

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Activity Mapping the Leech Nervous System

Permalink

<https://escholarship.org/uc/item/2rh0z6m0>

Author

Frady, Edward Paxon

Publication Date

2014

Supplemental Material

<https://escholarship.org/uc/item/2rh0z6m0#supplemental>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Activity Mapping the Leech Nervous System

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

in

Neuroscience
Specialization in Computational Neuroscience

by

Edward Paxon Frady

Committee in charge:

Professor William B. Kristan, Chair
Professor Gert Cauwenberghs, Co-Chair
Professor Timothy Gentner
Professor Massimo Scanziani
Professor Terrence Sejnowski

2014

Copyright
Edward Paxon Frady, 2014
All rights reserved.

The dissertation of Edward Paxon Frady is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Chair

University of California, San Diego

2014

TABLE OF CONTENTS

Signature Page	iii
Table of Contents	iv
List of Supplemental Files	v
List of Figures	vi
List of Tables	viii
Acknowledgements	ix
Vita	xi
Abstract of the Dissertation	xii
1 Introduction	1
2 Computation with Population Codes	6
3 Population Coding and Gain Control	25
4 Experimental Methods	50
5 The Imaging Computational Microscope	74
6 Scalable semi-supervised cell identification reveals canonical swim and preparatory networks.	101

LIST OF SUPPLEMENTAL FILES

SupplementaryAppendixA.pdf: Supplementary appendix for optically monitoring voltage in neurons by photoinduced electron transfer through molecular wires.

SupplementaryAppendixB.m: Supplementary code for the imaging computational microscope.

SupplementaryAppendixC.pdf: Supplementary appendix for scalable semi-supervised cell identification reveals canonical swim and preparatory networks.

LIST OF FIGURES

Figure 2.1:	The standard model of population coding.	9
Figure 2.2:	Populations encode multiple modalities of information.	16
Figure 2.3:	Transformation of population codes.	19
Figure 2.4:	Population coding in time.	21
Figure 3.1:	Local bend is a directional response that employs a population code.	27
Figure 3.2:	Models of additive vs. multiplicative synapses for controlling gain.	31
Figure 3.3:	Modulation of DE-3’s input-output function.	39
Figure 3.4:	Shunting inhibition predicts synaptic structure of DE-3.	42
Figure 3.5:	Multiplication via shunting in two-compartment model.	44
Figure 4.1:	Leech Behaviors	51
Figure 4.2:	Leech Preparations	52
Figure 4.3:	Leech Nervous System	53
Figure 4.4:	Mechanisms of fluorescent voltage sensing.	56
Figure 4.5:	Synthesis of VF probes.	57
Figure 4.6:	Characterization of VF sensors in HEK cells.	59
Figure 4.7:	Characterization of the speed, wavelength sensitivity, and capacitance of the VF2 fluorescence response.	61
Figure 4.8:	VF2 dyes resolve action potentials in neurons.	62
Figure 5.1:	ICM screenshot with labeled regions.	78
Figure 5.2:	ICM screenshot in pre-processing stage.	79
Figure 5.3:	ICM screenshot in PCA stage.	80
Figure 5.4:	ICM screenshot in ICA stage.	82
Figure 5.5:	ICM screenshot in segment stage.	83
Figure 5.6:	Example components	84
Figure 5.7:	ICM screenshot in visualization stage.	85
Figure 5.8:	ICM visualization reveals patterns in large-scale imaging data.	86
Figure 5.9:	Examples of ctICA components.	88
Figure 5.10:	Segmentation and Clustering	89
Figure 5.11:	Visualization of multiple trials with ctICA.	91
Figure 5.12:	Calculation of Signal-to-Noise Ratio (SNR).	93
Figure 5.13:	Sub-pixel motion correction with ECC.	94
Figure 5.14:	Comparison of ROI and ICA based signal extraction.	95
Figure 5.15:	SNR comparison of different methods.	96
Figure 6.1:	Voltage-sensitive dye imaging during several behaviors.	103
Figure 6.2:	Semi-supervised extraction of cell signals with ctICA	107
Figure 6.3:	Feature Extraction	108
Figure 6.4:	Feature Variability	111
Figure 6.5:	Canonical learning and visualization of medium-dimensional feature space.	114
Figure 6.6:	Table of cells in canonical swim and prep networks	116
Figure 6.7:	Canonical swim and preparatory networks	119

Figure 6.8:	Electrophysiological verification of targeted swim oscillator neurons. . . .	122
Figure 6.9:	Preparatory network shows rapid responses regardless of behavior. . . .	123
Figure 6.10:	The S cell is functionally coupled to the preparatory network.	125
Figure 6.11:	Electrophysiological validation of predicted S cell connections.	127

LIST OF TABLES

Table 4.1: Summary of VSD attributes	64
--	----

ACKNOWLEDGEMENTS

I have had an amazing experience during my time in graduate school, and I would like to thank everyone that has made this possible. My tenure as a TA for bootcamp each year has granted me a wonderful opportunity to grow with each new incoming graduate class. This has been an unforgettable experience, it has given me an incredible set of friends, a phenomenal community, and I have come to care deeply for all of my “kids” and the Neuroscience program.

My time here has further been enhanced by a great group of close friends and I would like to thank them all for making life during graduate school a fun experience. I will always look fondly back at my times here, exploring San Diego breweries and debating everything with my college friends, Brian Kearns, Nicholas Kramer, and all of the waiters and beavers; throwing Mesa parties and playing poker with Stephen Larson, Dave Matthews, Jeff Gauthier, Lee Lovejoy, Sam Nummella, Doug Rubino, Adam Calhoun and Dan Hill; going on adventures, watching crazy movies, discovering San Diego, trivia, doing Science and learning who I am with Cristina Nigro, Alison Boyd, Krista Perks, Kim Reinhold, Camille Toarmino, Cailey Bromer, Jason Pipkin, Mindi Summers, Bassam Attallah, Shawn Olsen, Fitz Sturgill.

I would like to thank William Kristan for being a wonderful advisor and mentor. Bill has been everything I could want in an advisor, his great knowledge and experience has been a guiding light and he has given me so many wonderful opportunities. I want to thank my committee, Gert Cauwenberghs, Tim Gentner, Massimo Scanziani, and Terry Sejnowski, for giving me tremendous feedback and support over the years as well as being a great source of inspiration for my science and as a scientist.

I would also like to thank and recognize all of the people who have worked with me on these projects and taught me new techniques and ideas: Joshua Yudice, Vaibhav Konanur, Krista L. Todd, Kathy A. French, Evan W. Miller, John Y. Lin, Paul A. Steinbach, Roger Y. Tsien, Cliff Woodford, Daniel Wagenaar, JOsh Poole, Eric Horvitz, Ashish Kapoor, and Kevin Briggman.

Chapter 2, in full, is a reprint of material as it appears in *Computation with Population Codes*. 2013. Frady, E.P., Kristan, W.B., In: Jaeger D., Jung R. (Ed.) *Encyclopedia of Computational Neuroscience*: Springer Reference (www.springerreference.com). Springer-Verlag Berlin Heidelberg, 2013. DOI: 10.1007/SpringerReference_348400. The dissertation author was the author of this paper.

Chapter 3, in part, is under preparation for submission as a publication as Frady, E.P., Yudice, J., Konanur, V., Todd, K.L., French, K.A., Kristan Jr, W.B. “Shunting and inhibitory currents co-regulate the input-output function of an identified leech neuron”.

Chapter 4, in part, is a reprint of material as it appears in Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires. 2012. Miller, E.W., Lin, J.Y., Frady, E.P., Steinbach, P.A., Kristan, W.B., Tsien, R.Y. PNAS 109(6): 2114-2119. The dissertation author was a co-author of this paper.

Chapter 5 and 6, in full, are under preparation for submission as a publication as Frady, E.P., Kristan W.B., Jr “The Imaging Computational Microscope.” and Frady,E.P., Kristan W.B., Jr “Scalable semi-supervised cell identification reveals canonical swim and preparatory networks.”

VITA

2004-2008	B. S. with Honors in Computation and Neural Systems, California Institute of Technology
2008-2014	Ph. D. in Neuroscience, University of California, San Diego
2009-2013	Neurosciences Bootcamp Teaching Assistant
2009-2013	Bootcamp Computational Special Project
2010	Advanced Imaging and Electrophysiology Lab Teaching Assistant
2012-2014	Analytical Methods in Computational Neuroscience
2013-2014	Neural Systems & Behavior Teaching Faculty

PUBLICATIONS

Frady, E.P., Kristan, W.B. (2013). Computational with Population Codes. In: Jaeger, D., Jung, R. (Ed.) *Encyclopedia of Computational Neuroscience*: Springer Reference. Springer-Verlag Berlin-Heidelberg, 2013.

Frady, E.P., Palmer, C.R., Kristan, W.B. (2012) Sexual Attraction: Sex-Specific Wiring of Neural Circuitry. *Current Biology* 22(22): R953-R956.

Miller, E.W., Lin, J.Y., Frady, E.P., Steinbach, P.A., Kristan, W.B., Tsien, R.Y. (2012). Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires. *PNAS* 109(6): 2114-2119.

Cerf, M., Frady, E.P., Koch, C. (2009). Faces and text attract gaze independent of task: Experimental Data and Computer Model. *Journal of Vision* 9(12): 1-15.

Cerf, M., Frady, E.P., Koch, C. (2008). Using semantic content as cues for better scanpath prediction. *ETRA* 143-146.

Einhauser, W., Rutishauser, U., Frady, E.P., Nadler, S., Konig, P., Koch, C. (2006). The relation of phase noise and luminance contrast to overt attention in complex visual stimuli. *Journal of Vision* 6:1148-58.

ABSTRACT OF THE DISSERTATION

Activity Mapping the Leech Nervous System

by

Edward Paxon Frady

Doctor of Philosophy in Neuroscience
Specialization in Computational Neuroscience

University of California, San Diego, 2014

Professor William B. Kristan, Chair
Professor Gert Cauwenberghs, Co-Chair

Neural circuits represent and process information using large numbers of component neurons. In this thesis, I describe the current theories of how information is processed by nervous systems, biophysical models of the basic mechanisms of population coding, and experimental techniques and computational algorithms to collect and synthesize the data necessary to understand neural computation in simple nervous systems. The leech nervous system is an ideal system for study, because it is readily accessible and consists of core neural structure, the ganglion, built of only a few hundred neurons. Even though this system is simple compared to mammalian nervous systems, it is still remarkably complex. To understand this complexity and make sense of the patterns of neurons, large-scale recordings of neural activity are required. We have developed a new Voltage-Sensitive Dye to record from a significant fraction of the

nervous system simultaneously during behavioral states. This type of large-scale data presents entirely new challenges to overcome, and we developed computational tools to visualize and stitch-together these large-scale recordings. By imaging from a ganglion in several animals and recording neural responses during several different behaviors, we have developed a system for scalable and rapid identification of dozens of individual neurons. With these tools and techniques, we have mapped out the activity of a significant fraction of the leech's nervous system and have identified dozens of novel cells. Many of these cells are part of two canonical networks in the leech nervous system: the swim and preparatory networks. We further show that the preparatory network is mediated by the S cell, and we use computationally guided electrophysiology to target and verify that the S cell drives the activity of other cells in the preparatory network. These tools enable us to study the nervous system at scale for the first time, and we have mapped out the roles of a significant fraction of the neurons during the production of behaviors. These are just the first steps necessary to build a complete picture of how leech neurons produce behavior and make decisions.

1 Introduction

The general question “How do neurons produce behavior?” is the foundation of neuroscience. This is the question that drives the following work. We are now beginning to understand that computation is the underlying force behind how nervous systems produce behaviors, but we have only just begun to unravel the mystery. For years, neuroscientists have continually thought that we have been on the precipice of unlocking the brain’s mysteries; that the great breakthroughs revealing the nature of the mind are tantalizingly close; that we are right around the corner from knowing how the brain works. But with each great breakthrough comes a deeper understanding of the complexity that truly underlies the brain’s processes. With each breakthrough, we realize that the precipice of understanding is further away than we naively thought.

Today, Neuroscience is a gargantuan field, reflecting the immense complexity of the brain. Molecular biologists, physicist, psychologists, computer scientists, physiologists and more are spread out studying brain phenomena across many scales of both space and time. The complexity of the brain is one of the biggest obstacles to overcome, and understanding this complexity is a new frontier for science. What does it take to understand complex systems? How do we extract, synthesize and make sense of the large-scale data needed to crack the complexity of the brain? These are fundamental problems that need solutions before nature will reveal her secrets of the mind. The science of complexity is something new and different than traditional science. Reductionism has been our main weapon for uncovering the Universe, probing and proving hypotheses with carefully controlled experiments. But reductionism naturally breaks apart the whole into constituent pieces; it seeks to make sense of the pieces individually, which then reveals the nature of the system. But with complex or emergent systems, the pieces don’t fully explain the whole – the whole is greater than the sum of its parts. The mind is an emergent phenomenon, and to understand the brain means we have to understand emergence; we have to not only understand the pieces, but also how the pieces work together. In order to understand the system, we have to see every aspect of the system, we have to map out every interaction and put

it all together into a coherent framework.

The leech is attractive as an experimental system, because its nervous system is relatively simple and the ability to see every aspect of neural activity was actually feasible. This simplicity is only in comparison with more complex organisms, and the complexity of even just one part of the leech brain is still mind-boggling. New techniques in Neuroscience, such as Voltage-sensitive dye (VSD) imaging, are allowing us to collect data on an extraordinary scale, and this is finally enabling us to collect enough information to begin to look at nervous systems in a holistic fashion. However, even though imaging allows us to view the nervous system at unprecedented scale and resolution, it is still not enough to get a complete holistic view. We need to piece together and synthesize data from many sources – we need to map out the function of every cell, we need to be able to identify the same patterns and cells across animals, we need to connect function and anatomy, and we need to guide reductionist experiments with emergent observations. In the leech, we can get closer to this than perhaps any other organism, but we are just beginning to put the puzzle together.

There is a fundamental connection between complexity, information and computation. The nervous system performs computation in an entirely different way than our silicon computers. In chapter 2, this thesis explores the basic foundations of neural computation. Population coding is a promising theory of the basis of information processing by nervous systems. I introduce the basic concepts of using a population code for encoding and decoding information, encoding high-dimensional information, and computing with this information.

The computational theory of population coding highlights the importance of both linear and non-linear modulation of neural activity. Multiplication is prominent across various nervous systems, and has appealing properties for neural computation. This theoretical work spurred us to look for inhibitory neurons which have qualitatively different modulatory properties, and we show in chapter 3 how DI-1 and 116 additively and multiplicatively modulate their post-synaptic target neuron, DE-3. For decades, we have been theorizing and studying the mechanisms of these basic operations, but a long-standing idea, which suggested that the mechanism of multiplication was shunting-inhibition, had been debunked by recent theoretical and experimental observations. We develop a new theoretical framework for the mechanism of these qualitative types of responses, and illustrate how shunting inhibition can multiplicatively modulate spiking. This is possible, because recent work has shown that neurons do not act like “point” integrators, but are rather elaborate multi-compartmental structures. When the constraint of a point neuron is removed, and neurons are modeled with many compartments, the mechanism of shunting-

inhibition implementing multiplication is possible.

