positive and negative ROIs used for training the discriminative model.

Synapse Marker Detection

Once the image channels had been preprocessed the main synaptic marker channel was passed to the core machine learning-guided detection algorithm. In the case of the pre- and post-synaptic images, this was the post-synaptic marker channel. This algorithm is at the heart of this thesis and does the majority of the important work in detecting synapses in the images. This was implemented by me as an ImageJ plugin that takes the synaptic image channel, training ROI data for the machine learning model, a BIOCAT algorithm chain, and various

thresholding and file I/O parameters. Once initiated, the machine learning model specified in the BIOCAT chain file was trained on the training ROI data. Once this was complete, a binary image copy of the synapse channel was constructed using code adapted from the open-source ImageJ RATS algorithm with generous thresholding parameters on every slice of the 3D image. When invoking RATS thresholding on an image, three parameters were required: the estimated global intensity of the noise in the image, a scaling factor, and a minimum pixel size for the quad-tree division of the image. For the neuron images, a noise estimate value of 3 or 4 was used, a scaling factor of 3 was set according to the recommendation of the original plugin author, and the minimum leaf size was automatically calculated to provide a five-level deep quad-tree. The purpose of this mask is to avoid analyzing voxels that are clearly backgr-und, greatly improving the overall performance of this stage of processing. A sample section of this mask is depicted in Figure 6. The synapse voxels make up a relatively small percent of the large 3D volume. Applying the machine learning model to all of these many millions of obviously background areas is inefficient. Instead, the algorithm considers only voxels that correspond to the foreground of the binary mask created by RATS.

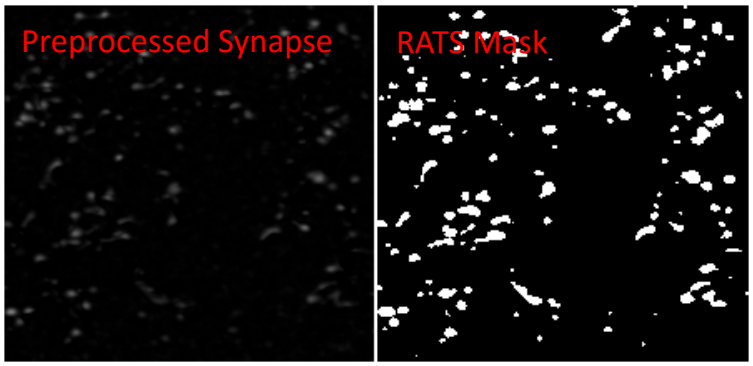

Figure 6: RATS used to supervise machine learning analysis. If the RATS mask is lenient, it serves to exclude the majority of background voxels from the machine learning classification.

For every voxel corresponding to the foreground of the mask, a region around the voxel that roughly corresponded with the average size of a synapse was extracted and passed to the trained machine learning model. The model performed feature extraction and classification according to the BIOCAT algorithms supplied in the chain file and then classified each voxel as positive or negative. Positive voxels are taken to be near a synapse center, and these were collected and output from the synapse annotator ImageJ plugin.

The result of the plugin was a set of 3D voxel aggregates that were determined to be close to synaptic centers by the machine learning model. The center of these objects was then found by a connected component analysis and center of mass calculation using the 3D Object Counter ImageJ plugin (Bolte & Cordelières, 2006). This was done to reduce the small clusters of positively identified synapse center candidates to a single point per synapse. This reduction in candidate voxels is demonstrated in Figure 7.

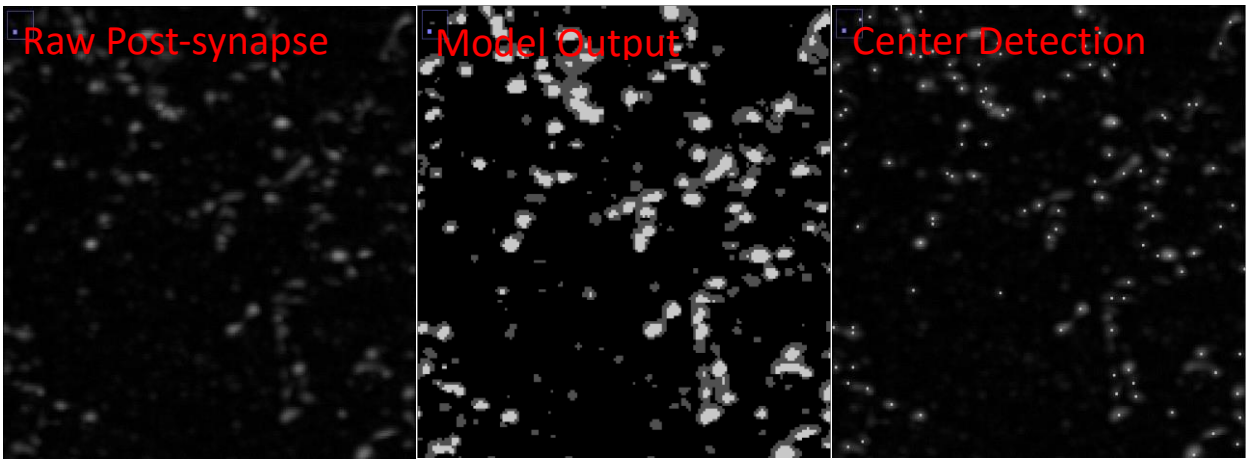

Figure 7: Application of the machine learning model and center detection.

Multi-Channel Comparison

Once the main synapse centers were been calculated, the set of candidate synapses was reduced by comparison to the other image channels to improve the accuracy of the detection. Staining artifacts in the post-synaptic channel and intracellular noise resulting from stain present inside the neuron can result in false-positive post-synaptic markers. By comparing the results of the synapse detection to the other channels, these false results can be screened out. Comparison to the other image channels was performed by proximity analysis between the previously detected synaptic markers and the structures present in the other image channels, as described below.

Pre-Synaptic Marker Comparison

The pre-synaptic channel, if present in an image, was not stained on only the target neuron but was instead stained on all neurons that fell within the imaging area. This resulted in a great number of pre-synaptic markers relative to post-synaptic, the bulk of which were not related to the target neuron at all. Rather than try to extract each of these pre-synaptic markers, it was decided to use them as a logical check for the already-detected post-synaptic markers. The pre-synaptic marker channel was preprocessed more aggressively using ImageJ tools to extract a binarized mask of the approximate locations of pre-synaptic marker stain.

The process for creating the binarized comparison mask of the pre-synaptic marker channel was composed of several traditional image processing steps. The pre-synaptic marker channel was subjected to a rolling-ball background subtraction of radius 5, despeckled, contrast enhanced by 3% globally, and then despeckled again to highlight the pre-synaptic markers while reducing the noise present in the image. After that, the partially processed image was segmented using a 3D watershed algorithm for ImageJ (Legland, Arganda-Carreras, & Schindelin, 2016) to isolate the approximate locations of the pre-synaptic markers.

The result of this preprocessing was a binarized image representing the locations of the pre-synaptic markers in the image. Using a custom ImageJ plugin developed by me, the results from the main synapse detection were compared to the mask by proximity. A synaptic marker was only output from this step if it was within a 3D space of 9x9x5 voxels of any pre-synaptic channel. This space is the approximate size of a synapse within these images. An example of the pre-synaptic marker channel preprocessing is shown in Figure 8.

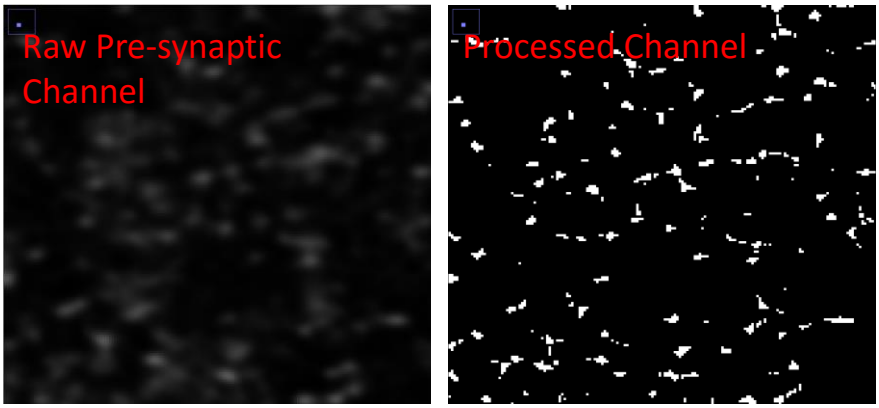

Figure 8: (left) An example of the raw pre-synaptic channel. (right) After background subtraction and 3D watershed processing of the channel.

Morphology Comparison

The morphology channel is present in all of the images. This channel is the result of staining target proteins on the cell membrane of the neuron. For synapse detection, this channel is used to exclude areas of synaptic stain that are not located near the neuron structure. These areas are considered to be artifacts from the image staining and preparation process in the lab and can't easily be eliminated fully before imaging occurs. Two different approaches were attempted for morphology colocalization, traditional preprocessing and comparison to a 3D neural reconstruction generated by another member of our lab. Due to the promising progress of our lab's neural reconstruction research, it was decided to incorporate this technique into synapse detection. The neural reconstruction comparison was used for the final reported results.

The 3D neural reconstructions were used to further reduce the set of candidate synapses by removing those that were not proximal to the dendrite morphology. This was accomplished using a Java-based tool designed by former ILAAL members who took a 3D neurite file (the morphology), a synaptic marker file, and image size parameters as input and outputted the markers that were sufficiently close to the neuron structure while omitting those that were not. The output of this step was considered to be the final synapse detection.

Density Analysis

At the same time as morphological colocalization, each synapse was tagged according to the diameter of the nearest neuron morphological segment. The Java tool used to perform the morphological colocalization was originally designed to perform the comparison of synapses to neurite diameter, but the accurate detection of the possible synapses by the methods proposed in this thesis was required before this could be done. Because this tool was also capable of screening out candidate synapses that are not close to dendrite morphology, the density analysis and colocalization steps were accomplished with this one tool. This tool took a candidate synapse file, a neural reconstruction file, and image size parameters as input. The neuron branches were binned into categories based on a size range, and the proximity of each candidate synapse was compared to the neuron morphology. The output of this tool was a file of synapses labeled based on the size of their associated neurites and distribution statistics including density per length of each size category and density per surface area of each size category. Little is known about the density distribution of synapses in complex neurons such as LPTC (Defilipe et al., 2002; Menon et al., 2013). This analysis, while still in its preliminary stages, is one of the first to explore the density distribution in these neurons.

Validation

There are no other official gold standard counts for the synapses in these images, nor is there very much available information about the subcellular distribution of the synapses in the

dendritic tree. In order to perform a quantitative analysis of the synapse detection methods, a set of sample small regions from the images to be tested was constructed. These regions comprise only a tiny fraction of the image, but they are representative of the whole while still being manageable for manual annotation. Manual annotation of the full-size images that contain many hundreds of synapses is prohibitively difficult, while smaller regions can be more carefully considered. Several of these regions were selected and manually annotated by members of the lab with input from our collaborators at Ye Lab, University of Michigan Life Sciences Institute. These test regions formed a gold standard set for comparison that was used to judge the effectiveness of synapse detection on the larger images. The development of these test regions alone is considered to be a large contribution to the synapse detection research efforts.

The results of the synapse detection algorithms were compared to the test sub-regions using a specialized program that compared the manually annotated synapses in each test region with the automatically detected synapses and calculated useful quantitative metrics describing the effectiveness of the detection algorithms for each test region. This tool reported the fraction, precision, recall, and F-measure for each region. Fraction was defined as the number of synapses detected within the bounds of the test region divided by the number of gold standard synapses manually annotated for that region. This metric provided a straightforward way of determining if the synapse detection algorithm was detecting an appropriate number of synapses. Precision was defined as the number of synapses detected by the algorithm that were proximal to a manually annotated synapse divided by the number of

detected synapses. This metric determined if the detection algorithms were returning false-positive results such as noise or staining artifacts. Recall was the number of manually annotated synapses that were proximal to a detected synapse divided by the number of manually annotated synapses. This represented the relevance of the detected synapses, providing information about how many manual synapses were missed by the detection methods. Finally, the F-measure was the harmonic mean of the precision and recall statistics. It provided a useful unified statistic to judge the overall effectiveness of the detection on each region.

Chapter 7

Results

Data Sets

The images used to test synapse quantification were partial and whole dendritic trees from *Drosophila melanogaster* lobula plate tangential cell (LPTC) horizontally sensitive neurons provided by the Bing Ye Lab at University of Michigan Life Sciences Institute, seen below in Figure 9 and Figure 10 respectively. Though both LPTC neurons, these images are not stained for the same structures and are not of the same exact neuron. These neurons possess complex structures and are related to the optical processing and flight control systems in *Drosophila m.* The process of creating these images is difficult and time intensive, as it can take upwards of 72 hours of work by lab scientists to properly prepare and image a neuron sample. The true number of synapses in the dendritic tree is unknown, and the method proposed in this thesis is one of the first automated methods for estimating the total amount. These images were captured via a laser scanning confocal microscope system and then reconstructed into 3D. The whole dendritic tree image was stitched together using Amira software (Stalling, Stalling, Westerhoff, & Hege, 2005) before it was sent to our lab. The majority of the visualizations in this section were created using the Vaa3D multi-dimensional data visualization and analysis tool