These chapters serve to highlight the general principles and mechanisms of neural computation, and reveal how the nervous system utilizes the properties of chaotic systems and high-dimensional space. Now that we have a basic understanding of the mechanisms of computation, we need to begin to address the algorithms that are being used by the nervous system. If we make a metaphor to silicon-computing, then the theory of population coding is like the theory of boolean algebra; we understand how ‘0’s and ‘1’s are used to perform computation on a computer, but we have yet to understand what exactly is the nervous system computing. What is the software that is running on the neural hardware?

To begin to understand the software of the brain, we must explore the properties of the high-dimensional space that nervous systems utilizes to perform computation. In the context of the nervous system, we identify the dimensions of the high-dimensional space as the activity of individual neurons. To explore this space, we must collect a tremendous amount of information. The higher the dimensionality of this space the more information that we must collect. This is one of the primary reasons that we study relatively simple nervous systems – the dimensionality of the human brain is so vastly larger that we are hopeless to collect enough information to understand how the system works. But the dimensionality of even a few hundred neurons, like in the leech, is immense. To make sense of it requires advanced techniques in large-scale data collection and analysis.

In chapter 4, we describe the techniques for exploring the high-dimensional computational space of the leech’s nervous system. The simplicity and accessibility of the leech nervous system gives us a feasible chance at exploring this high-dimensional space, yet a tremendous amount of data still needs to be collected and synthesized. We utilized VSD imaging to collect all of this data, and recent progress in these dyes has given us an unprecedented view of the nervous system in action.

But handling and making sense of this data is a challenge all on its own. We still need to synthesize, visualize, and understand data from multiple experiments to piece together the leech’s nervous system. Chapter 5 describes the computational tools that we have built in order to achieve these goals. This tool utilizes principal and independent component analysis to automatically extract cellular signals from imaging data and remove artifacts, and can be used for imaging in any modality. We illustrate how to use this tool and how to create “activity maps” to visualize large-scale neural data by extracting “features” from neural responses. This enables rapid visualization of large-scale neural activity, which can then be used to target and guide

experiments with real-time analysis.

With the tools of the computational microscope, we have mapped out the swim oscillator network and the preparatory network in the leech's nervous system. In Chapter 6, we go over the experiments and computational techniques used to piece together these networks. This chapter illustrates a basic problem in all large-scale recording efforts – the problem of identification. Our visual systems are so capable at recognizing and identifying patterns, that we often take this task for granted, and we usually trust that a targeted neuron can be identified through visual inspection. However, with large-scale data, the task of identification becomes the task of identifying many things simultaneously. This is a much more challenging problem, and we can no longer rely solely on human pattern recognition to solve this problem. As neuroscience progresses and the vast data sets needed to explore high-dimensional space are collected, we will have to get computers and machines to look at the data for us.

Currently, however, machine-learning is nowhere close to the sophistication of the human visual system. For virtually every problem in data analysis, we can make a computational algorithm that is correct most of the time, but not all of the time, and this lost information is intolerable. This is why the framework of the analysis system is “semi-supervised”, which means that machine-learning and algorithms are used to do the majority of the grunt work and human oversight refines the final results. Semi-supervised allows the system to learn from human intervention, and produce high-quality results fast.

This work sets the stage for understanding and experimenting on a nervous system at scale. The complexity of nervous systems demands new techniques for collecting and synthesizing vast amounts of data, as well as the ability to use this synthesis to guide experiments and ask novel questions about the behavior of systems of neurons. We illustrate how our efforts have enabled computationally guided electrophysiology – experiments guided by synthesized data in real time. These experiments require collecting large-amounts of data and putting this high-dimensional information in context with previous work. The computational tools we have built enable this process to occur in real-time, and we verify predictions and test new hypotheses that could only be tested with large-scale data put in context.

The long-standing dream of this work is to set the stage for collecting the highest resolution and most complete data set connecting structure and function of any nervous system. In collaboration with my graduate colleague Jason Pipkin and the Ellisman lab, we have been working to link the high-resolution functional recordings made with VSD imaging with the high-resolution structural data obtained by serial block-face electron microscopy (SBFEM). Develop-

ing a scalable-identification system and mapping out the activity of a large fraction of neurons in the leech nervous system is a necessary first step towards this dream. With this activity map, we can now identify dozens of cells in the leech nervous system, characterize their activity, and then obtain detailed information about their anatomy and connections to each other using SBFEM. Our hope is to carry out both the VSD and SBFEM experiments in a single animal. This, perhaps, will provide enough information to understand the link between structure and function in the nervous system, and to finally answer the deepest mysteries about how neurons compute, generate behaviors, and make decisions.

2 Computation with Population Codes

Abstract

Population coding is a computational theory postulating that information is represented and processed by a large number of neurons. In such a coding scheme, each neurons on its own encodes only a small amount of the information that is distributed across the population. Population coding provides robustness, because even if individual neurons are noisy or fail altogether, the population code still processes information effectively.

Fundamentals of computation

To understand how a population code computes, we will first start by describing the basics of computation by analyzing silicon-based computers. There are two essential ingredients to computation: representation and transformation. Information in computers is represented using discrete binary digits, called bits. A bit is often thought of as a “yes” or “no” response, or a “1” or “0”. A bit represents 2 states and is the simplest form of information. A computer is built by utilizing a large number of bits to represent information. The number of elements that can be represented grows exponentially with the number of bits that are used. Eight bits (a “byte”) can represent up to 256 (2^8) different values. Large numbers are represented by more and more bits: an integer is typically represented by 32 bits, and a “double” is a decimal number that is represented by 64 bits. A computer can represent anything by using more and more bits of information. A 2-hour movie, for example, is currently represented with about 8 billion bits.

Computation is in essence the transformation of information from one representation to another. For instance, a computer represents—or “encodes”—each letter in this text by a byte (8 bits). For you to read this text the computer must transform each letter into an image, which is made of pixels that are also represented by bits. The computer transforms the byte that represents each letter into a different sequence of bits that represents the pixel colors. The specific rules of the transformation are called the “algorithm”. The monitor then “decodes” the bits that represent

the pixel colors and creates an image. A computer can perform “universal” computation, which means that it has the ability to transform any string of bits into any other arbitrary string of bits. This ability is what makes a computer so powerful, because it can transform any type of information in one representation into any other representation.

The brain, too, must represent a large number of variables produced by the sensory environment or intrinsic states. To perform their functions, neurons in the brain must be able to encode, transform, and decode this information. However, unlike the bits on a computer, neurons are intrinsically noisy. The failure of a single bit in a computer could have catastrophic effects on the entire system, but by design there is virtually no noise in a computer, so these types of failures are rare. Biological systems, on the other hand, have a great deal of noise of many sorts—both from the intrinsic properties of the component neurons and from ongoing neural activity from sources other than the signal of interest. The neurons must perform their computations in the face of many potential failures caused by this noise. The brain appears to solve this noise problem by distributing information across a large population of neurons. Having more neurons entails robustness to the encoded information, because a large number of independent neurons represent the information. By spreading the information about different variables across many neurons, very little information is lost if a single neuron misfires or even dies.

Recent theoretical work in systems neuroscience has developed the “Population Coding” framework of neural computation, based on this premise that information is processed across a large population of neurons. Population coding computes differently than a computer, but it still must perform the same fundamental operations of information processing: representation and transformation. The brain receives information from the outside world and must represent this information in the activity of neurons. The brain then transforms this information in different brain regions to other representations that are useful for perception or behavior. For instance, to grab a baseball, your brain transforms visual information that is stored in an eye-centered coordinate system into a body-centered coordinate system. The same information is stored in both coordinate systems, but with different representations. The motor cortex then uses the body-centered coordinate system to generate a motor sequence that contracts and relaxes the correct muscles to reach out while opening the hand, then stop the reaching and close the hand as it envelops the baseball. For the brain to perform such a behavior, it must have a representation of all the relevant information, as well as the ability to transform one representation into another.

Representation of single variable in population codes

The first step for a population code is to represent information through the activity of neurons. A good way to understand how a population code can represent information is to decode a single variable from neural responses. A classic example of population coding is provided by determining how a simple cell in visual cortex responds to bars oriented at different angles. A given simple cell responds maximally to a bar that is oriented at the neuron's "preferred" angle, but this neuron also responds to orientations that are close to this preferred angle. The relationship between the stimulus and the firing rate of the neuron creates the neuron's "receptive field", which peaks at the preferred angle and decays as the stimulus angle moves away from the preferred angle. The receptive field is typically a bell-shaped curve, often taken to be Gaussian for nonperiodic variables and a cosine function for periodic variables (Deneve et al. 1999).

A large number of neurons make up the population, where each neuron has its own receptive field. The preferred orientations of each neuron are spread out across the full stimulus space. Figure 2.1A shows several example receptive fields from 8 different neurons; each bell-shaped curve is the receptive field of a different neuron. These receptive fields partially overlap and are centered at different preferred angles. When a stimulus is presented, all neurons respond based on their receptive fields. When an oriented bar is presented that has an angle of $\theta = 90^\circ$, then several neurons with preferred orientations close to 90° will fire (Fig. 2.1B). The closer the neuron's preferred orientation is to 90° , the higher the neuron's firing rate, but noise can make the firing rate variable for a given stimulus presentation. The activity of the population is described by:

$$r_j = f_j(\theta | \Theta_j) + \varepsilon \quad (2.1)$$

where r_j is the firing rate of neuron j , Θ_j is the preferred orientation of neuron j , f_j is a function that describes the neuron's receptive field and ε is noise. This is the "standard model" of population coding.

There are various approaches used to come up with an estimate, $\hat{\theta}$, of the orientation of the stimulus encoded by the firing rates of the population of neurons. All approaches obtain an estimate by computing a function from the activity across the population R , where $R = (r_1, r_2, \dots)$. Because of noise, the response across the population varies during different presentations of the same identical stimulus. This variability means that the estimate, $\hat{\theta}$, is a random variable. Different techniques for decoding the estimate from the population activity have been developed, each with different tradeoffs: in bias (whether the $\hat{\theta}$ is right on average), variance (how close $\hat{\theta}$

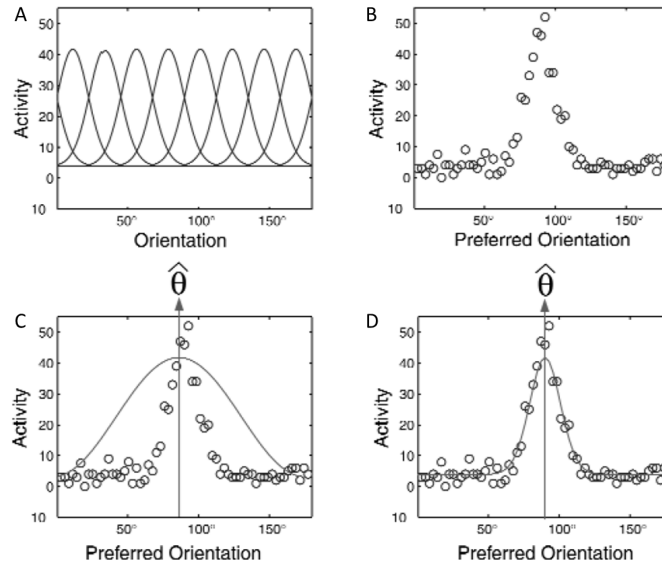


Figure 2.1: The standard model of population coding. (A) A population of neurons have receptive fields that are sensitive to different orientations. Each bell-shaped curve is from a different neuron, and the neuron’s preferred orientations are spread out across the stimulus space. y-axis: spikes per second. (B) When an orientation of $\theta = 90^\circ$ is presented, the neurons respond based on their receptive fields shown in A. (C) The population vector estimator is a voting method that decodes the orientation from the neural activities by fitting a cosine function to the responses. (D) The maximum likelihood estimator fits a template derived from the receptive field shapes. From: Deneve, Latham & Pouget, 1999.

is to θ on average), computational efficiency, and biological plausibility (Pouget et al. 1998).

One interpretation of the population code representation is that each neuron “votes” for its preferred stimulus. Such schemes are typically computationally efficient and biologically plausible, but have higher variances than other decoding methods (they are only optimal under specific conditions as discussed below). Three commonly used voting schemes are:

1. Population vector estimator (PVE). Bar orientation is a spatially periodic variable and the PVE decodes such periodic variables by fitting a cosine function (Fig. 2.1C) to the activity patterns, R :

$$\hat{\theta}_{PVE} = \text{phase}(z) \quad (2.2)$$

where

$$z = \sum_j r_j e^{i\Theta_j} \quad (2.3)$$

2. Optimal linear estimator (OLE). Because the PVE works only for periodic variables, other measures must be used for nonperiodic variables. A standard algorithm for decoding a nonperiodic variable from a population code is the OLE:

$$\hat{\theta}_{OLE} = \sum_j \Theta_j r_j \quad (2.4)$$

3. Center of mass estimator (COM). A simple extension of the OLE, the COM can account for the background firing (one form of noise) by comparing the neural responses when the stimulus is presented to the background rate when the stimulus is absent:

$$\hat{\theta}_{COM} = \frac{\sum_j \Theta_j (r_j - b)}{\sum_j (r_j - b)} \quad (2.5)$$

where b is the spontaneous activity of the neurons. All three of these estimators are “voting methods”, because each spike from a neuron is like a vote for that neuron’s preferred stimulus feature, and the estimate is a simple combination of all of the votes. Unlike these voting methods other estimators can decode the population activity optimally, which means the variance of the estimate is minimized and there is no bias. These estimators are less biologically plausible, however, and less computationally efficient than the voting methods:

4. Maximum likelihood estimator (ML). This estimator finds the value of θ that maximizes the conditional probability of the population response, $P(R|\theta)$, and is optimal when θ is uniformly distributed:

$$\hat{\theta}_{ML} = \underset{\theta}{\operatorname{argmax}} P(R|\theta) \quad (2.6)$$

The ML estimator is similar to the PVE in that it is effectively fitting a template to the population response (Paradiso, 1988) (Fig. 2.1D). However, the PVE assumes that the shape of the template is a cosine, whereas the ML derives the template directly from the receptive fields of the neurons, which means that the decoder must have knowledge of the receptive field shapes of the encoding neurons. This may be a biological implausibility, but it is also possible that the receptive field shapes are learned over time (Sanger, 2003).

5. Maximum a posteriori estimator (MAP). This estimator is an extension of the ML estimator that can incorporate the prior distributions of R and θ . The MAP estimate maximizes the full

posterior distribution of the encoded variable, $P(\theta|R)$, which is derived from the conditional probability through Bayes's theorem:

$$\hat{\theta}_{MAP} = \underset{\theta}{\operatorname{argmax}} P(\theta|R) \quad (2.7)$$

$$P(\theta|R) = \frac{P(R|\theta)P(\theta)}{P(R)} \quad (2.8)$$

This is a useful modification if certain values of θ are more likely than others; e.g. if horizontal orientations are seen more frequently than vertical orientations. If the prior distribution is flat – meaning that all values of θ are equally likely, then the MAP estimate is identical to the ML estimate.

High dimensional population codes

The activity of a population of neurons can be considered to form a high-dimensional response space: two neurons would form a two-dimensional space (a plane), three neurons would form a three-dimensional space, etc. If the activity of three neurons are plotted as the x, y, and z coordinates in a three-dimensional Cartesian system, then a particular stimulus activates these neurons in a particular firing pattern, which would correspond to a point, called the “population vector”, in three-dimensional space. Each additional neuron would add a dimension to this response space. As the firing rate of neuron j increases, the population vector would move up in the j th dimension. If we had 10 neurons representing the 1-dimensional orientation, then each presentation of a particular orientation would activate the 10 neurons in a specific way that would correspond to a population vector in the 10-dimensional response space. Each orientation would cause a different pattern of activation, which would correspond to a different population vector in the 10-D space. Together, the responses to all possible orientations would form a 1-D “manifold” (a curvy line) in the 10-D response space.

In addition to representing information about one stimulus variable, a population of neurons can encode multiple variables simultaneously. A signal could have several feature dimensions, all of which are encoded by the brain. For instance, a population of neurons could simultaneously encode both the orientation and the thickness of a bar stimulus. In this population, each neuron would have a two-parameter receptive field that has both a preferred orientation and a preferred thickness:

$$r_j = f_j(\theta, \phi | \Theta_j, \Phi_j) + \varepsilon \quad (2.9)$$

where ϕ is the thickness of the bar, and Φ_j is the thickness preferred by neuron j . In this representation, each combination of orientation and thickness would correspond to a single point in the high-dimensional response space (with each dimension again corresponding to the activity of one neuron). All possible combinations of bar orientation and bar thickness would form a 2-dimensional manifold (a curvy plane). In one direction along this manifold, the information about orientation is represented, and in another direction information about thickness is represented. In general, the population can encode multiple dimensions of the stimulus, where each neuron can be sensitive to any combination of stimulus dimensions:

$$r_j = f_j(\theta, \phi, \psi, \dots | \Theta_j, \Phi_j, \Psi_j, \dots) + \varepsilon \quad (2.10)$$

An individual neuron can have a preference for an arbitrary combination of each of the stimulus dimensions. The preferred orientation and thickness of a given neuron can be chosen independently. For instance, a neuron with a preferred orientation of $\Theta_j = 90^\circ$ could be just as likely to respond to a thick bar as a thin bar.

The information about a particular feature of the stimulus does not necessarily correspond to the firing of a particular neuron. In the population code, the neurons represent stimuli through their joint activities. When representing multiple stimulus dimensions, the neurons are activated by all features of the stimulus with different strengths. The orientation of a stimulus would thus be only one of many components of the response for a given neuron. The information about orientation is encoded in the combination of activities across the population. This can be thought of as a particular direction in the high-dimensional space, but not as the firing of a particular neuron.

Just as adding more neurons adds more dimensions to the population code, adding more features to the stimuli would increase the number of dimensions in the manifold representing those features. In general, the larger the number of dimensions (i.e., neurons), the larger the number of features that a particular neuronal population can represent. This relationship between the number of neurons and the number of features that can be encoded requires a consideration of “basis sets”.

Basis sets

For the same neurons to encode more than one feature, the activity of the neuronal population must create a “complete basis set” for these features. To understand what such a requirement means, consider a particular set of neurons with receptive fields that respond to both orientation and thickness, but these receptive fields are actually perfectly correlated: neurons that prefer 90° orientations also prefer thick bars, and neurons that prefer 0° also prefer thin bars. If the receptive fields are correlated in such a way, then there are different combinations of orientations and thicknesses that activate the population in exactly the same way. If the population response is the same to two different stimuli, then there is less information about the stimulus available in the activity. Technically, this inability of the population to fully encode the two features arises because we have intentionally made the receptive fields “linearly dependent”. For the neurons to represent all the dimensions of the stimulus they must form a complete basis of the stimulus space: the manifold determined by the neurons’ firing rates in the response space must have as many dimensions as the number of features. Because we have made the receptive fields perfectly correlated, every combination of thickness and orientation no longer forms a 2-dimensional manifold, but instead forms a 1-dimensional manifold. Every point in the 1-dimensional manifold represents different combinations of the two dimensions, which means that the information about both features is not present in the population activity.

To form a complete basis of a stimulus space with d dimensions, there must be at least d neurons with receptive fields that are “linearly independent”. This means that no receptive field could be made out of a linear combination of the other receptive fields. The number of stimulus dimensions that a population code could possibly represent is less than or equal to the number of neurons. Typically a signal can be fully characterized by only a few features, and the number of neurons responding to a stimulus is usually much greater than the number of stimulus dimensions. If there are d stimulus dimensions, and N neurons in the population, and $N > d$, then the population code is a d -dimensional manifold in an N -dimensional space. To form a complete basis of a d -dimensional stimulus, there must be at least d neurons with receptive fields that are linearly independent. Not all of the neurons need to be linearly independent, however; with N neurons, some of the receptive fields can be linear combinations of others, as long as there is a subset of at least d linearly independent receptive fields.

If the stimulus dimensions (d) are low compared to the number of neural dimensions (N), the neurons can form an “overcomplete basis set” of the stimulus space. There are at least d neurons with linearly independent receptive fields in an overcomplete basis set, but the other

neurons can be linearly dependent. The overcomplete basis gives robustness to the population code because redundant information about the stimulus is duplicated across many neurons in the N -dimensional space. Even if one of the N neurons completely failed, the $N - 1$ neurons would still form a complete basis of the d -dimensional stimulus space, and no information about the stimulus would be lost. An interesting feature of this robustness is that none of the $N - 1$ receptive fields needs to be exactly the same as the neuron's that failed. The information carried by the failed neuron is also represented by some linear combination of the $N - 1$ neural responses. This property means that no two neurons are necessarily “redundant”—every neuron could have a unique response to a stimulus—but there is redundant information in the population vector. This redundancy helps the neuronal population maintain all of the information even in the presence of large amounts of noise.

The receptive field functions (f_j) play a critical role in the formation of a basis set. The receptive field shapes considered so far are Gaussian (aka, bell-shaped curves), because these types of functions match many biological receptive fields and because these types of functions are capable of forming a basis set. There are other types of receptive field functions (such as sigmoids or threshold-linear) that can also form basis sets, and it is possible to have different types of basis functions together in a single population code. The key property is that the basis set must be able to represent (or closely approximate) any type of non-linear function by a linear combination of the basis functions:

$$g(\theta) = \sum_j w_j f_j(\theta | \Theta_j) \quad (2.11)$$

where w_j are the coefficients and $g(\theta)$ is some non-linear function. This requirement of a basis set is similar to the decomposition of a signal by a Fourier transform, where a linear combination of sines and cosines can closely approximate any non-linear function. Some receptive field functions do not have this property. For instance, if all the receptive fields were linear, then any linear combination of them would also be linear. This means that non-linear functions cannot be fully represented, which means that a population code based entirely on linear receptive fields could not fully represent any signal.

Nervous systems can utilize the population coding framework beyond representing extrinsic features of the stimulus. The brain must represent intrinsic information about decisions and motor-commands as well. The nervous system of the leech has provided insights into how the same population of neurons can encode different modalities of information. The leech has a very simple nervous system, where the activity of a significant fraction of cells (> 100) can be

recorded simultaneously with voltage-sensitive dyes (VSDs). These recordings show that almost every cell in the leech nervous system is multi-functional: each neuron responds to many different stimuli and during several different behaviors. In effect, each neuron has a high-dimensional receptive field that is sensitive to many of the sensory, motor and intrinsic variables. Neurons in the leech encode the sensory environment using a population code. The local bend response is a reflex in which the leech bends its body away from a stimulus. Local bending requires information about the stimulus location to be propagated through the nervous system. The neurons involved in this behavior have overlapping receptive fields each with preferred touch locations, and encode the information about touch location just like the orientation information in the standard model (Lewis & Kristan, 1998; Fig. 2.1A). Neurons are also encoding information about the motor-patterns. For instance, virtually all the neurons that respond while the leech swims also respond while the leech crawls (Briggman & Kristan, 2006; Fig. 2.2A). Swimming and crawling are both oscillatory behaviors, and many of the neurons' voltages oscillate with the muscle contractions at different phases (Fig. 2.2B,C). There is no correlation between the phase at which a neuron oscillates during swimming and the phase at which it oscillates during crawling. This may be a consequence of the population forming a complete basis of all the motor patterns, because uncorrelated responses can arise from linearly independent receptive fields.

The swim and crawl motor patterns are represented as different population vectors in the high-dimensional response space of neural activity (Fig. 2.2D). The swim motor pattern is off in one direction, and the crawl motor pattern is off in another. To visualize the population code recorded through VSD imaging, the high-dimensional neural activity was reduced to three dimensions using principal components analysis. Each principal component corresponds to different combinations of activity patterns across all the neurons (Fig. 2.2E). The combinations of activity that are most correlated are chosen as the first three principal components, and the population vector can be visualized as a point in three dimensional space. By plotting the population vector over time, the two behaviors can be distinguished in the neural activity (Fig. 2.2F). In fact, the neural activity predicts the future behavioral state before the swim or crawl behavior actually begins. If the leech is stimulated at a particular site, it "chooses" to swim or crawl with equal probability. This decision can be decoded from the population activity before the motor-pattern can be seen. The information about the decision is encoded by the same population of neurons, and can be decoded as a particular direction in the high-dimensional space. By finding the point in time where the population vectors during swim and crawl separate, the decision direction can be decoded as the linear discriminant that maximally separates the two classes of behavior. The

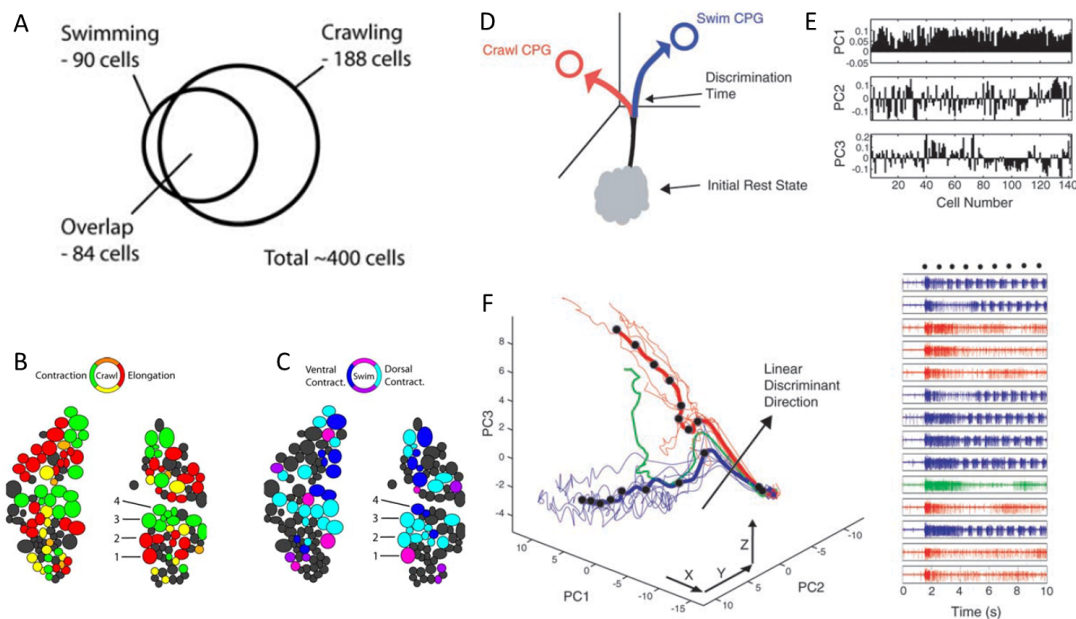


Figure 2.2: Populations encode multiple modalities of information. (A) Neurons that oscillate with the swimming and crawling behaviors were observed with voltage sensitive dyes. Virtually all neurons are multi-functional: those involved in swimming were also involved in crawling. (B and C) Image of the dorsal side of the leech ganglion where neuron somas are shown as circles. Neurons involved in crawling (B) and swimming (C) are colored based on the relative phase of their oscillations. (D) The swimming and crawling central pattern generators (CPG) are encoded as different population vectors in the high dimensional neural space. The population begins at an initial rest state, and then makes an excursion to one of the CPGs after a stimulus. (E) The first three principal components of neural responses. This allows the high-dimensional neural space to be reduced to 3-dimensions, which can be visualized. (F) Trajectories of the neural responses in PCA space during swimming and crawling behaviors (left). The black dots indicate equally spaced intervals in time. These behaviors can be predicted from the population response using linear discriminant analysis, which chooses the direction that maximally separates the population vectors during the two behaviors (arrow). The decision to swim or crawl can be determined before the motor patterns start (right): the traces show the firing of several motor neurons during swimming (blue) or crawling (red). In one outlier, the population vector appeared to head in the swim direction, but changed course and went to the crawl pattern (green). From: Briggman et al. 2005; Briggman & Kristan, 2006.

decision direction is actually a different direction than the information about either motor pattern (Briggman et al. 2005). This shows that the leech nervous system encodes all of the pertinent information as different directions along a manifold in high dimensional space, regardless of whether the information is sensory, motor, or intrinsic.

Transformations with population codes

Each area of the brain represents information using its own population code by combining information from the population codes of its inputs. Decoding a population code with an estimator is a strategy for an experimentalist to get an understanding of how information is represented by neuronal activity. In the brain, however, the stimulus dimensions are never directly extracted from a population code of neural activity. Instead, one population code is mapped directly onto another population code. This mapping is still effectively an estimation problem, because the direct mapping is equivalent to decoding the relevant information from the inputs, then encoding them into the local code.

Mapping one population code onto another requires non-linear transformations. An excellent example of this problem is the transformation of a retinotopic map onto a head-centered map (Pouget & Sejnowski, 1997; Fig. 2.3). As the eyes move around, neurons that receive information from the retina have receptive fields that move with the eyes. These neurons encode the visual scene in retinotopic coordinates. Another population of neurons encode the position of the eyes. With the information from these two populations, the brain can encode the visual input in reference to the head position, which would be useful if an object was moving towards your head and you wanted to move your head out of the way. The goal is to transform the population code in the retinotopic coordinates into a population code in head-centered coordinates by combining the retinotopic information with the eye position information. For simplification, we will consider a model of only the horizontal spatial dimension of the stimulus, x . In this model, each neuron in the retinotopic population has a receptive field that is based on where the stimulus falls on the retina. For these receptive fields, the relevant stimulus dimension is the position on the retina, r_x , and each neuron has a preferred retinal position (Fig. 2.3A). A second population encodes the position of the eyes, e_x (Fig. 2.3B). The eye position neurons were given sigmoidal receptive fields, which were chosen because of their biological relevance; each of these neurons has a different inflection point (similar to a preferred stimulus) in their receptive fields, and they too form a population code.

To combine the information from the retina (r_x) and the eye position (e_x), an intermediate

(“hidden”) layer was added (Fig. 2.3C) that forms a population code to represent the combined information about retinal activity and eye position. The key to the transformation was to make the population code of the hidden layer have a complete basis set of the information from both input populations. This projected the two dimensions of the stimulus (retinal position and eye position Fig. 2.3A,B) onto a 2-dimensional manifold in high-dimensional space (Fig. 2.3C). This transformation allowed the relevant dimension of head-centered coordinates (Fig. 2.3D) to be a simple linear combination of the population code formed in the hidden layer, which could easily be read out by the next neuronal population. The advantage of using a population code with a complete basis set is that any reference frame is just a linear combination away.