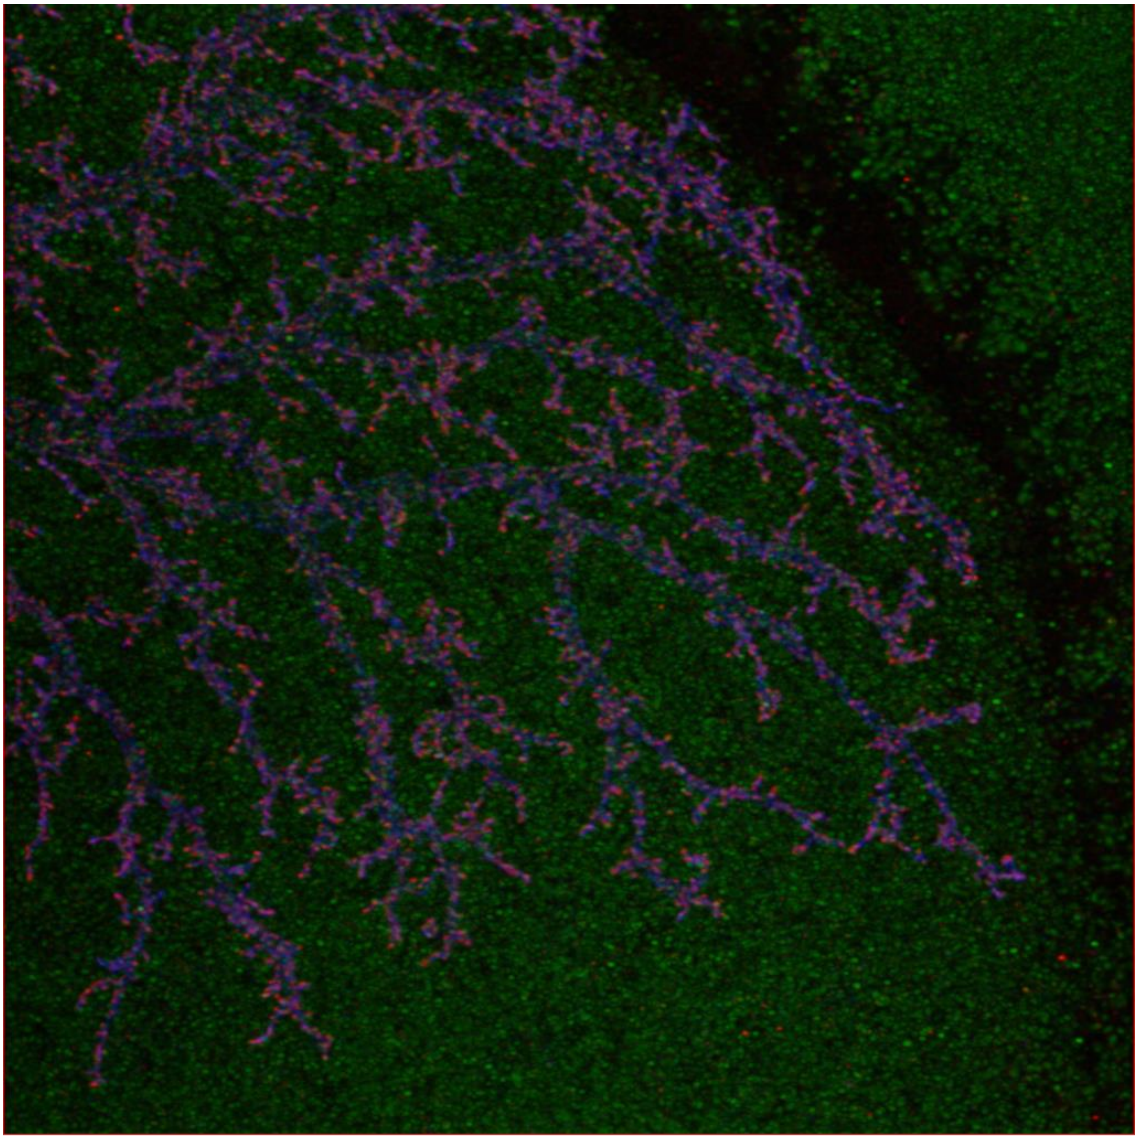

Figure 9: A maximal intensity projection of a partial dendritic tree image. The blue channel is the neuron morphology, red is post-synaptic markers, and green is the pre-synaptic markers.

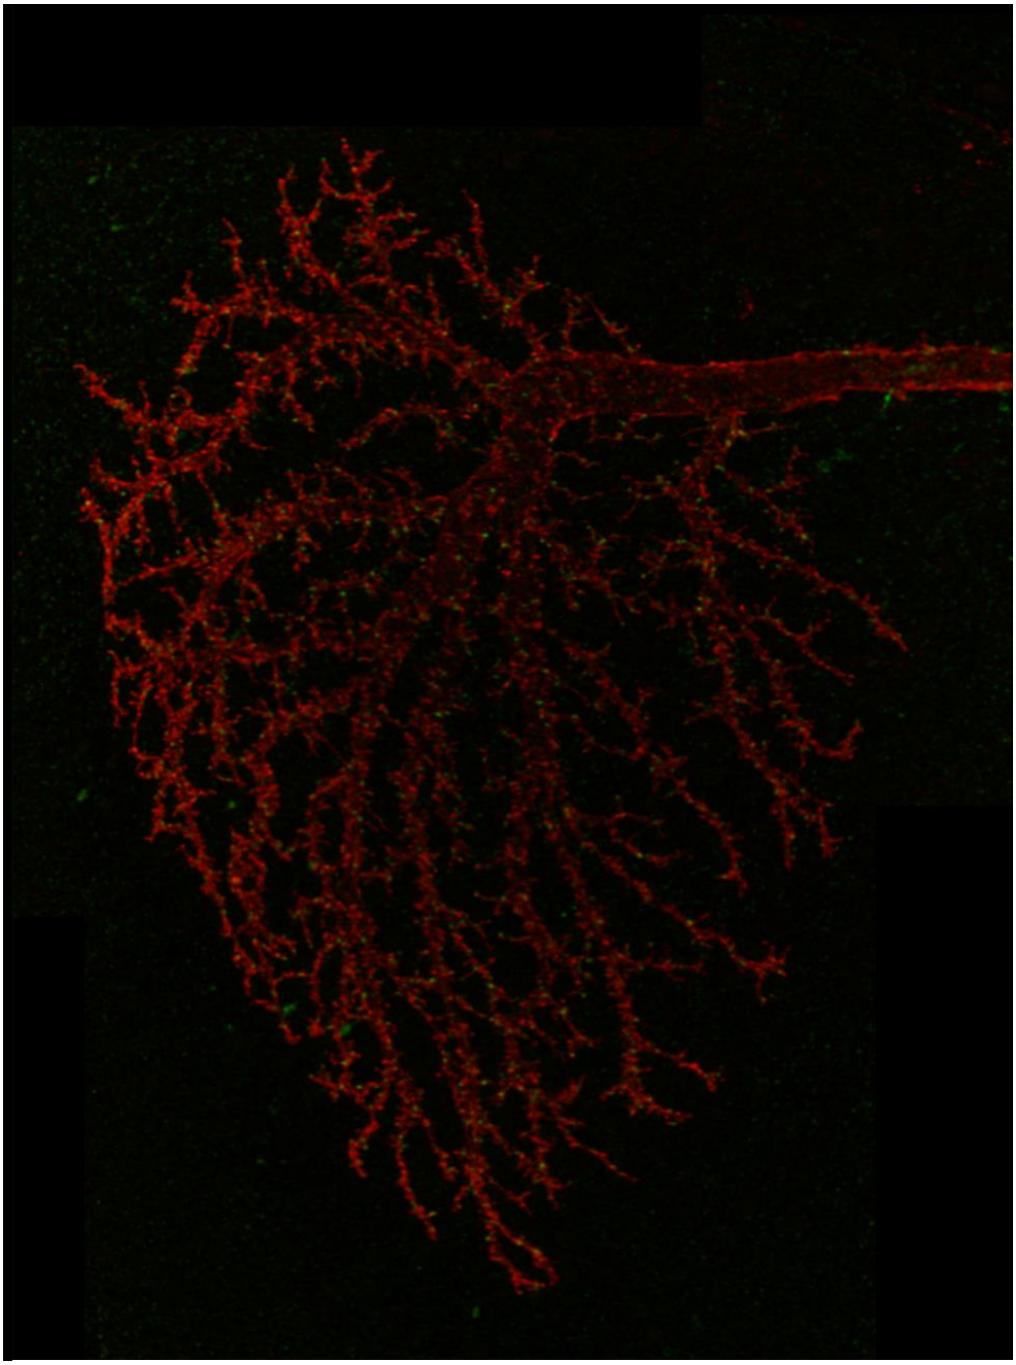

Figure 10: A maximal intensity projection of the complete dendritic tree image. The red channel is the morphology. The green channel is the possible synapses.

developed by the Peng Lab at the Allen Institute for Brain Science (Peng, Ruan, Long, Simpson, & Myers, 2010). This GUI-driven image visualization application allows for easier examination and manipulation of 3D images than a slice-based application such as ImageJ.

Model Selection

Before a full analysis of the images was performed, an appropriate machine learning model was chosen. The main interface of BioCAT was used to efficiently compare algorithm chains consisting of feature extractors and classifiers and select the most appropriate chain for synapse detection. Figure 11 shows the comparison and testing of chains performed by BIOCAT. The model was chosen after training on a set of 20 positive and 20 negative ROIs of size 9x9x3 voxels from the Kibra 004 partial dendritic tree along with a 5-fold cross validation. The most powerful feature extractor was determined to be the 3D anisotropic wavelet, or Haar wavelet transform, for both image types tested. The most effective classifiers were the multi-layer perceptron (MLP) and the support vector machine (SVM). Combining the 3D Haar wavelet extractor with either of these classifiers yielded a recognition rate of 95-100% depending on the variance in cross validation. As can be seen in Figure 11, the Haar wavelet + MLP and the Haar wavelet + SVM are the highest performing overall chains. Due to the shorter training time and negligible accuracy differences between it and the MLP, the SVM was used for testing.