In fact, information about eye-position, head-position, body-position, joint-position, etc can all be combined into a single population code. With such a code, different motor system can read out the information about visual space with different reference frames all from the same population (Pouget & Sejnowski, 1997) with just a simple linear transformation. Further, the information about each reference frame can be decoded by these motor systems in parallel, because the population code is representing all of the necessary information simultaneously.

The key to forming this basis set was to find how the activity in the two population codes should be combined. It is standard procedure that synaptic connections can compute linear transformations of the inputs ($r_j = \sum_i w_{ij} r_i$), i.e. the synaptic strengths multiply the input firing rates to produce the downstream firing rates. But this modification alone was not sufficient to form a complete basis set of the data; gain control was also needed. Gain control is a scaling of the post-synaptic firing rate, and is seen in many different biological neural circuits. Gain control is a non-linear operation, which meets the necessary conditions for forming a basis set (other non-linearities are also possible, but gain-control is the most biologically plausible (Pouget & Sejnowski, 1997)). Figure 2.3E,F shows how neurons in the hidden layer are gain-modulated by the eye position, where the receptive field shape is maintained, but the amplitude is altered. These neurons are inspired by the activity of neurons in parietal cortex.

The mechanism of gain control in neurons is still controversial. Early studies postulated that shunting inhibition could be the mechanism of such multiplicative transformations (Blomfield, 1974; Vu & Krasne, 1992). Shunting inhibition alters the resistance of the neuron by opening inhibitory channels with a reversal potential that is the same as the resting potential. Changing the resistance of the neuron scales the voltage (and thus the firing rate) through Ohm’s law ($V = IR$). More recent studies have suggested that this mechanism is not sufficient when spiking dynamics are considered and that shunting inhibition acts additively (Holt & Koch,

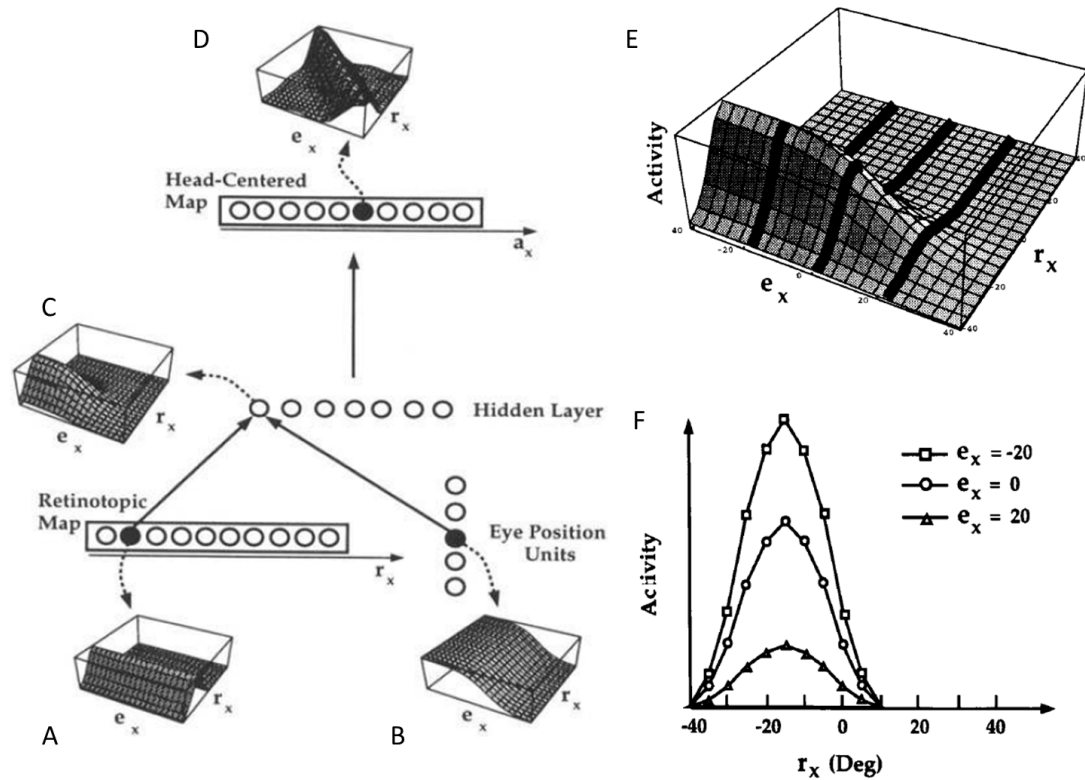


Figure 2.3: Transformation of population codes. (A) A population of neurons forms a retinotopic map of the visual world, with individual neurons having receptive fields based on where the stimulus falls on the retina (r_x). (B) Another population of neurons encodes the position of the eyes (e_x). (C) The hidden layer combines information from the retinotopic map and eye position units and forms a population code that is a complete basis set of the input information. (D) The head-centered reference frame can be decoded from the hidden layer with a linear transformation. (E and F) The receptive field of a neuron in the hidden layer (same as panel C) has a receptive field that is gain modulated by the eye-position units. This allows the hidden layer to form a complete basis set of the inputs. (E) Full 2-D receptive field. (F) Slices of the receptive field shown in black from (E). From: Pouget & Sejnowski, 1997

1997). Many alternatives use neural circuitry to implement gain control (Baca et al. 2008, Salinas & Abbott, 1996). However, shunting inhibition is receiving renewed interest as the mechanism for gain control, as models show that adding noise (Doiron et al. 2000; Chance et al., 2002) or using multi-compartmental neurons (Capaday & van Vreeswijk, 2006) can implement multiplicative operations in spiking neuron models.

Population coding in time

Many signals that the brain processes change in time. This temporal information can also be encoded by population codes. To encode temporal information, the receptive fields must also vary in time. An illustration of how temporal population coding is implemented by the nervous system comes from studies of olfactory processing in the insect antennal lobe (Geffen et al. 2009). While randomly puffing pulses of odor, several neurons were recorded as they responded to the fluctuating odor stimulus. Neurons respond with different activity patterns during such an odor stimulus; some neurons respond immediately when the odor turns on and then go silent, some neurons stay on while the odor is present, and still other neurons respond only when the odor turns off. The firing rates of these neurons were approximated by a “linear, non-linear model”, and closely matched the recorded rates. This model found a “temporal receptive field” that describes a preferred temporal fluctuation in the odors where the neuron fires maximally. Just like receptive fields considered before, temporal fluctuations in the stimulus that are close to the preferred temporal fluctuation also cause the neuron to fire. The temporal receptive field integrates the fluctuation pattern of the odor stimulus over time (this is the linear part of the model), and this sum is then passed through a non-linear function (which is a threshold-linear filter).

The temporal receptive fields found for all of the neurons can be well described by only two main components (Fig. 2.4A). These components are determined by calculating the principal components across time for all of the temporal responses of the neurons. Each neuron’s receptive field is effectively explained as a linear combination (i.e., a weighted sum) of the two components. Figure 2.4B shows the weights of the two components that best match the derived filters from the data, where the filters are spread out at every angle, mainly along the edge of the unit circle. Figure 2.4C shows the shapes of all of the filters sorted based on the angle. These neurons represent temporal variations of the odor signal as a two-dimensional manifold in the high-dimensional neural space. The neurons’ receptive fields are spread out almost uniformly, which creates an overcomplete basis set that is minimally redundant and maximally

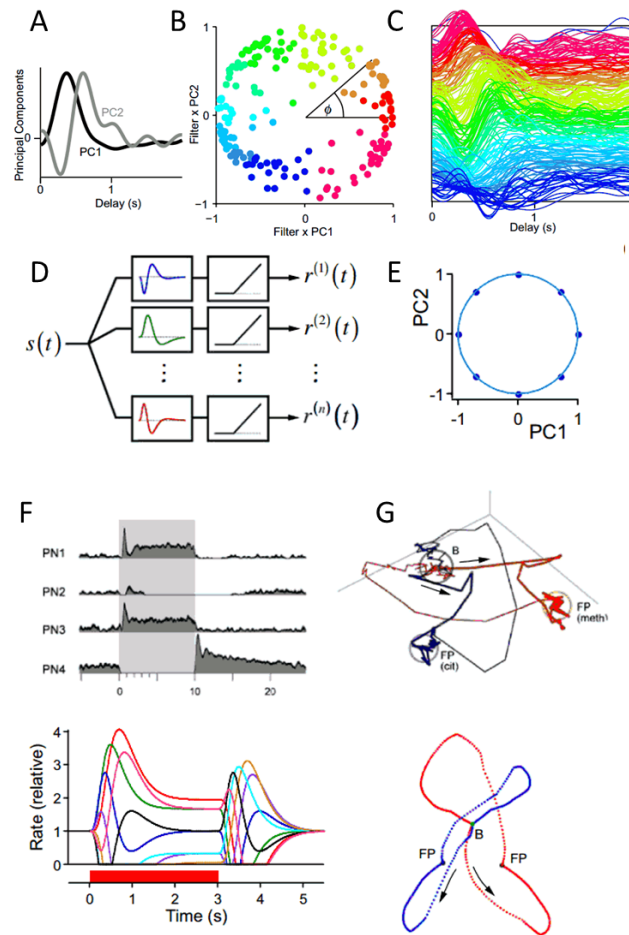


Figure 2.4: Population coding in time. (A) The first two principal components of the temporal receptive fields of many neurons in the insect antennal lobe. (B) The coefficients of the first two components for each receptive field. The receptive fields are spread out uniformly around the unit circle as different combinations of the two main components. The colors are based on the angle the coefficients are from the center. (C) All receptive fields plotted vs. the angle from (B). (D and E) A simple model of the insect antennal lobe population code. (D) The signal ($s(t)$) is encoded by the responses of a population of neurons ($r^{(n)}(t)$) with different temporal receptive fields and a threshold-linear filter. (E) The receptive fields were chosen based on 8 different combinations of the principal components derived in (A). (F) Responses of antennal lobe neurons to a step odorant (top), compared with the responses of the model (bottom) shows similar patterns of activation. (G) The population vector trajectories in PCA space of the recorded neurons from F (top) when two different odors were presented. The population vector encodes the odorant presence and recent fluctuation history. The model neurons show a similar pattern (bottom). From: Geffen et al. 2009.

robust. These types of receptive fields allow the neurons to encode a representation of some temporal history of the signal in the population code.

This encoding of temporal history in the population response can be seen in a simple model of these neurons (Fig. 2.4D). In the model, a population of neurons are randomly given one of the 8 receptive fields made out of different combinations of the two main components (Fig. 2.4E). When the population is given a step odor stimulus, the neurons respond differently based on their receptive fields. Just like the responses seen in the antennal lobe, the model neurons go through different fluctuations when the stimulus comes on, and when it goes off (Fig. 2.4F). The stimulus information is encoded as a point in the high-dimensional activity space, which in this case moves in time (Fig. 2.4G). Before the onset of the odor step, the population is at some equilibrium (denoted as B). After the step, the population vector makes a rapid excursion from baseline and eventually reaches a fixed point (FP), which indicates that the firing rates have reached a steady state. When the odor is turned off, the population vector leaves the fixed point, makes another excursion and settles back to the original baseline. If a different odor is given, then the neurons progress through a similar pattern, but along a completely different trajectory.

Each point along these trajectories encodes the recent history of the odor fluctuations. Since the odor filters are responsive only to a window of a few seconds, the information in the population code contains only a few seconds of history. At the fixed point after the odor turns on, the population vector represents the presence of a particular odor and the fact that the odor has not changed in the last few seconds. During the excursion, the population vector continuously encodes the odor presence and the sudden change in odor that occurred when the odor step turned on. Different odor fluctuations would move the population vector along different trajectories in the high-dimensional space, and the recent history of the stimulus fluctuations can be decoded by the immediate position of the population vector.

Alterations in the neural responses over longer time-scales would be required to encode a longer history of temporal information. Other types of temporal fluctuations seen in nervous systems may be serving this purpose. A full range of different time scales could be implemented by various neural mechanisms. For instance, adaptation or neuro-modulators can encode information on the order of seconds, short-term plasticity could be involved in representing information on the order of tens of seconds, and long-term plasticity could encode even larger temporal scales.

Conclusion

Population coding is a powerful framework for neural computation. By spreading information across a large population of neurons, the brain solves many of the problems posed by biology. Noise is mitigated by coding information in the joint activities of many neurons, instead of individual neurons that can easily fail. Further robustness is achieved by coding with an overcomplete basis set, which contains redundant information that can compensate for information lost due to noise. An overcomplete basis set enables a population of neurons to represent all of the features of the stimulus, the intrinsic states, and/or motor activity as well as the recent history of these states simultaneously. This facilitates parallel processing from downstream systems, since they can decode different reference frames from the same population. Population codes can be arbitrarily transformed into other representations and combined with information from other population codes in a biologically plausible manner. Population coding is a promising theoretical basis for universal computation in the nervous system, a necessary foundation for generating the wide variety of perceptual, cognitive, and behavioral states seen in the brain.

Acknowledgements

This work was published in Frady, E.P., Kristan, W.B. (2013). Computation with Population Codes. In: Jaeger D., Jung R. (Ed.) *Encyclopedia of Computational Neuroscience*: Springer Reference (www.springerreference.com). Springer-Verlag Berlin Heidelberg, 2013. DOI: 10.1007/SpringerReference_348400

References

- Baca, S.M., Marin-Burgin, A., Wagenaar, D.A., Kristan Jr., W.B. (2008) Widespread inhibition proportional to excitation controls the gain of a leech behavioral circuit. *Neuron* 57, 276-289.
- Blomfield, S. (1974). Arithmetical operations performed by nerve cells. *Brain Res.* 69, 115-124.
- Briggman, K.L., Abarbanel, H.D.I., Kristan Jr., W.B. (2005). Optical Imaging of Neuronal Populations During Decision-Making. *Science* 307: 896-901.
- Briggman, K.L., Kristan Jr., W.B. (2006). Imaging Dedicated and Multifunctional Neural Circuits Generating Distinct Behaviors. *Journal of Neuroscience* 26(42):10925-10933.
- Capaday, C., Van Vresswijk, C. (2006). Direct control of firing rate gain by dendritic shunting inhibition. *J. Int. Neurosci.* 5(2), 199-122.

- Chance, F.S., Abbott, L.F., Reyes, A.D. (2002). Gain modulation from background synaptic input. *Neuron* 35, 773-782.
- Deneve, S., Latham, P.E., Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neuroscience* 2(8): 740-745.
- Doiron, B., Longtin, A., Berman, N., Maler, L. (2000). Subtractive and divisive inhibition: effect of voltage-dependent inhibitory conductances and noise. *Neural Comp.* 13, 227-248.
- Geffen, M.N., Broome, B.M., Laurent, G., Meister, M. (2009) Neural encoding of rapidly fluctuating odors. *Neuron* 61, 570-586.
- Holt, G.R., Koch, C. (1997). Shunting inhibition does not have a divisive effect on firing rates. *Neural Comp.* 9, 1001-1013.
- Lewis, J.E., Kristan Jr., W.B. (1998). A neuronal network for computing population vectors in the leech. *Nature* 391: 76-79.
- Pouget, A., Sejnowski, T.J. (1997). Spatial Transformations in the Parietal Cortex Using Basis Functions. *Journal of Cognitive Neuroscience* 9(2): 222-237.
- Pouget, A., Zhang, K., Deneve, S., Latham, P.E. (1998). Statistically efficient estimation using population coding. *Neural Computation* 10, 373-401.
- Sanger, T.D. (2003). Neural population codes. *Current Opinion in Neurobiology* 13:238-249.
- Salinas, E., Abbott, L.F. (1996). A model of multiplicative neural responses in parietal cortex. *Proc. Natl. Acad. Sci.* 93, 11956-11961.
- Vu, E.T., Krasne, F.B. (1992). Evidence for a computational distinction between proximal and distal neuronal inhibition. *Science* 255(5052), 1710-1712.

3 Population Coding and Gain Control

Perspectives of Gain Control

Additive and multiplicative interactions have been observed widely throughout various nervous systems. Since Hubel & Wiesel (1962) described orientation selectivity as the summation of presynaptic activity from thalamic neurons with aligned receptive fields, additive models of neuronal interactions have served to describe the firing properties of neurons, typically both the excitatory and inhibitory components (Movshon et al. 1978; Bonin et al. 2005). The mapping of receptive fields based on methods such as reverse correlation (Chichilnisky, 2001) relies on the premise that presynaptic activity combines additively (a weighted sum) to produce the receptive field shape. Multiplicative interactions have also been described in many nervous systems, including the vertebrate retina (Shapley & Victor, 1979; Solomon et al. 2006), LGN (Bonin et al. 2005; 2006) and cortex (Heeger, 1992; Carandini et al. 1997; Simoncelli & Heeger, 1998; Zoccolan et al. 2005; Andersen & Mountcastle, 1983; Andersen et al. 1985) as well as in invertebrates (Baca et al. 2008; Olsen & Wilson, 2008; Olsen et al. 2010; Gabbiani et al. 1999). Divisive normalization and gain control are two types of multiplicative interactions that are involved in many sensory (Treue & Trujillo, 1999; Shenoy et al. 1999; Ohshiro et al. 2011), cognitive (Reynolds & Heeger, 2009; Salinas & Abbott, 1997; Beck et al. 2008; Louie & Glimcher, 2010) and motor functions (Baca et al. 2008; Prochazka, 1989). Together, addition and multiplication are seen as canonical neural computations (Carandini & Heeger, 2012) observed across many nervous systems.

A large body of theoretical research has produced a strong basis for population coding as an underlying principle of neural computation (Sanger, 2003; Averbach et al., 2006). Information processing using population codes is based on the relative firing rates among a population of neurons, and requires both additive and multiplicative modulation of neuronal responses (Zipser & Andersen, 1988; Salinas & Thier, 2000). Modulation of gain is an important component of

the information-theoretic principles of population coding (Pouget & Sejnowski, 1997; Deneve et al. 1999; Pouget et al. 2000; 2003; Ringach, 2010), because gain control does not alter the information stored within a population code by preserving the shape of the receptive fields (Salinas & Sejnowski, 2001). Although population coding has been linked with rate models of neuronal input-output functions, its link to spiking biophysical properties has yet to be resolved. This is because the mechanism of multiplication in neurons with spiking dynamics is controversial (Gabbiani & Knopfel, 1994; Holt & Koch, 1997; Chance et al. 2002; Mitchell & Silver, 2003; Shu et al., 2003; Graham & Schramm, 2009; Fernandez & White, 2010). Two main classes of mechanisms, both involving GABAergic inhibitory synapses, have been proposed to explain how neurons could implement multiplication. One class posits that synaptic mechanisms directly implement multiplication through shunting inhibition (Blomfield, 1974; Vu & Krasne, 1992). Shunting alters the conductance of a neuron by opening inhibitory channels, causing the multiplication of neural activity through Ohm's law ($I = Vg$). A second class of hypotheses suggests that neural circuitry causes multiplication without a direct nonlinear synaptic mechanism (Salinas & Abbott, 1996; Hahnloser et al. 2000; Capaday, 2002; Baca et al. 2008; Ayaz & Chance, 2009).

To test these two hypotheses of gain control, we set out to unify physiological results with computational principles of population coding in the leech local bend reflex. The neural activity underlying local bending is well characterized (Kristan et al., 1982; Lockery & Kristan, 1990a,b) and is described by a model that encodes information using a population code (Lewis & Kristan, 1998a,b,c; Lewis, 1999). Local bending is a withdrawal response from a mechanosensory stimulus, which the leech achieves by bending its body away from the stimulated location (Kristan, 1982; Lockery & Kristan, 1990a; Fig. 3.1A, B). Touch information that drives local bending is encoded by the firing of 4 pressure-sensitive sensory neurons (P cells) that have receptive fields centered in each quadrant around the animal (Lockery et al. 1989; Lewis & Kristan 1998b; Lewis, 1999; Fig. 3.1C). All four P cells connect to a population of local bend interneurons (LBI cells), with the strength of the connection depending upon the difference between the preferred touch location of the P cell (i.e., the most sensitive location in the receptive field, Fig. 3.1D) and the local bend interneuron. This wiring encodes the touch location as a population vector in the LBI cells (Lewis & Kristan, 1998a). The local bend interneurons activate excitatory and inhibitory motor neurons to contract and relax the correct longitudinal muscles, creating a 3-layered, feed-forward neural network that “computes” the reflex (Lewis, 1999; Fig. 3.1E).

The current study begins by combining the model of population coding (Lewis & Kris-

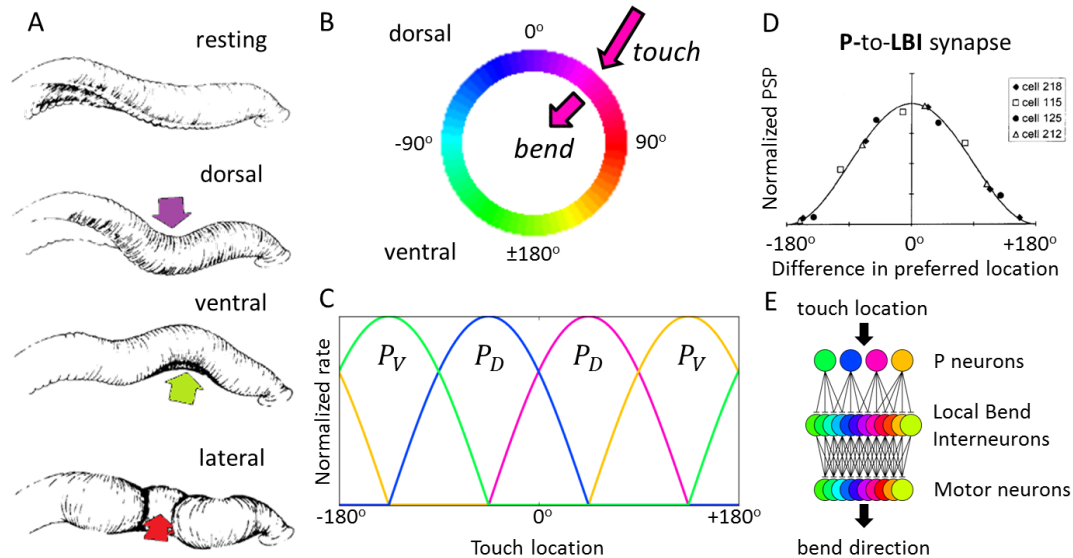


Figure 3.1: Local bend is a directional response that employs a population code. (A) Touching the leech in different locations around the circumference of the body results in a bend away from the stimulated site. (B) In cross-section, the leech's body is idealized as a circle. A touch applied at a particular location causes longitudinal muscles to contract centered at that location. (C) Four mechanosensory neurons responsive to pressure (P cells) have overlapping receptive fields that cover the full circle of the body wall. These receptive fields form rectified-cosine functions, with each P cell firing maximally when a touch is at the cell's preferred location. The preferred locations are at $\pm 45^\circ$ and $\pm 135^\circ$, with the dorsal midline defined as 0° . The receptive fields are colored based on the preferred location and the color code in panel B. (D) All four P cells make excitatory connections to a set of interneurons. The strength of the connection is based on the difference between preferred location of the P cell and that of the interneuron. (E) The behavior is computed in a 3-layer neural network where the touch location information is transmitted to subsequent layers based on a population code. The interneurons connect to longitudinal motor neurons, which produce the bend. (Adapted from Lockery & Kristan, 1990a; Lewis, 1999).

tan, 1998a) with a proposal that generalized inhibition controls the gain of the neural circuit responsible for local bending (Baca et al. 2008). First, we developed a simplified rate-coded model to handle additive and multiplicative inhibition in a mathematically tractable way. This simple model highlights a major problem in using a circuit mechanism for computing division compared with a nonlinear synaptic mechanism: a complex and finely balanced circuit is required. Second, we derived a two compartment spiking model, where inhibitory shunting channels are separated from the spike initiation zone, which implements multiplication through shunting inhibition in a biologically plausible manner. Using this two-compartment model, we then show that a spiking, conductance-based model can robustly match the essential features of the local bend reflex and can compute with a population code. Finally, we provide new electrophysiological data indicating that two different inhibitory neurons (cells 116 and DI-1) modulate the input-output function of a local bend motor neuron (DE-3) in ways predicted by the model.

Rate Models of the Local Bend Population Code

We built upon a previous model of local bending (Baca et al., 2008) that explored how GABAergic inhibition influences this behavior. The magnitude of the local bend reflex increases with the intensity of the stimulus, and blocking GABA transmission scales the response magnitude proportionally across stimulus intensities. To explain this, the previous study postulated a simple rate-coding model (Fig. 3.2A) with a single (or single class of) additive GABAergic inhibitory neuron (G_+) that increases its activity in proportion to the total activity of the four mechanosensory P cells:

$$G_+ = w * P \quad (3.1)$$

where G_+ is the firing rate of G_+ , the additive GABA neuron, P is the firing rate of the P cell, and w is the synaptic weight of P onto G_+ . In turn, G_+ inhibits the local bend interneurons (LBI) with a synaptic connection that adds linearly to the excitatory inputs from the P cells. The additive influence of excitation and inhibition is described as:

$$LBI = \nu * P + \gamma * G_+ \quad (3.2)$$

Where LBI is the firing rate of the interneuron, ν is the weight from P onto LBI , and γ is the (negative) weight from G_+ onto LBI . The interneuron activity (LBI) is scaled by GABAergic inhibition, since the G_+ cell activity grows in proportion to the sensory inputs. This

relationship can be seen by substituting Eq. 3.1 into Eq. 3.2:

$$\begin{aligned} LBI &= v * P + \gamma * w * P \\ LBI &= (v + \gamma * w) * P \end{aligned} \tag{3.3}$$

Figure 3.2B shows this scaling of *LBI* when the GABA inhibition is blocked by 70% (modeled by reducing γ by 70%, Eq. 2). For simplicity, we have removed the sigmoidal nonlinearities of the original model. The linear model is qualitatively similar to the sigmoidal model and, in the Discussion, we elaborate on how this model could be generalized to include nonlinearities. We do not know whether blocking GABAergic transmission through BMI (a GABA competitive antagonist) blocks transmission in all types of leech GABA synapses. Blocking GABA in our models is essentially a stand-in for manipulation of the inhibitory neurons, and highlights the regulatory effects of these neurons on the circuit.

This model of local bending (Baca et al., 2008) considered the scaling of the neural responses across the physiological range of intensities of the stimulus, but did not consider the spatial dimension, i.e. the responses across receptive fields. To test the full role of inhibition in local bending, we considered both the intensity and the spatial dimensions by combining the “general inhibition model” (Fig. 3.2A) with the “population coding model” that encodes the touch information across a population of local bend interneurons (Lewis & Kristan 1998a; Fig. 3.1E).

The most straight-forward combination of these two models is to connect G_+ uniformly to all the *LBI* cells (Fig. 3.2D). Using such a model, we simulated a touch at 0° (Fig. 3.2C) at increasing intensities by increasing the *P* cell firing rates according to the population coding model (Fig. 3.1C). This stimulus activated just the two dorsal *P* cells (Fig. 3.1C), which fired at equal frequencies for each intensity. Figure 3.2E shows the firing rate of two of the interneurons, *LBI* (0°) and *LBI* (90°) (with receptive field peaks centered at 0° and 90° , respectively). In the intensity dimension, both neurons’ firing rates are scaled up by a constant in response to blocking GABA transmission (Fig. 3.2E). This scaling was seen regardless of the *LBI* cell’s preferred touch location. However, these cells’ firing rates were scaled up by different amounts: *LBI*(90°)’s activity increased almost four-fold when GABA was blocked, whereas *LBI*(0°)’s activity was increased by less than two-fold (as seen in Eq. 3.3, because the amount the *LBI* is scaled depends on v). This different scaling propagated across the receptive field dimension as an additive effect on the interneuron population: the responses increased by a constant amount

across the spatial dimension (Fig. 3.2F). Hence, manipulating this additive inhibition alters the receptive field shape of all the *LBI* cells. This is contradictory with gain control, because reshaping receptive fields alters the information that the neurons represent in the population code (Salinas & Sejnowski, 2001).

Other circuit solutions have been suggested to implement gain control that preserve the shape of the receptive fields (e.g. Salinas & Abbott 1996; Hahnloser et al. 2000), but they do not fit with the known circuitry of the local bend reflex. We did find a way to obtain gain control with no alteration of the receptive fields using an additive synaptic mechanism that fit within the known constraints of the circuitry: we set the G_+ -to-*LBI* synaptic weights proportional to the *P*-to-*LBI* synaptic weights. Because each *P* cell has a different weight distribution onto each local bend interneuron, however, we needed to have at least four G_+ cells, each of which had the same distribution of connections onto the *LBI* cells as one of the four *P* cells (Fig. 3.2G). This configuration produced gain control that was consistent across both the intensity and spatial dimensions (Fig. 3.2H-I). However, this solution makes the system quite complex: it took four cells rather than one, and it required a precise matching of synaptic weights of the *P* and G_+ cells onto each of the *LBI* cells. Further, this solution would require a G_+ neuron for every excitatory neuron that synapses onto the *LBI* neurons, which could be quite large when all the other types of sensory input are considered.