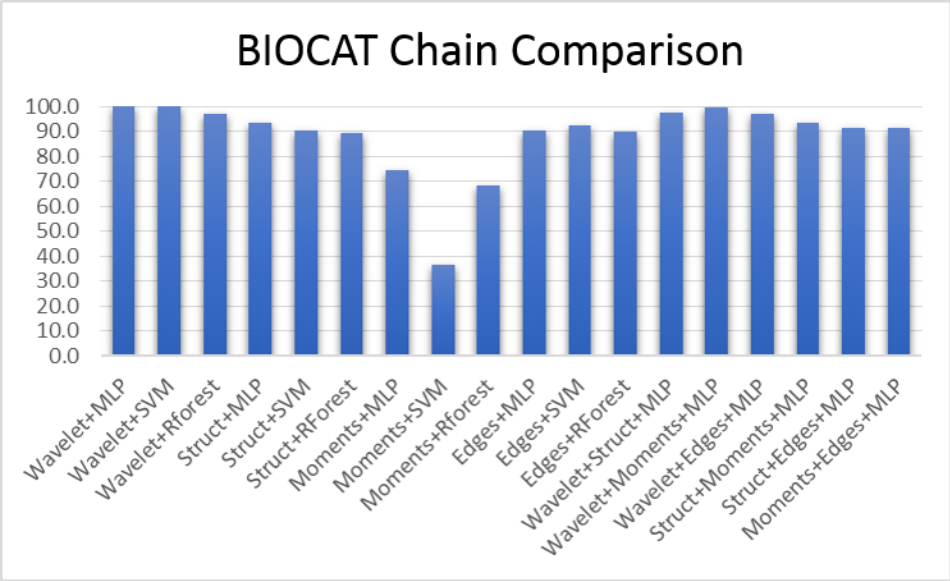

Figure 11: Comparison of 18 algorithm chains for 3D synapse detection tested in BIOCAT.

Synapse Quantification

Partial Dendritic Tree Images

There were four pre- and post-synapse partial dendritic tree images tested using these methods. These were all 1024x1024 images with between 128 and 155 slices in the z direction. The partial dendritic tree images were composed of three channels: cholinergic post-synaptic Dα-7 markers stained with green fluorescent protein (GFP), pre-synaptic nc82markers stained with DyLight649, and overall morphology stained with red fluorescent protein (cD8-RFP).

Together the three-channel image has a bit depth of 24, with each channel having a bit depth of 8. These images are a subset of a series of images representing an entire neuron, but the images designated as 004, 006, 008, and 010 are of the complex dendritic tree and were the most relevant for synapse detection. Despite not being a complete dendritic tree, these images are still very dense with synapses and are a good representation of the overall complex neuron structure in 3D. Each of these images was fully processed including pre-synaptic marker detection and morphology colocalization. The learning model was trained on ROIs selected in Image 004, and the same learning model was applied to the other images in the set. A small section of a final image can be seen in Figure 12. Some areas that appear to be false negatives are excluded by one of the processing steps. These regions could be false negatives, or they may be correctly removed based on a failure to meet the criterion of the detection methods. Visually, it can be very difficult to tell. In the Validation section below this is discussed in more detail.

The synapse detection algorithm performed well on these dense images, giving a plausible estimate for synapse count and locations within the dendritic tree sections. Areas of synaptic stain that were not near to the neuron morphology and pre-synaptic markers were successfully excluded. Even regions of high synapse density were usually successfully counted. Both very bright and very dim synapses are represented in the results.

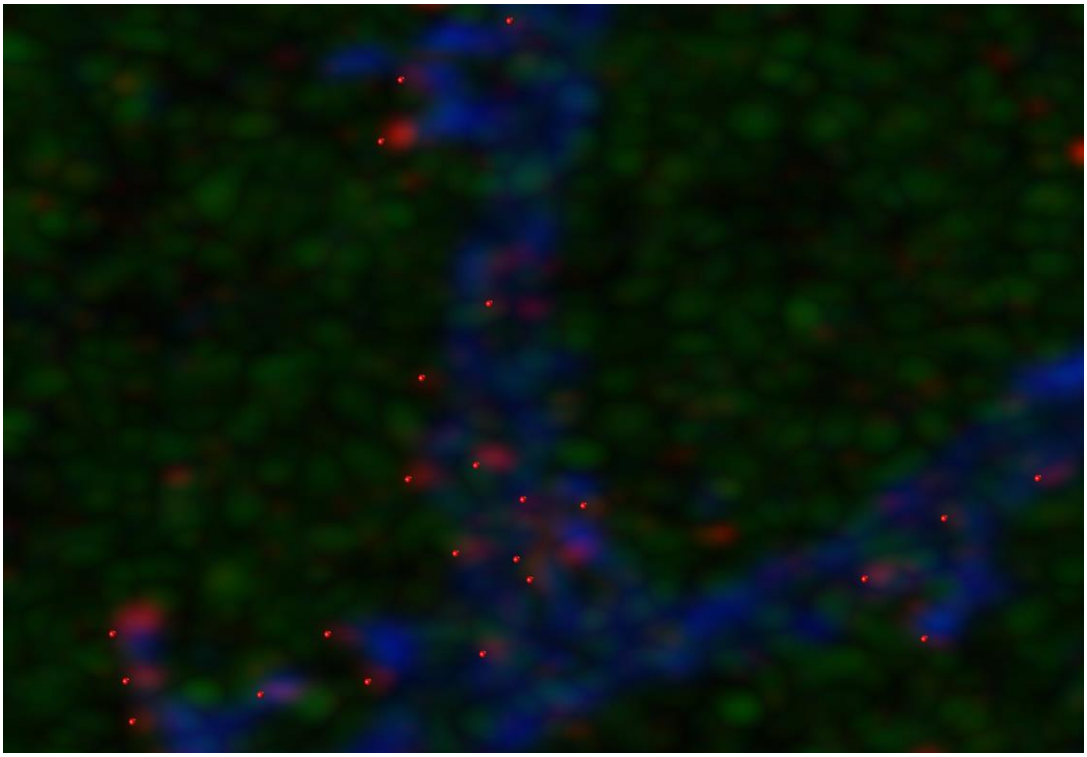

Figure 12: Final synapse detection of a small region of partial dendritic image 004. The red channel is the post-synaptic channel, green is pre-synaptic, and blue is the morphology.

For the four partial dendrite images, synapses were quantified using the succession of reductive steps described above, starting with machine learning-guided post-synaptic marker detection and then performing multi-channel comparison to reduce the candidate set of synapses until a final result was reached. Table 1 shows the detected synapses undergoing reduction at each major step for these four images.

Table 1 : Synapse detection in the partial dendritic trees using three channels

Image	ML guided	Pre-synapse	Morphology	Final Count
004	2488	-255	-190	2043
006	2100	-875	-225	1000
008	2997	-1443	-411	1143
010	1219	-680	-154	385

Whole Dendritic Tree Image

One complete dendritic tree image was processed using these methods. This image was actually a composite stitched together in the image providers lab (Ye Lab at University of Michigan) using the Amira data visualization software. The resulting image was 1296x2333x59 voxels in size. The whole dendritic tree image has two channels: the GABAergic post-synaptic RDL stained with Cy5 and the overall morphology again labeled by cD8-RFP.

The complete dendrite image was quantified using almost the same steps as above. There was no pre-synaptic channel for this image, so only morphological colocalization was performed. Table 2 shows the detected synapses at each step for the complete dendrite image.

Table 2: Synapse detection in the complete dendritic tree using two channels

Image	ML guided	Morphology	Final Count
whole dendrite	1328	-283	1045

Density Analysis

Density analysis of the synapses within the dendritic branch structure is an ongoing undertaking of the ILAAL, and the results of the methods in this thesis directly enabled work to proceed in this area. The determination of quantitative results for synapse density is still being developed, but the preliminary qualitative results are striking and worth presenting here. As mentioned above, the final step of morphological colocalization also groups synapses based on the diameter of the nearest synapse branch. In Figure 13, a visual of this grouping overlaid on top of the morphological reconstruction can be seen. In order to facilitate a qualitative assessment of the synapse distribution, the detected synapses are colored on a scale from blue to red. Synapses near larger branches are blue, while those near narrower branches shift increasingly towards red. Figure 14 represents the distribution of synapses in the complete dendritic tree image with the same coloration.

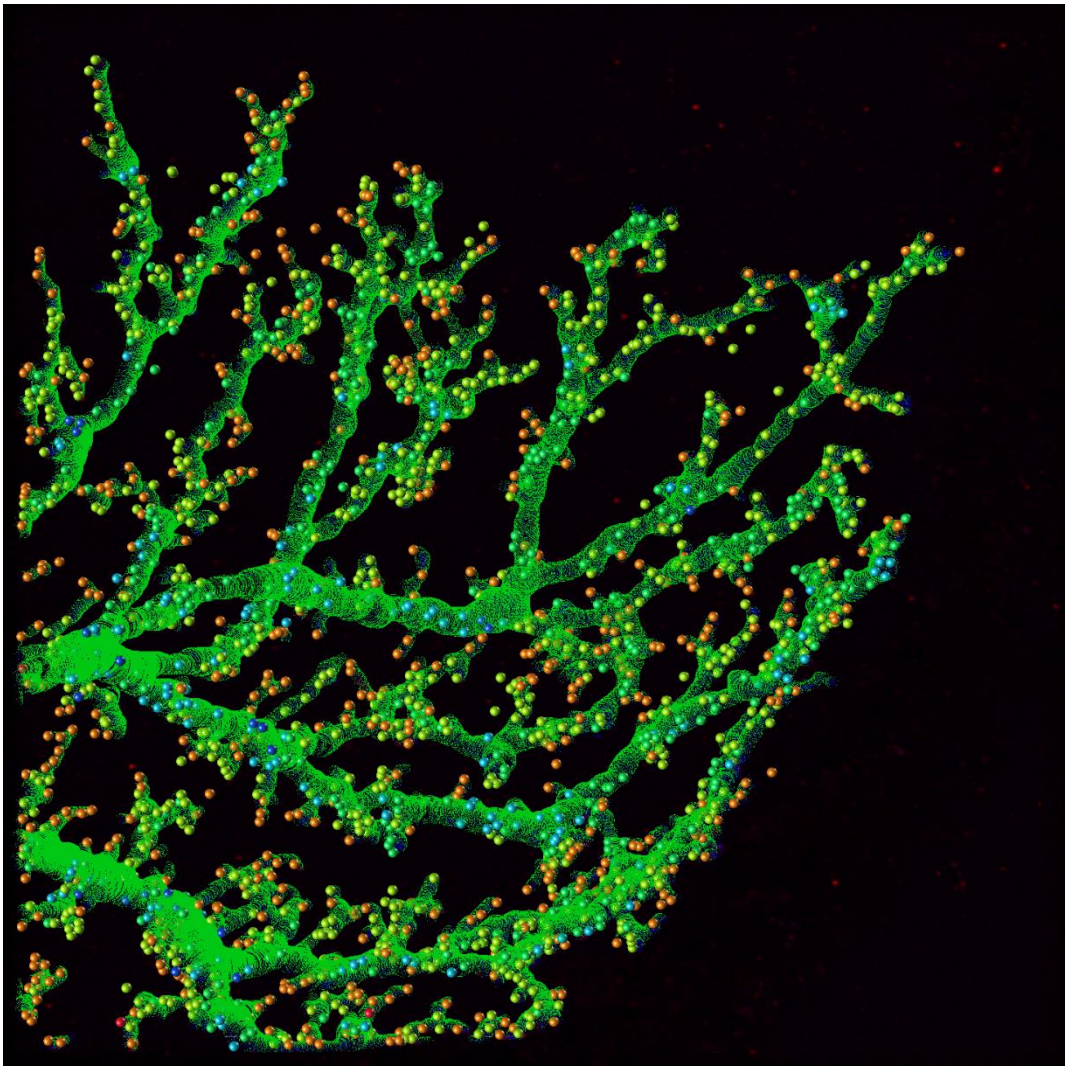

Figure 13: Partial dendritic image 004 with reconstruction and density color-coded synapses. Synapses near larger branches are blue, while those near narrower branches shift increasingly towards red.