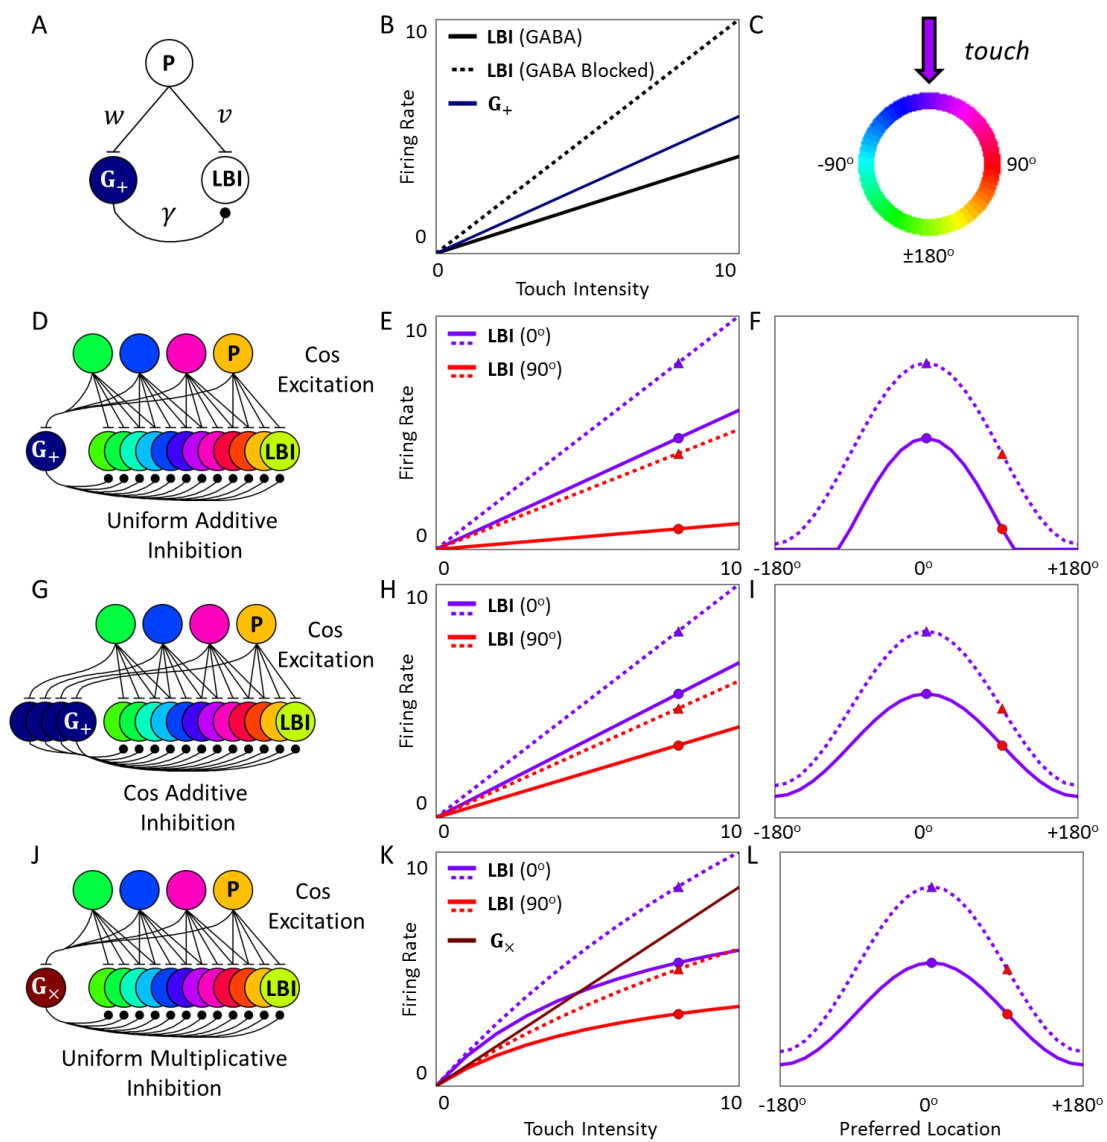
We found a way to avoid both the spiraling complexity and fine balancing of synaptic weights by using a multiplicative synapse to control gain across the receptive field. This multiplicative synapse directly scaled the activity of the interneurons, which altered the basic *LBI* cell equation (Eq. 3.2) by dividing the *P*-to-*LBI* excitation ($v * P$) by the synaptic input from G_\times ($\gamma * G_\times$):

$$LBI = \frac{v * P}{\gamma * G_\times + 1} \quad (3.4)$$

(1 was added to the denominator to assure that when there was no inhibition (i.e. $\gamma * G_\times = 0$), *LBI* would equal $v * P$). In this model a single G_\times cell received input from all four *P* cells and uniformly inhibited all *LBI* cells (Fig. 3.2J). The multiplicative inhibition maintained constant relative firing rates across receptive fields (Fig. 3.2L), but because the G_\times cell's activity increased in proportion to the *P* cell activity, the *LBI* cell response did not scale evenly across stimulus intensities (Fig. 3.2K).

These modeling results show that multiplicative and additive inhibition affect the circuit in complementary ways. Multiplicative inhibition allows the population activity to be scaled

Figure 3.2: Models of additive vs. multiplicative synapses for controlling gain. (A) General inhibition model. This model produces a change in gain across intensity dimension by blocking additive GABAergic inhibitory transmission (from Baca et al., 2008). (B) Results from a linearized version of the general inhibition model. Both the G_+ cell and the LBI cell activities increase linearly with the touch intensity (dark blue line and black line, respectively). (C) For the combined model, a touch was given at 0° at different intensities. This stimulus activated the two dorsal P cells equally, and does not activate the ventral P cells. (D-F) Combined model with uniform additive inhibition on the interneuron population. (D) Circuit schematic. Each P cell excites a single G_+ cell, which additively inhibits the local bend interneurons uniformly. The P cells excite the local bend interneurons as described in Fig. 3.1D (Cos Excitation). (E) The firing rates of the interneurons with receptive fields at 0° ($LBI(0^\circ)$) and 90° ($LBI(90^\circ)$) are plotted against touch intensity. (F) The firing rates of the interneuron population plotted against their preferred touch locations. The circular and triangular markers indicate identical data points used in panels E and F. (G-I) Same measures as in panels D-F, but with cosine distributed G_+ -to- LBI synapses that mirror the P cell synapses. With this circuitry, altering additive GABA inhibition scaled the local bend interneuron activity in the intensity dimension (H) without affecting the shape of the receptive fields (I). (J-L) Same measures as in panels D-F, but with multiplicative inhibitory neurons that uniformly inhibit the local bend interneurons. A multiplicative mechanism does not alter the shape of the receptive fields of the local bend interneurons (panel L), but with this circuitry the gain in the intensity dimension is not constant (K).



without affecting the receptive field shapes, which is an essential property for processing information in a population code (Salinas & Sejnowski, 2001). Additive inhibition combines with excitation to modulate the receptive field shape and plays a role in influencing the responses across the intensity dimension. Experimental evidence has revealed that both multiplicative-like and additive-like inhibition do indeed exist in other brain structures, such as cortex (Wilson et al. 2012; Carandini & Heeger, 1994).

We addressed these theoretical results by developing a simple spiking model that can explain a simple mechanism for multiplication, shunting inhibition. Shunting inhibition as a mechanism for implementing multiplication of neural firing rates is an old idea. However, theoretical modeling work (Holt & Koch, 1995) suggested that shunting inhibition was invalid as a mechanism. This was because spiking models illustrated that the axon acts as a kind of voltage-clamp due to the after-hyperpolarization, and modeling results showed that shunting inhibition caused an additive shift in firing rate. Several alternative solutions emerged, but were complex and required mechanistic assumptions about multiplication that were not valid in all systems. The elaborate architecture of leech neurons and theoretical work of sub-cellular processing (Poirazi et al. 2004; Häusser and London, 2005; Larkum et al. 2009) inspired a multi-compartmental architecture. The next section is a manuscript in preparation for publication, and in it we present a theoretical basis for shunting inhibition and experimental evidence for differential implementation of additive and multiplicative modulations in the leech nervous system. These simple experiments were inspired by both the theoretical importance of multiplication and the need to precisely control and modulate gain, as well as a large body of literature recognizing multiplication and addition as canonical operations across a wide-variety of nervous systems (Carandini & Heeger, 2012).

Shunting and inhibitory currents co-regulate the input-output function of an identified leech neuron

Abstract

The interplay of linear and non-linear modulation of neuronal responses has a key role in neural computation and population coding. Multiplication is a non-linear operation that is prevalent throughout a wide variety of nervous systems, but the biophysical mechanism is still controversial. We characterize two inhibitory neurons in the leech, DI-1 and 116, which influence a post-synaptic cell, DE-3, additively (linearly) and multiplicatively, respectively. We

hypothesize that the difference is a result of the location of the synaptic connections on DE-3's arbor, and show that 116 contacts DE-3 in the necessary location to produce non-linear multiplicative modulation via shunting inhibition. We derive a two-compartment spiking model that implements multiplication using shunting inhibition without the need for noise or complex circuit mechanisms. This model restores shunting inhibition as a plausible mechanism for multiplicative operations in a range of nervous systems.

Introduction

Neurons can form targeted synaptic connections to specific parts of the dendritic tree (Spruston, 2008), but the functional role of such localization is not fully understood. Theoretical studies have suggested a computational distinction between proximal and distal inhibitory synaptic connections (Vu & Krasne, 1992). Distal inhibitory inputs can combine linearly with excitatory inputs, whereas proximal connections can act in a non-linear fashion through shunting inhibition. The need for non-linear modulation has been recognized in a diverse set of computational theories (Pouget & Sejnowski, 1997; Deneve et al. 1999; Pouget et al. 2000; 2003; Ma et al. 2006; Ringach, 2010), and non-linear phenomenon such as contrast-scaling, gain control, and normalization are seen in a variety of nervous systems (Prochazka, 1989; Carandini et al. 1997; Zoccolan et al. 2005; Solomon et al. 2006; Baca et al. 2008; Louie & Glimcher, 2010; Olsen et al. 2010; Ohshiro et al. 2011).

The leech nervous system appears to “compute” the local bend reflex using a population code. Theoretical insights have pointed to the importance of both linear and non-linear modulation for utilizing the full computational power of a population code (Lewicki & Sejnowski, 2000; Salinas & Thier, 2000; Beck et al. 2008). We therefore investigated inhibitory neurons in the leech to assess if they linearly or non-linearly affected DE-3, a motor neuron involved in the local bend reflex. DE-3 receives input from a known inhibitory motor neuron, cell DI-1 (Lytton & Kristan, 1989), as well as cell 116, a previously uncharacterized inhibitory interneuron. We show that 116 influences DE-3 with a non-linear multiplicative component, while DI-1's influence is purely additive. We provide evidence that 116 contacts DE-3 in the necessary location to perform multiplicative modulation through a shunting mechanism.

There is still much debate as to the mechanisms of such non-linear responses. Conductance based synaptic connections were thought to enable multiplication of neural responses through shunting inhibition. However, several theoretical (Holt & Koch, 1997) and experimental (Chance et al. 2002) studies suggest that shunting inhibition does not act in a multiplicative man-

ner. This is thought to be because the dynamics of the action potential interfere with the driving forces of the shunting channels. We created a spiking model of DE-3 that can implement both additive and multiplicative operations by separating the spiking mechanism from the incoming synaptic input. We derived a robust approximation to the spiking model, which reveals that the shunting conductance must be within the same order of magnitude of the coupling conductance in order for shunting inhibition to have a multiplicative effect.

Materials and Methods

Electrophysiology. An isolated leech ganglion (leeches are hermaphrodites) was dissected and desheathed with DP nerves cleaned using standard procedures (Briggman & Kristan, 2006). We used suction electrodes to record extracellular action-potentials from the DP nerve and sharp microelectrodes to record intracellularly and to inject current. Cell DE-3 was identified by matching spikes one-for-one in the microelectrode recordings and the DP nerve recordings. Cell 116 and DI-1 have somata with a stereotypical location and size, and can be identified from their neighbors through their action-potential shape and the fact that they inhibit cell DE-3. Confirming that they were inhibitory required only a brief depolarization, which would reduce the firing of cell DE-3 in the DP nerve recording. Occasionally cells were slightly hyperpolarized to hold the cell and mitigate damage caused by impaling the electrode. Although the inhibitory neurons were typically not spiking when no current was being injected, a hyperpolarizing pulse was always used just to make sure that there were no inhibitory inputs onto cell DE-3 in the control condition.

The input-output (IO) function of DE-3 was measured by injecting several current pulses of different strengths into DE-3 and counting the resulting number of action-potentials. The pulses were presented for 1 second with 4-5 seconds between each pulse. This inter-pulse period was to ensure that DE-3 returned to its baseline firing rate, as DE-3 adapts to the current input. Simultaneously with the DE-3 current pulses, the inhibitory neurons 116 or DI-1 were injected with a 1 second depolarizing or hyperpolarizing pulse to assess their effects on DE-3's IO function. The pulses were presented in a random order to prevent any systematic biases in the measurement of the IO function. The firing rates of DE-3 were computed as the number of spikes during the 1-second current pulse.

Linear fits were made to DE-3's IO function ignoring firing rates above 80 Hz and below 1 Hz to avoid artifacts from saturation and non-negative firing rates. The measurement of the IO functions were made 3-5 times in a single animal. Eight different animals were examined for

116 experiments, and six different animals were studied for DI-1 experiments (two animals were overlapping – both 116 and DI-1 experiments were done). The slopes of the linear fits were then normalized by the slope of the linear fit in the hyperpolarizing case. Each trace was individually averaged, but averaging across animals first had little difference in the results. Because DE-3 adapts, we also analyzed the IO functions by only considering the first and second half of the 1-second period. However, this did not produce any qualitative differences in the results.

Neuronal morphology. We verified the identities of cells DE-3 and 116 with electrophysiological recordings, then determined their morphologies. For the examples shown in Fig. 2, we injected 5% AlexaFluor-568 Dextran 10000 MW (Life Technologies-D22912) into cell DE-3, and injected a mixture of 5% Biotin-Dextran 10000 MW and 2.5% AlexaFluor 488 Dextran 10000 MW (both Life Technologies-D1956, D22910) into cell 116. The markers were allowed to diffuse for 1 hour, and the tissues were fixed in 2% paraformaldehyde in phosphate buffer, pH 7.4 overnight at 4C. We then washed the tissue with PBS and amplified the biotin signal by overnight incubation in 2ug/mL Streptavidin-AlexaFluor 488 conjugate (Life Technologies -S11223) in 0.3% TritonX-100 in PBS (PBX). The tissue was serially dehydrated in ethanol, cleared in methyl salicylate, and mounted in Gurr DePex mounting medium (Electron Microscopy Sciences). A Z-stack was made with a Leica SP5 confocal microscope. The image shown in Figure 2B is a maximum projection of the intensity values across the depth of the image stack.

Modeling. Spiking simulations were implemented with the Brian Simulator (Goodman & Brette, 2008). The dynamics of the two-compartment spiking model are described by:

$$C_{Sh} \frac{dV_{Sh}}{dt} = g_l(E_l - V_{Sh}) + g_\gamma(E_\gamma - V_{Sh}) + g_C(V_{Ax} - V_{Sh}) + I_D \quad (3.5)$$

$$C_{Ax} \frac{dV_{Ax}}{dt} = g_l(E_l - V_{Ax}) + g_C(V_{Sh} - V_{Ax}) \quad (3.6)$$

Where C_{Sh} and C_{Ax} are the capacitances of the ShC and axon, E_l is the resting potential (defined as 0), E_γ is the shunting reversal potential, g_l is the leak conductance, g_C is the coupling conductance, and V_{Sh} and V_{Ax} are the voltages of the ShC and axon. When the voltage of the axon (V_{Ax}) goes above the threshold voltage (V_θ), then the neuron fires an action-potential and V_{Ax} is set to the reset potential (V_R). I_D indicates the currents coming into the ShC from the dendrites. Figure 3.5A shows a schematic of the Equations 3.5 and 3.6, and the values of the parameters are shown in Table 1. The code for the simulations can be found online.