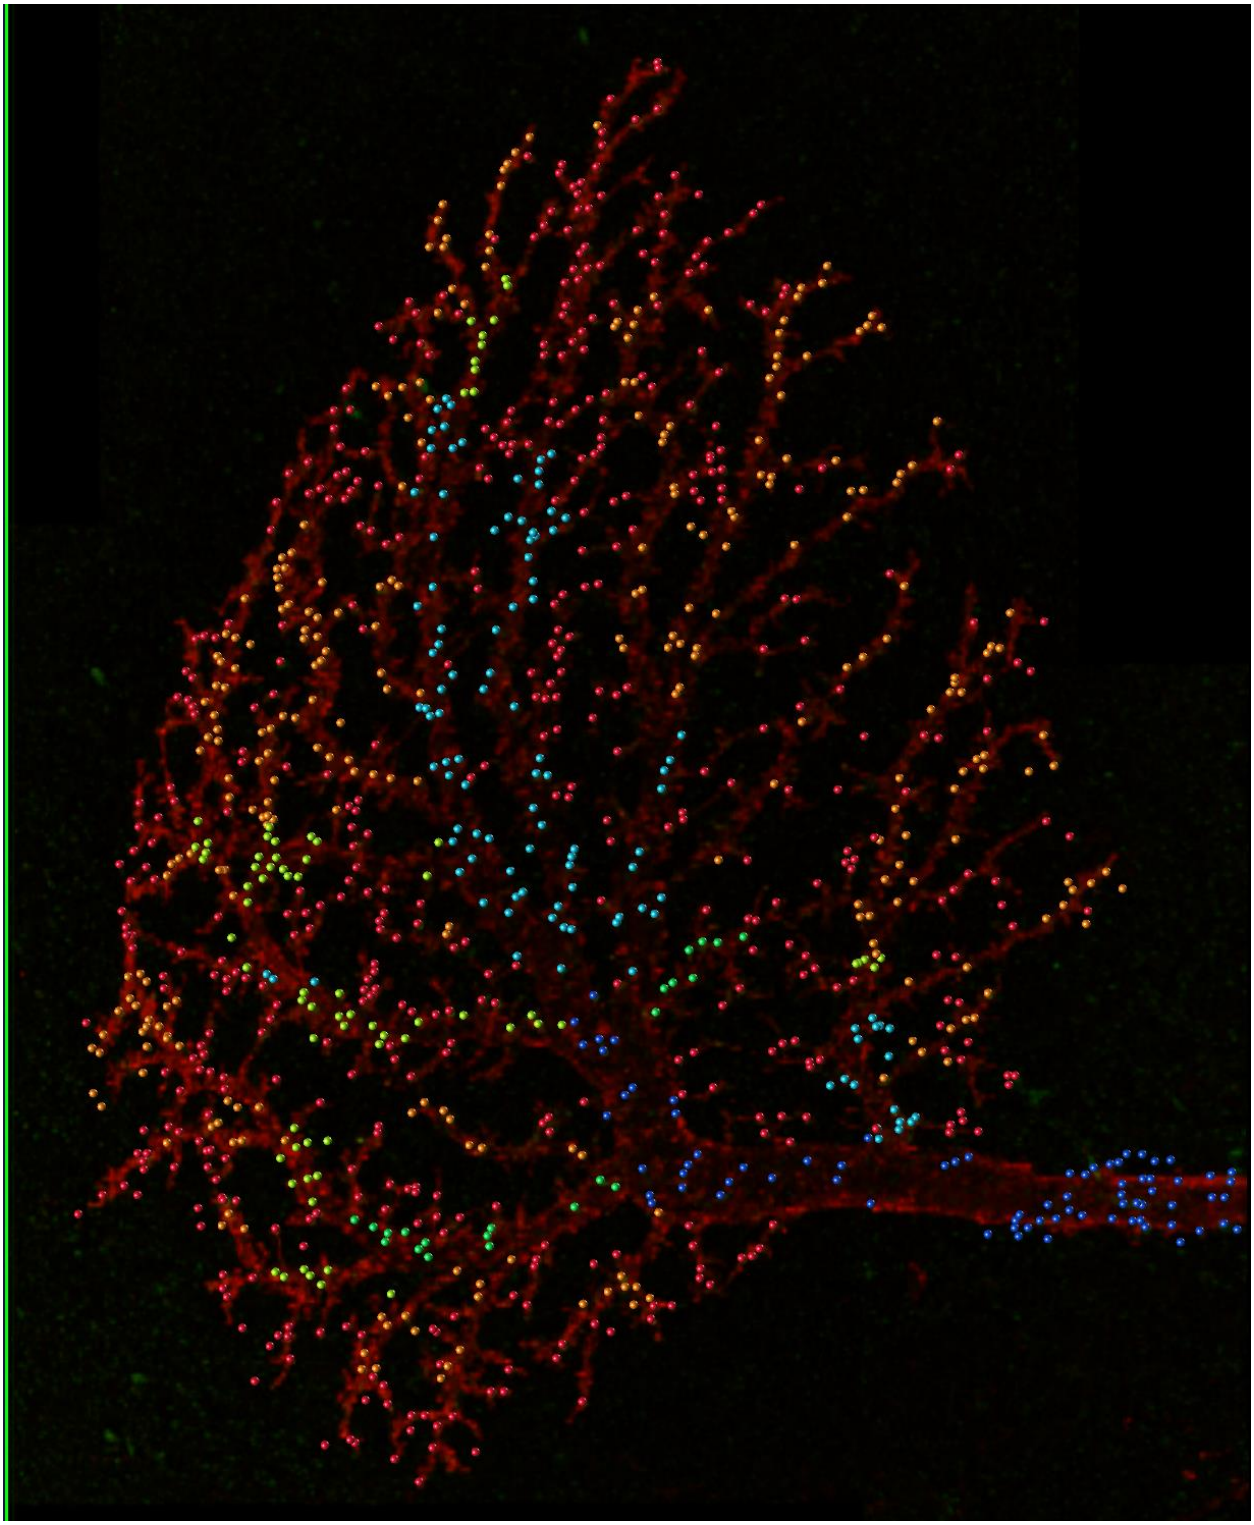

Figure 14: The complete dendritic image shown with only the color-coded synapses. Again, blue synapses are proximal to large branches, moving to red synapses at narrow branches.

Validation

Quantitative determination of the success of these methods was a significant challenge on its own. Partial dendritic image 004 was used to construct a comparison set and validate the effectiveness of the synapse detection. Ten sub-regions of size 150x150x10 voxels were randomly selected from the image, such that each region contained at least 20 synapses. These regions can be seen as a 2D projection below in Figure 15. These were manually annotated using the Vaa3D tool by three different lab members, and the final results were combined and agreed upon by each annotator. In order to provide a direct comparison with the synapse detection algorithms, the manual annotations were performed after the morphology channel subtraction described above in Methods. Without performing this preprocessing step first, the intracellular artifacts present in the image greatly obfuscate the synapses and render manual annotation prohibitively difficult. After this gold standard set was chosen and agreed upon, the specialized validation program was used to automatically calculate the fraction, precision, recall, and F-measure for the ten test regions (Table 3). It should be noted that while this set was used as the gold standard set in order to quantify the overall success of this algorithm, that doesn't mean that the manual annotation was perfect or that the synapses chosen for the set were absolutely correct. Even our neuroscientist collaborators are not entirely sure of the location of each synapse in these images. It is necessary to have a basis for comparison, however, and the combined opinion of several human annotators was deemed acceptable.

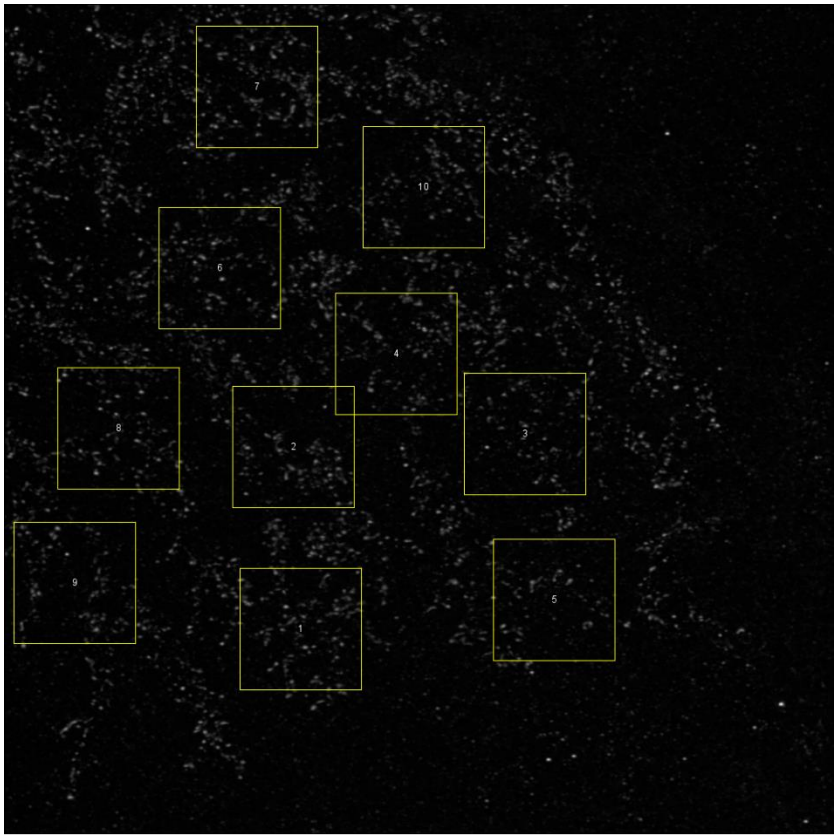

Figure 15: Ten validation regions selected for an image. Note that the regions that appear to overlap are in fact on different Z slices.

Table 3: Validation results for ten test regions

Region	Gold Standard	Detected	Fraction	Precision	Recall	F-measure
1	32	30	93.75%	80.00%	78.12%	79.05%
2	42	42	100.00%	76.19%	83.33%	79.60%
3	39	29	74.36%	79.31%	66.67%	72.44%
4	31	24	77.42%	100.00%	74.19%	85.19%
5	28	18	64.29%	77.78%	67.86%	72.48%
6	36	30	83.33%	70.00%	72.22%	71.09%
7	46	37	80.43%	70.27%	65.22%	67.65%
8	29	27	93.10%	81.48%	86.21%	83.78%
9	28	24	82.14%	91.30%	78.57%	84.46%
10	25	23	92.00%	82.61%	84.00%	83.30%
AVERAGE	**336**	**284**	**84.08%**	**80.89%**	**75.64%**	**77.90%**

Comparison to Other Methods

In order to demonstrate the effectiveness of these proposed methods, the results were compared to the traditional 3D Object Counter ImageJ plugin (Bolte & Cordelières, 2006). Both the global threshold-based Object Counter and its RATS threshold variant were tested on partial dendritic tree image 004. Additionally, the purely machine-learning driven BIOCAT can be used for comparison by calculating the centers of the objects produced by the 3D ROI annotation mode. Before performing any of these tests, the image was treated to the same general preprocessing techniques that were used to prepare the image for machine learning-guided detection. The morphology channel was subtracted from the post-synaptic channel, and rolling-ball background subtraction of radius 5 was applied.

Table 4 contains an overview of the comparison methods. Table 5 contains the full validation results for comparison to the Object Counter 3D ImageJ plugin. Table 6 contains the comparison to OC3D used with adaptive (RATS) thresholding, and finally Table 7 shows the comparison to direct annotation with BIOCAT.

Table 4: Overview of the comparisons performed on partial dendrite image 004

Test	Gold Standard	Detected	Fraction	Precision	Recall	F-measure
OC3D	336	418	125.21%	49.78%	75.40%	59.40%
OC3D-RATS	336	374	111.32%	61.05%	80.80%	69.16%
BIOCAT	336	320	95.26%	71.61%	75.60%	72.90%
Proposed	**336**	**284**	**84.08%**	**80.89%**	**75.64%**	**77.90%**

Table 5: Object Counter 3D validation. Performed with a global threshold

Region	Gold Standard	Detected	Fraction	Precision	Recall	F-measure
1	32	33	103.12%	57.58%	62.50%	59.94%
2	42	50	119.05%	62.00%	73.81%	67.39%
3	39	48	123.08%	54.17%	84.62%	66.05%
4	31	41	132.26%	60.98%	87.10%	71.73%
5	28	26	92.86%	34.62%	53.57%	42.06%
6	36	58	161.11%	43.10%	77.78%	55.47%
7	46	48	102.17%	44.68%	60.87%	51.53%
8	29	44	151.72%	45.45%	82.76%	58.68%
9	28	31	110.71%	51.61%	75.00%	61.15%
10	25	39	156.00%	43.59%	96.00%	59.96%
AVERAGE	**336**	**418**	**125.21%**	**49.78%**	**75.40%**	**59.40%**

Table 6: Object Counter 3D + RATS comparison

Region	Gold Standard	Detected	Fraction	Precision	Recall	F-measure
1	32	30	93.75%	53.33%	53.12%	53.23%
2	42	52	123.81%	59.62%	78.57%	67.79%
3	39	42	107.69%	64.29%	89.74%	74.91%
4	31	33	106.45%	84.85%	93.55%	88.99%
5	28	23	82.14%	56.52%	75.00%	64.46%
6	36	49	136.11%	57.14%	86.11%	68.70%
7	46	46	100.00%	65.22%	78.26%	71.15%
8	29	36	124.14%	58.33%	86.21%	69.58%
9	28	30	107.14%	56.67%	71.43%	63.20%
10	25	33	132.00%	54.55%	96.00%	69.57%
AVERAGE	**336**	**374**	**111.32%**	**61.05%**	**80.80%**	**69.16%**

Table 7: BIOCAT comparison

Region	Gold Standard	Detected	Fraction	Precision	Recall	F-measure
1	32	32	100.00%	71.88%	75.00%	73.40%
2	42	42	100.00%	71.43%	83.33%	76.92%
3	39	30	76.92%	70.00%	58.97%	64.02%
4	31	26	83.87%	88.46%	74.19%	80.70%
5	28	17	60.71%	64.71%	50.00%	56.41%
6	36	41	113.89%	68.29%	86.11%	76.17%
7	46	45	97.83%	57.78%	71.74%	64.01%
8	29	34	117.24%	67.65%	89.66%	77.11%
9	28	23	82.14%	82.61%	75.00%	78.62%
10	25	30	120.00%	73.33%	92.00%	81.61%
AVERAGE	**336**	**320**	**95.26%**	**71.61%**	**75.60%**	**72.90%**

Discussion

The proposed methods for synapse quantification are very promising. Even among neuroscientists, the validity of a proposed synapse is open to debate (Burette et al., 2015). Current synapse quantification is more often defined by the probability of a location being a synapse rather than a strict binary determination (Isaacson & Walmsley, 1995). The use of multiple channels to confirm the initial detection is in agreement with this ideology. By confirming with multiple factors, the final detected synapses are more trustworthy. Future investigations into automatic synapse detection could explore the benefits of assigning a probability to each candidate synapse, rather than the simple binary determination yielded by the current methods.

In the partial dendritic images, visual inspection as well as validation using the gold standard set indicates that the synapse detection is a success. There seem to be false negatives in the above Figure 12, but it is possible that these areas fail to meet one of the multi-channel comparison criteria or that the machine learning-guided extraction step determined these areas to be not similar enough to the training synaptic ROIs. In any case, the careful construction of the validation gold standard set was created to try to determine the overall performance.

As shown in Table 3 above, the average F-measure for the validation performed on partial dendritic image was 77.9%. The precision of these methods is on average higher than the recall by around 5%, indicating that the synapse detection is generally conservative in its

output. The synapses that are detected are 80.89% in agreement with the gold standard, but there are some false negatives that lower the recall. These validation statistics indicate that there is a general agreement between the human annotation and the results of the synapse detection methods, even if there is not complete certainty of the location of each synapse within these test regions.

The initial machine learning-guided detection step is the most critical and most complex component of the methods. This step was very successful and was instrumental in overcoming the challenges of intensity variance between synapses. This step was also very effective at eliminating background noise from the estimation. By restricting the machine learning algorithms to only a subset of the total image using a very lax RATS masking, the strengths of both segmentation and machine learning identification were leveraged while remaining efficient and effective. The learning model itself proved to be stable and reliable after training, and was able to successfully annotate multiple images within the partial dendritic image set with a good outcome.

Considering the comparisons to other methods, it can be seen in Table 4 that the proposed methods outperform both the naïve Object Counter 3D and simple BIOCAT annotation. This is likely due to several key factors. The multi-channel analysis that the proposed methods are capable of performing helps to eliminate false positives and increase overall precision of the detection. Indeed, the proposed methods yielded an average precision of 80.89%, outpacing the other comparison methods by a large margin. The OC3D-RATS test produced very high recall, but this is due to the gross overestimation of total synapses in the

image. The accuracy of this test was therefore low, at only 61.5%. The comparison to BIOCAT annotation is especially helpful because it highlights the strengths of combining the machine learning core with the other steps such as multi-channel comparison and adaptive thresholding. In the proposed methods, the RATS mask serves to narrow the focus of the machine learning detection, and then comparison to other channels increases overall precision by reducing the impact of false positives from noise in the synapse channel. It should also be noted that the restriction of the machine learning analysis by the RATS mask to only a subset of the total image voxels allows the proposed methods to act many times faster than BIOCAT's whole-image annotation.

Chapter 8

Conclusions

This thesis presents automatic methods for quantification of synapses in 3D confocal images. The need for automated image analysis is a major opportunity for both computer scientists and biologists to collaborate in order to develop solutions that are uniquely suited to this task. The methods presented here were ultimately successful in extracting and quantifying these synapses in a way that is unique among other biological image analysis methods. The combination of classic image processing techniques and machine learning algorithms for the analysis of biological images is an exciting research pursuit. The time investment it would take to manually perform a tally of synapses in one of these 3D images is staggering, and the proposed methods only take a few minutes to perform synapse detection, not including selecting training data for the discriminative model. Even with that factored in, only a tiny fraction of the synapses in an image need be counted by hand in order to produce a satisfactory outcome.

The specialized tools developed for this thesis will continue to be used for the exploration of synaptic images. These tools represent a new degree of interdisciplinary collaboration between biologists and computer scientists, a rising trend that promises to help

break down roadblocks impeding biological research while allowing many computer science techniques to find new life in new applications.

Additionally, the creation of a validation scheme and testing set for the partial dendritic image can be considered to be a success on its own, as there are no commonly available annotated 3D synapse image sets. Despite the validation regions being much smaller than the total image, they satisfactorily represent the whole and are a useful starting point for a quantitative assessment of the success of these synapse detection methods. While it would be better to have larger and more varied validation data, at the current time, those developed for this thesis are among the only such sets available.

Future Work

The methods presented in this thesis have demonstrated their usefulness in detecting synapses in complex 3D neuronal images, but an even more streamlined implementation for the complete pipeline is desirable. The actual implementation of these methods as a set of Java-based plugins and standalone tools could benefit from packaging into a single more user-friendly tool. Certainly these methods are many, many times faster than manual annotation already, even with the necessary user interaction to operate the separate programs. The major work of synapse detection is still automated, thus decreasing the impact of human error on the detection, which is one of the primary goals of these methods.

Bibliography

Abramoff, M. D., Magalhaes, P. J., & Ram, S. J. (2004). Image Processing with ImageJ. *Biophotonics International*, *11*(7), 36–42. Retrieved from http://imagej.nih.gov/ij/

Bjornsson, C. S., Lin, G., Al-Kofahi, Y., Narayanaswamy, A., Smith, K. L., Shain, W., & Roysam, B. (2008). Associative image analysis: A method for automated quantification of 3D multi-parameter images of brain tissue. *Journal of Neuroscience Methods*, *170*(1), 165–178. http://doi.org/10.1016/j.jneumeth.2007.12.024

Bolte, S., & Cordelières, F. P. (2006). A guided tour into subcellular colocalization analysis in light microscopy. *Journal of Microscopy*, *224*(3), 213–232. http://doi.org/10.1111/j.1365-2818.2006.01706.x

Burette, A., Collman, F., Micheva, K. D., Smith, S. J., & Weinberg, R. J. (2015). Knowing a synapse when you see one. *Frontiers in Neuroanatomy*, *9*, 100. http://doi.org/10.3389/fnana.2015.00100

Carpenter, A. E., Kamentsky, L., & Eliceiri, K. W. (2012). A call for bioimaging software usability. *Nature Methods*, *9*(7), 666–670. http://doi.org/10.1038/nmeth.2073

Chapelle, O., Haffner, P., & Vapnik, V. N. (1999). Support vector machines for histogram-based image classification. *IEEE Transactions on Neural Networks*, *10*(5), 1055–1064. http://doi.org/10.1109/72.788646

Chen, X., Velliste, M., Weinstein, S., Jarvik, J. W., & Murphy, R. F. (2003). Location proteomics - Building subcellular location trees from high resolution 3D fluorescence microscope images of randomly-tagged proteins, *Proc. SPIE, San Jose, CA, June* 2003, (Vol. 4962, pp. 298–306). http://doi.org/10.1117/12.477899

Defilipe, J., Alonso-Nanclares, L., & Arrellano, J. (2002). Microstructure of the neocortex : Comparative aspects. *Journal of Neurocytology*, *31*, 299–316. http://doi.org/10.1023/A:1024130211265

Fiala, J. C., Spacek, J., & Harris, K. M. (2002). Dendritic Spine Pathology: Cause or Consequence of Neurological Disorders? *Brain Research Reviews*, *39*(1), 29–54. http://doi.org/10.1016/S0165-0173(02)00158-3

Ghauharali, & Brakenhoff. (2000). Fluorescence photobleaching-based image standardization for fluorescence microscopy. *Journal of Microscopy*, *198*(2), 88–100. http://doi.org/10.1046/j.1365-2818.2000.00683.