Derivation of Linear Approximation. The two-compartment spiking model cannot be fully analytically characterized, but the spiking rate can be approximated with a linear function by making some simplifying assumptions. We will break the model down by analyzing the axon and shunt compartments separately. The instantaneous firing rate (IFR) of the axon is the inverse of the spiking period (T_{spike}), the time that the neuron takes to reach threshold from the reset potential:

$$IFR = \frac{1}{T_{spike}} = \frac{1}{C_{Ax}/g_l \ln\left(\frac{I_{Sh}-g_l V_R}{I_{Sh}-g_l V_\theta}\right)} \quad (3.7)$$

The IFR converges to a straight line as I_{Sh} increases (Fig. 3.5C). We used this convergence to linearly approximate the IFR-current relationship:

$$I\tilde{F}R = mI_{Sh} + b \quad (3.8)$$

By taking the derivative of the IFR and setting I_{Sh} to a large value, the slope of this line, m , is:

$$m = \lim_{I_{Sh} \rightarrow \infty} \frac{dIFR}{dI_{Sh}} = \lim_{I_{Sh} \rightarrow \infty} \frac{g_l^2 (V_\theta - V_R)}{C_{Ax} (I_{Sh} - g_l V_R) (I_{Sh} - g_l V_\theta) \ln^2\left(\frac{I_{Sh}-g_l V_R}{I_{Sh}-g_l V_\theta}\right)} \rightarrow \frac{1}{C_{Ax} (V_\theta - V_R)} \quad (3.9)$$

By substituting this value of m into Equation 3.8 and setting I_{Sh} to a large value, we can solve for b :

$$b = \lim_{I_{Sh} \rightarrow \infty} IFR - \frac{1}{C_{Ax} (V_\theta - V_R)} I_{Sh} \rightarrow -\frac{g_l (V_\theta + V_R)}{2C_{Ax} (V_\theta - V_R)} \quad (3.10)$$

Next we will consider the shunt-compartment. The ‘‘steady-state’’ currents of this compartment are described by these equations:

$$I_D = g_l (V_{Sh}^{EQ} - E_l) + g_C (V_{Sh}^{EQ} - V_{Ax}) + g_\gamma (V_{Sh}^{EQ} - E_\gamma) \quad (3.11)$$

$$I_{Sh} = g_C (V_{Sh}^{EQ} - V_{Ax}) \quad (3.12)$$

where V_{Sh}^{EQ} is the steady-state voltage in the shunt compartment to a constant injection of I_D .

The nature of the axon spiking allowed us to ignore any feedback effects of the axon

on the ShC. Because the axon membrane potential varied only between the threshold and the reset potentials (most IF models do not include the voltage changes due to the action potential), the spiking mechanism effectively acts as a kind of voltage clamp on a long time scale (Koch et al., 1995). In the simulation, we made the threshold and reset potentials small (about 2mV) to ensure that the axon voltage fluctuations remained close to the resting potential and therefore had minimal effects on the ShC. This manipulation ensured that the amount of current that flowed into the axon from the ShC (I_{Sh}) was predominantly dependent on the ShC voltage (Eq. 3.12). Even if we had incorporated currents from the action potential (Capaday & Van Vreeswijk, 2006) or had allowed larger axon voltage fluctuations, we could compensate for these changes by altering the parameters derived from the linear approximation and still achieve the same qualitative results.

Since V_{Ax} is close to 0 on average in the simulation, we approximated can assume it is 0 to derive the linear approximation. The differences caused by the actual value of V_{Ax} in the simulation could be countered by altering the parameters of the linear approximation, but they are sufficiently small that our approximation of $V_{Ax} = 0$ is valid. This is especially true when the axon is in the “spiking-regime”, and essentially V_{Ax} is oscillating around 0 between the threshold and reset potentials. For convenience, we have defined the resting potential, E_l , to be 0, and thus the shunting reversal potential, E_γ , was also set to 0. With these assumptions we derived a simple formula for the gain, Γ , based on the three conductances from Equation 3.11 and 3.12:

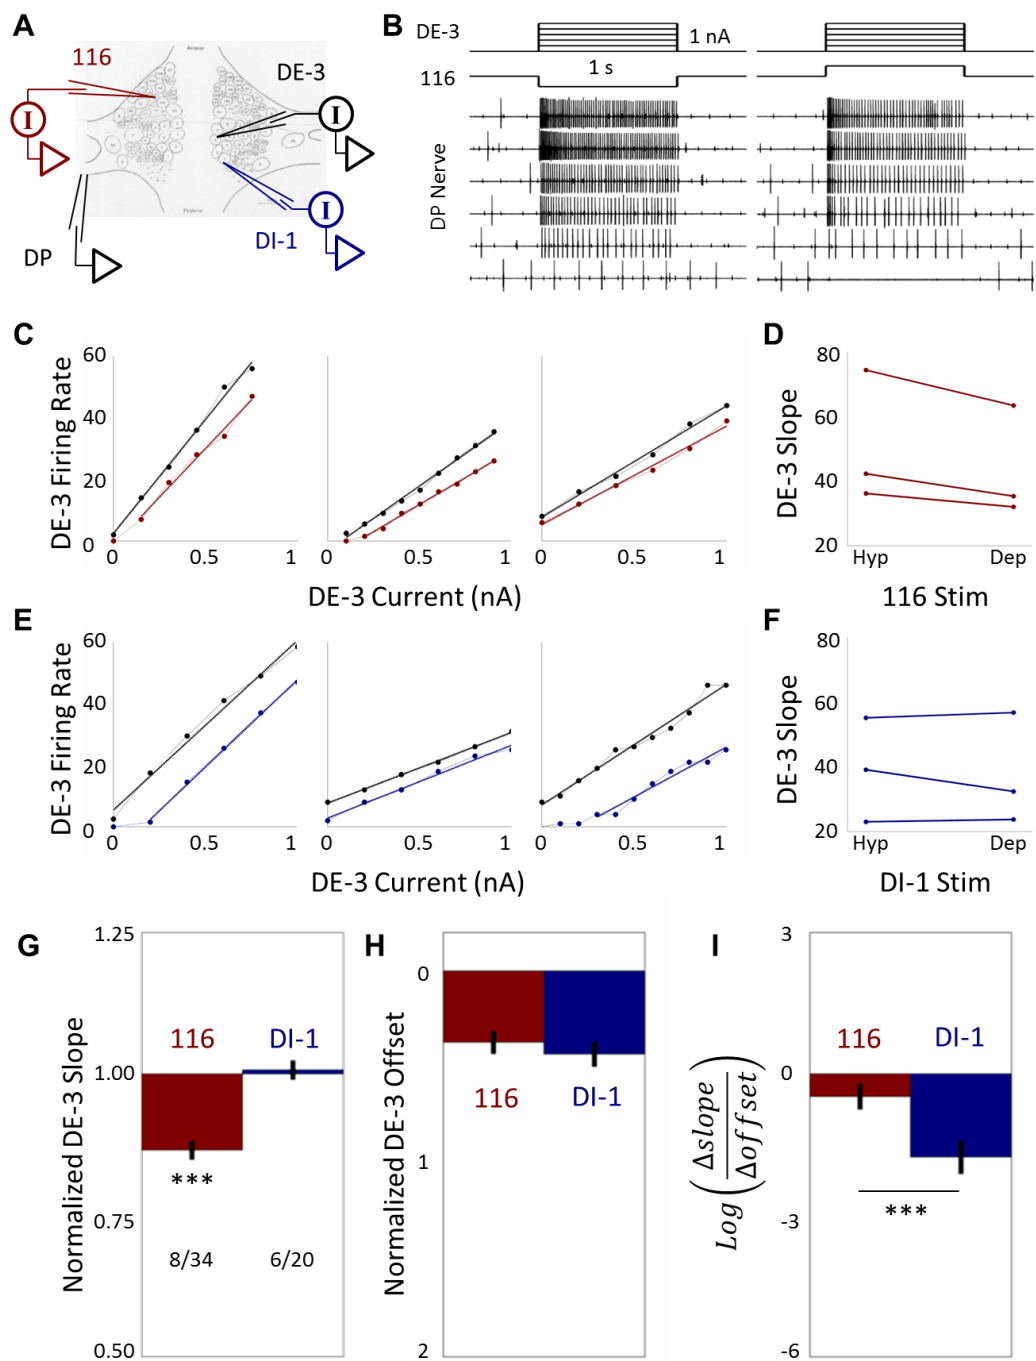
$$\Gamma = \frac{I_{Sh}}{I_D} = \frac{g_C}{g_l + g_C + g_\gamma} \quad (3.13)$$

Results

We investigated two inhibitory neurons that synapse onto cell DE-3, a motor neuron involved in the local bend reflex. This neuron receives input from a known inhibitory motor neuron, cell DI-1, as well as from cell 116, a previously uncharacterized inhibitory interneuron (Fig. 3.3).

To assess the effect of inhibition on their post-synaptic targets, we injected several levels of depolarizing current pulses into cell DE-3 while simultaneously depolarizing or hyperpolarizing cell 116 (Fig. 3.3B) or cell DI-1. We plotted DE-3’s input-output (IO) function when each inhibitory neuron was hyperpolarized (Fig. 3.3C,E, black lines) and depolarized (Fig. 3.3C,E, colored lines). Under all conditions, DE-3’s IO function was very close to linear over a large dynamic range, saturating only with firing frequencies above 80 Hz, the upper end of the phys-

Figure 3.3: Modulation of DE-3's input-output function. (A) We recorded intracellularly from an inhibitory neuron (cell 116 or cell DI-1) and cell DE-3 simultaneously while injecting current pulses into them. Extracellular electrodes recorded the action potentials of DE-3 from the DP nerve. (B) Different levels of depolarizing current steps (0 to +1nA) were passed into cell DE-3 (top) to measure its input-output function. Cell 116 was simultaneously either hyperpolarized (left) or depolarized (right). The largest class of spikes in the DP nerve are from cell DE-3 (bottom); each trace corresponds with a different current step into DE-3. (C) The firing rate of cell DE-3 is plotted against the current step size when cell 116 was silenced by hyperpolarization (black line) or when cell 116 was activated by depolarization (red line). The curves are fit with a straight line, ignoring firing rates less than 1 Hz or greater than 80 Hz. Three examples from different animals are shown. (D) The slope of the linear fits from C are plotted against cell 116 stimulus. In all three examples shown, the slope decreases when 116 is depolarized. (E and F) The same experiments were done with a different inhibitory neuron, cell DI-1. The slope of the IO function decreases in only one of the examples shown. (G) Several experiments were combined by normalizing the slope of each cell DE-3's IO function by its slope when the inhibitory neuron was hyperpolarized. The numbers indicate the number of different animals and the total number of experiments. (H) The offsets of the linear fits were normalized by the DE-3 slope and show that 116 and DI-1 have similar effects of the offset. (I) Trial-by-trial comparison of the relative effects of slope and offset from activation of the inhibitory neurons. Cell 116 exhibits a greater effect on the slope relative to the offset than does cell DI-1.



iological range of cell DE-3 (Mason & Kristan, 1982). We made linear fits to cell DE-3's IO function, ignoring firing rates below 1 Hz and above 80 Hz, and plotted the slope of the fit vs. the current injected into the inhibitory neuron (Fig. 3.3D,F). Activation of cell 116 appeared to decrease the slope of DE-3's IO function (Fig. 3.3C,D). In contrast during activation of cell DI-1, DE-3's IO function appeared parallel to when DI-1 was inhibited (Fig. 3.3E,F).

To test whether these apparent differences were statistically significant, we compared the combined results of recordings from many different cell pairs from several animals by normalizing the slope of each DE-3's IO function to the slope when the inhibitory neuron was silenced by hyperpolarization (black lines in Fig. 3.3C,E). Averaging data from eight different animals (34 total trials), activation of cell 116 showed a significant effect on the normalized slope of cell DE-3's IO function (14% decrease in slope, $p < 1e - 10$, t-test, Fig. 3.3G). In contrast, averaging data from six different animals (20 total trials) activation of cell DI-1 showed no significant effect on the slope of cell DE-3's IO function (1% increase in slope, $p > 0.60$, t-test). Both neurons had a similar effect on the y-intercept of the linear fits (Fig. 3.3H). To be sure that the effect was not a result of averaging or differences in stimulation efficacy, we compared the change in slope with the change in intercept on a trial-by-trial basis (Fig. 3.3I). Cell 116 has a significantly larger effect on the slope relative to offset than does DI-1 ($p < 1e - 6$).

Based on the synaptic location and reversal potentials, several theoretical studies have suggested conductance-based mechanisms that inhibitory synapses could use to implement multiplicative and additive operations (Blomfield, 1974; Vu & Krasne, 1992). One prediction is that additive inputs (like that from cell DI-1) propagate from remote dendrites as currents before reaching the spike initiation zone. Consistent with this prediction, previous work has shown that the synaptic connection from cell DI-1 to cell DE-3 is on one of the two most medial second-order processes of cell DE-3 on the contralateral side of the ganglion (Fig. 3.4A, blue box), and that cell DE-3's spike initiation zone is on the primary process beyond the secondary branching sites (Lytton & Kristan, 1989) (Fig. 3.4A, Axon). A second prediction, based on shunting inhibition, is that multiplicative inhibitory synapses should be situated on the path between the spike initiation zone and the additive inputs. Thus, the multiplicative connections would be predicted to form on the primary process of cell DE-3 lateral to the secondary branches (Fig. 3.4A, red box). To examine this, we filled cells DE-3 and 116 with dyes that fluoresce at different wavelengths (Fig. 3.4B). The bottom of Figure 3.4B shows four consecutive confocal planes of the two cells, which show that processes of cell 116 come into close apposition to cell DE-3 at several sites in the predicted location.

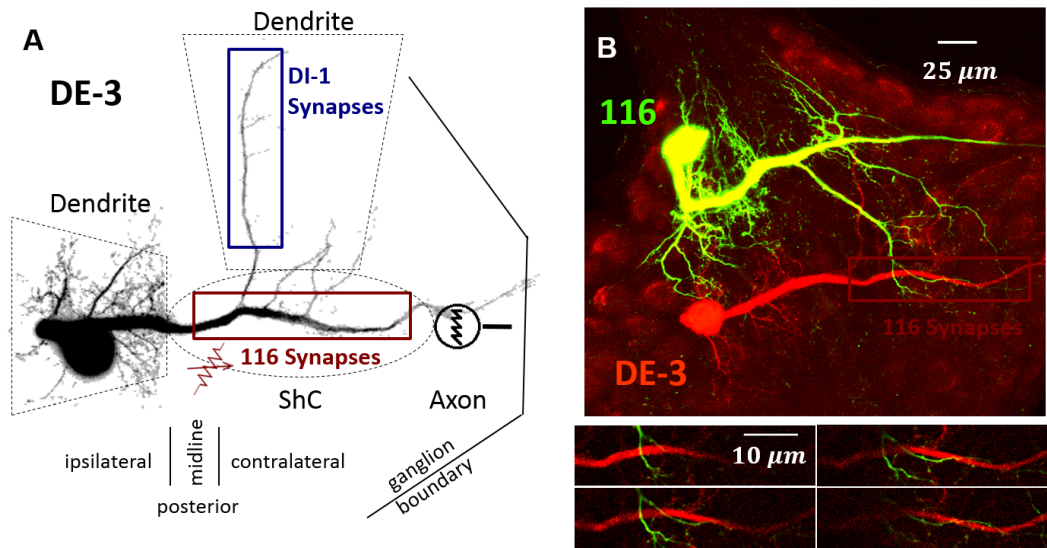


Figure 3.4: Shunting inhibition predicts synaptic structure of DE-3. (A) The structure of cell DE-3 after being filled with a fluorescent dye, Alexa 488. The components of the model are overlain on the predicted anatomical regions of cell DE-3. Cell DI-1 synapses are typically located on the first contralateral secondary dendrite (blue box). Previous models (Vu & Krasne, 1992) predict that cell 116 synapses would be on the main process of DE-3 (red box). (B) Cell DE-3 was filled with a red-fluorescing dye and cell 116 was filled with a green-fluorescing dye. The predicted site of 116 synapses from the model is shown as the red box. A projection at low magnification (top) and 4 consecutive slices of the image stack zoomed at the predicted location (bottom) are shown.