Huxley, A. F. (2002). Hodgkin and the action potential 1935-1952. *The Journal of Physiology*, *538*(Pt 1), 2. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11773311

Isaacson, J. S., & Walmsley, B. (1995). Counting Quanta: Direct Measurements of Transmitter Release at a Central Synapse. *Neuron*, *15*, 875–884.

Kerschensteiner, D., Morgan, J. L., Parker, E. D., Lewis, R. M., & Wong, R. O. L. (2009). Neurotransmission selectively regulates synapse formation in parallel circuits in vivo. *Nature*, *460*(7258), 1016–1020. http://doi.org/10.1038/nature08236

Kim, J., Zhao, T., Petralia, R. S., Yu, Y., Peng, H., Myers, E., & Magee, J. C. (2012). mGRASP enables mapping mammalian synaptic connectivity with light microscopy. *Nature Methods*, *9*, 99–102. http://doi.org/10.1038/nmeth.1784

Kuzirian, M. S., & Paradis, S. (2011). Emerging themes in GABAergic synapse development. *Progress in Neurobiology*, *95*(1), 68–87. http://doi.org/10.1016/j.pneurobio.2011.07.002

Legland, D., Arganda-Carreras, I., & Schindelin, J. (2016). MorphoLibJ: MorphoLibJ v1.2.0. http://doi.org/10.5281/ZENODO.50694

Libbrecht, M. W., & Noble, W. S. (2015). Machine learning applications in genetics and genomics. *Nature Reviews Genetics*, *16*(6), 321–332. http://doi.org/10.1038/nrg3920

Liu, G. (2004). Local structural balance and functional interaction of excitatory and inhibitory synapses in hippocampal dendrites. *Nature Neuroscience*, *7*(4), 373–379. http://doi.org/10.1038/nn1206

Lodish, H., Berk, A., Zipursky, S. L., Matsudaira, P., Baltimore, D., & Darnell, J. (2000). Overview of Neuron Structure and Function. W. H. Freeman.

Maity, B., Sheff, D., & Fisher, R. A. (2013). Chapter 5 – Immunostaining: Detection of Signaling Protein Location in Tissues, Cells and Subcellular Compartments. In *Methods in Cell Biology* (Vol. 113, pp. 81–105). http://doi.org/10.1016/B978-0-12-407239-8.00005-7

Menon, V., Musial, T. F., Liu, A., Katz, Y., Kath, W. L., Spruston, N., & Nicholson, D. A. (2013). Balanced synaptic impact via distance-dependent synapse distribution and complementary expression of AMPARs and NMDARs in hippocampal dendrites. *Neuron*, *80*(6), 1451–63. http://doi.org/10.1016/j.neuron.2013.09.027

Meseke, M., Evers, J. F., & Duch, C. (2009). Developmental Changes in Dendritic Shape and Synapse Location Tune Single-Neuron Computations to Changing Behavioral Functions. *Journal of Neurophysiology*, *102*(1), 41–58. http://doi.org/10.1152/jn.90899.2008

Miles, R., Tóth, K., Gulyás, A. I., Hájos, N., & Freund, T. F. (1996). Differences between Somatic and Dendritic Inhibition in the Hippocampus. *Neuron*, *16*(4), 815–823. http://doi.org/10.1016/S0896-6273(00)80101-4

Morgan, J. L., Schubert, T., & Wong, R. O. (2008). Developmental patterning of glutamatergic synapses onto retinal ganglion cells. *Neural Development*, *3*(1), 8. http://doi.org/10.1186/1749-8104-3-8

Myatt, D. R., Hadlington, T., Ascoli, G. a, & Nasuto, S. J. (2012). Neuromantic - from semi-manual to semi-automatic reconstruction of neuron morphology. *Frontiers in Neuroinformatics*, *6*(March), 4. http://doi.org/10.3389/fninf.2012.00004

Navlakha, S., Suhan, J., Barth, A. L., & Bar-Joseph, Z. (2013). A high-throughput framework to detect synapses in electron microscopy images. *Bioinformatics (Oxford, England)*, *29*(13), i9–17. http://doi.org/10.1093/bioinformatics/btt222

Peng, H., Ruan, Z., Long, F. F., Simpson, J. H., & Myers, E. W. (2010). V3D enables real-time 3D visualization and quantitative analysis of large-scale biological image data sets, *28*(4), 348–353. http://doi.org/10.1038/nbt.1612

Scott, E. K., Raabe, T., & Luo, L. (2002). Structure of the vertical and horizontal system neurons of the lobula plate in Drosophila. *The Journal of Comparative Neurology*, *454*(4), 470–481. http://doi.org/10.1002/cne.10467

Shamir, L., Delaney, J. D., Orlov, N., Eckley, D. M., & Goldberg, I. G. (2010). Pattern Recognition Software and Techniques for Biological Image Analysis. *PLoS Comput Biol*, *6*(11).

Shrestha, B. R., & Grueber, W. B. (2011). Generation and Staining of MARCM Clones in Drosophila. *Cold Spring Harbor Protocols*, *2011*(8), pdb.prot5659–pdb.prot5659. http://doi.org/10.1101/pdb.prot5659

Soto, F., Bleckert, A., Lewis, R., Kang, Y., Kerschensteiner, D., Craig, A. M., & Wong, R. O. (2011). Coordinated increase in inhibitory and excitatory synapses onto retinal ganglion cells during development. *Neural Development*, *6*(1), 31. http://doi.org/10.1186/1749-8104-6-31

Stalling, D., Stalling, D., Westerhoff, M., & Hege, H.-C. (2005). Amira: a highly interactive system for visual data analysis. *THE VISUALIZATION HANDBOOK*, 749–767. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.129.6785

Turbes, C. C., & Schneider, G. T. (1989). Directionality of neural signals in central nervous system neural networks. *Biomedical Sciences Instrumentation*, *25*, 1–5. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2545289

Wählby, C., Lindblad, J., Vondrus, M., Bengtsson, E., & Björkesten, L. (2002). Algorithms for cytoplasm segmentation of fluorescence labelled cells. *Anal Cell Patho*, *24*, 101–111.

Webb, R. H. (1996). Confocal optical microscopy. *Rep. Prog. Phys*, *59*, 427–471.

Wilkinson, M. H. F. (1998). Optimizing Edge Detectors for Robust Automatic Threshold

Selection: Coping with Edge Curvature and Noise. In *Graphical Models and Image Processing* (Vol. 60, pp. 385–401). http://doi.org/10.1006/gmip.1998.0478

Xiao, H., & Peng, H. (2013). APP2: automatic tracing of 3D neuron morphology based on hierarchical pruning of a gray-weighted image distance-tree. *Bioinformatics (Oxford, England)*, *29*(11), 1448–1454. http://doi.org/10.1093/bioinformatics/btt170

Zhou, J., Lamichhane, S., Sterne, G., Ye, B., & Peng, H. (2013). BIOCAT: a pattern recognition platform for customizable biological image classification and annotation. *BMC Bioinformatics*, *14*(1), 291. http://doi.org/10.1186/1471-2105-14-291

Zhou, J., & Peng, H. (2011). Counting cells in 3D confocal images based on discriminative models. In *Proceedings of the 2nd ACM Conference on Bioinformatics, Computational Biology and Biomedicine - BCB '11* (pp. 399–403). http://doi.org/10.1145/2147805.2147858