To address the differential modulation of DE-3, we created a multi compartmental model separated into three components: the dendrites, the “shunt-compartment” (ShC), and the axon (Fig. 3.5A). Since currents add together linearly, dendrites that transform synaptic activity into currents produce exclusively additive operations. A separate compartment with a conductance-based synaptic model can implement such linear operations, as long as the voltage of the dendritic compartment is sufficiently far from the reversal potential of the synapse. Further, the linearization of conductance-synapses can be enhanced with channels in the dendrites (Bernander et al. 1994). Therefore, rather than simulating an extra compartment to produce additive synaptic interactions, we simulated these interactions by simply adding the synaptic currents.

The axon is modeled with an integrate-and-fire (IF) spiking mechanism: if the voltage (V_{Ax}) reaches threshold (V_{θ}) a spike is triggered and the voltage is set to the reset potential (V_R). The spiking of the neuron is solely determined by the current that reaches the axon (I_{Sh}), and the instantaneous firing rate (I_{FR}) of the axon can be approximated as a linear function of this

current (Capaday & Van Vreeswijk, 2006) (see Methods):

$$I\tilde{F}R = mI_{Sh} + b \quad (3.8)$$

$$m = \frac{1}{C_{Ax}(V_{\theta} - V_R)} \quad (3.9)$$

$$b = -\frac{g_l(V_{\theta} + V_R)}{2C_{Ax}(V_{\theta} - V_R)} \quad (3.10)$$

Multiplicative inhibitory synapses increase a conductance (g_{γ}) that shunts the dendritic current (I_D) in the ShC (red channel in Fig. 3.5A). I_{Sh} flows to the axon through the coupling conductance (g_C). By setting the action-potential threshold relatively low (2mV above rest) and keeping the average voltage of the axon near the resting potential, the two-compartment model can be closely approximated. The gain, Γ , can then be derived based on the conductances (see Methods):

$$\Gamma = \frac{I_{Sh}}{I_D} = \frac{g_C}{g_l + g_C + g_{\gamma}} \quad (3.13)$$

This equation means that I_{Sh} is scaled relative to I_D by the ratio of the coupling conductance to the total conductance of the ShC. Equation 3.13 shows that the inhibitory input, which opens shunting inhibitory channels and increases g_{γ} , changes the ratio of I_{Sh} to I_D in a multiplicative fashion. Because the leak conductance, g_l , is typically much smaller than either the shunting or the coupling conductances, the shunting conductance affects the gain only if it is within the same order of magnitude as the coupling conductance. If the coupling conductance (g_C) is large compared to the shunting conductance (g_{γ}), for instance, then it dominates in Equation 3.13, and altering the shunting conductance has little effect on the gain.

To ensure that these approximations were valid, we simulated a full two-compartment spiking model (Fig. 3.5B, see Table 1 for parameters). We measured the steady-state firing rate in the model after injecting several constant levels of current, I_D , at four different levels of constant shunting conductance, g_{γ} . The results of these simulations (Fig. 3.5C, solid lines) and the corresponding linear approximations (dashed lines) closely matched. These data showed, despite the various approximations, that shunting inhibition multiplicatively scaled the firing rate in this spiking model as mathematically derived (Eq. 3.8, 3.13).

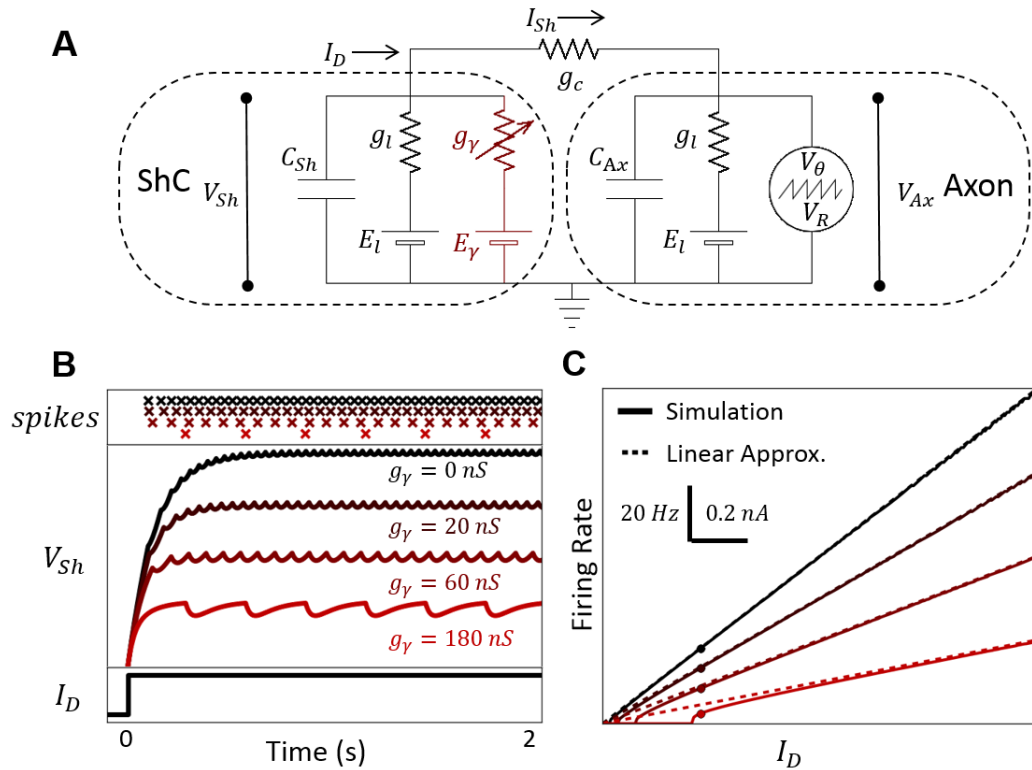


Figure 3.5: Multiplication via shunting in two-compartment model. (A) The dendrites transform synaptic activity into a current that reaches the ShC as I_D . This current is scaled multiplicatively by inhibitory shunting conductances (g_γ) before reaching the axon as I_{Sh} . (B) The response of the model to step current injection at four different levels of shunting inhibition. (C) The steady firing rate of the model is plotted versus the injected current for each level of shunting as solid lines, and the linear approximation of the model is plotted as dashed lines. Dots correspond to data in B. The linear fit is an excellent approximation except at low firing rates.

Discussion

Both DE-3 and inhibitory neurons have been previously noted for their role in the local-bend reflex (Kristan, 1982; Baca et al. 2008). The circuitry and physiology of the local-bend reflex has been well characterized, and computational models suggest that the information for the reflex is encoded in a population code (Lockery & Kristan, 1990; Lewis & Kristan, 1998). Theoretical and experimental work highlights the importance of multiplication for neural computation (Pouget & Sejnowski, 1997; Sussillo & Abbott, 2009; Carandini & Heeger, 2012). We identified cell 116 as an inhibitory neuron which could be implementing multiplicative operations in the leech. This contrasts with the purely linear nature in which DI-1 affects DE-3. Taken together this suggests that cell 116 and DI-1 are playing complementary computational roles in

the local bend reflex, and further investigation into their circuitry and physiology is warranted.

The shunting mechanism predicts that the multiplicative synapses would lie on the primary process between the axon and the secondary dendrites. Several putative contacts are seen at the light level in the predicted location (Fig. 3.4B), but further work needs to be done to validate that these contacts are actually synaptic connections. In mammalian cortical neurons, the cell soma is a prime candidate as the shunt compartment. Cortical pyramidal neurons receive targeted inhibitory input to their somas from PV+ interneurons (Markram, 2004). This finding reveals that PV+ synapses are situated in the necessary location to perform multiplicative operations via shunting inhibition, and several studies have shown that activation of PV+ neurons do indeed affect the slope of the input-output function of post-synaptic pyramidal cells (Attallah et al. 2012, Wilson et al. 2012). Attallah et al (2012) also show that PV+ neurons effect the offset of the IO function, consistent with the offset component of 116, and theoretical work suggests that pure multiplication is not necessary for neural computation. Rather, it is the non-linear effect that has important computational consequences.

The analytical approximation of the model reveals several important insights about this mechanism of gain control. Equation 1 of the approximation shows that the model is linear and we set the value of b to be zero. This makes $\hat{I}_{Input\ Gain}$ and $\hat{I}_{Output\ Gain}$ exactly the same (Silver, 2010), and this must be re-considered in the context of non-linearities. This mechanism of gain control could implement either of these types of gain depending on where the non-linear function arises. If non-linearities were produced by mechanisms in the dendritic tree, then shunting inhibition would result in output gain of this non-linear IO function, scaling it up and down (b would also need to be 0 for pure output gain). When the non-linearity arises in the axon, then shunting inhibition would be producing input gain, scaling the non-linear function left and right.

Our two compartment model is based on the assumption that there is some electrotonic separation between the synaptic inputs and the spike initiation zone. Because the conductivity of the intracellular medium is not infinitely large, descriptions of neurons as more than just “point” integrators are now seen widely, especially in models of the dendritic tree (Poirazi et al. 2003; Spruston, 2008; Larkum et al. 2009). Our shunting model requires that the total conductance of the shunt compartment is on par with the coupling conductance \hat{G} meaning that the current that flows out through all of the shunting channels in the ShC is about the same as the current that gets to the spike initiation zone. Grubb & Burrone (2010) show that the axon-initial segment (AIS) of cultured hippocampal neurons is regulated by homeostatic processes \hat{G} a neuron which

is over excited can shift the AIS further down the axon, which in turn decreases the neuron's excitability. Our model would produce this excitability change if the axon (comparable to the AIS in a hippocampal neuron) is moved away from the ShC (comparable to the hippocampal neuron's soma) resulting in a lower coupling conductance, which would decrease the gain of the neuron's input-output function (Equation refeq:Gamma). Homeostatic control of the coupling conductance and plasticity of shunting channels could combine to regulate a neuron's gain and the impact of multiplicative inhibition.

Shunting inhibition was previously hypothesized as the mechanism of multiplication, but several studies suggested that the shunting mechanism was not multiplicative when spiking dynamics are included (Holt & Koch, 1997; Chance et al. 2002). Many theoretical alternatives have arisen to address the mechanism of multiplication, but require complex circuitry (Hahnloser et al. 2000), sophisticated log-exp transformations (Gabbiani et al. 2002), or cumbersome noise manipulations (Chance et al. 2002, Mitchell et al. 2003, Ly & Doiron, 2009). The model that we propose here is the first to our knowledge that interprets utilizes the axon as a separate compartment within a multi-compartmental neuron to implement multiplication via shunting inhibition. Interestingly, the voltage-clamp nature of axon spiking (Koch et al. 1995) is the reason that shunting inhibition fails in a single compartment model (Holt & Koch, 1997), but we actually utilized this effect to control the gain in the two-compartment model. Further, this phenomenon suggests that firing rate is based only on the current that reaches the axon. By assuming there is some electrotonic separation between the ShC and the axon, our model restores shunting inhibition as a plausible mechanism for multiplicative operations in neurons with spiking dynamics. This model provides a simple, biologically plausible, and computationally efficient way to implement additive and multiplicative operations, and we have derived an approximation that allows this model to be easily fit to many neural phenomena.

Acknowledgements

Part of this chapter is presented as a manuscript that is in preparation as a publication as Frady, E.P., Yudice, J., Konanur, V., Todd, K.L., French, K.A., Kristan Jr, W.B. "Shunting and inhibitory currents co-regulate the input-output function of an identified leech neuron".

References

Atallah, B.V., Bruns, W., Carandini, M., Scanziani, M. (2012). Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron* 73, 159-170.

- Baca, S.M., Marin-Burgin, A., Wagenaar, D.A., Kristan Jr., W.B. (2008) Widespread inhibition proportional to excitation controls the gain of a leech behavioral circuit. *Neuron* 57, 276-289.
- Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M., Latham, P.E., Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142-1152.
- Bernander, O., Koch, C., Douglas, R. J. (1994). Amplification and linearization of distal synaptic input to cortical pyramidal cells. *J. Neurophysiol.* 72, 2743-2753
- Blomfield, S. (1974). Arithmetical operations performed by nerve cells. *Brain Res.* 69, 115-124.
- Briggman, K.L., Kristan, W.B. (2006). Imaging dedicated and multifunctional neural circuits generating distinct behaviors. *J. Neurosci.* 26(42), 10925-10933.
- Capaday, C., Van Vresswijk, C. (2006). Direct control of firing rate gain by dendritic shunting inhibition. *J. Int. Neurosci.* 5(2), 199-222.
- Carandini, M., Heeger, D.J., Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* 17(21), 8621-8644.
- Carandini, M., Heeger, D.J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience* 13, 51-62.
- Chance, F.S., Abbott, L.F., Reyes, A.D. (2002). Gain modulation from background synaptic input. *Neuron* 35, 773-782.
- Deneve, S., Latham, P.E., Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neurosci.* 2(8), 740-745.
- Gabbiani, F., Krapp, H. G., Koch, C., Laurent, G. (2002). Multiplicative computation in visual neuron sensitive to looming. *Nature* 420, 320-324.
- Goodman, D.F., Brette, R. (2008). Brian: a simulator for spiking neural networks in Python. *Front. Neuroinform.* 2(5), 1-10.
- Hahnloser, R.H.R., Sarpeshkar, R., Mahowalk, M.A., Douglas, R.J., Seung, H.S. (2000) Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature* 405(22), 947-951.
- Häuser, M., London, M. (2005). Dendritic computation. *Annual Review of Neuroscience* 28, 503-532.
- Holt, G.R., Koch, C. (1997). Shunting inhibition does not have a divisive effect on firing rates. *Neural Comp.* 9, 1001-1013.

- Koch, C., Bernander, O., Douglas, R.J. (1995). Do Neurons Have a Voltage or Current Threshold for Action Potential Initiation? *J. Comput. Neurosci.* 2(1), 63-82.
- Kristan Jr., W.B. (1982). Sensory and motor neurons responsible for the local bending response in leeches. *J. Exp. Biol.* 96, 161-180.
- Larkum, M.E., Nevian, T., Sandler, M., Polsky, A., Schiller, J. (2009) Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. *Science* 325(5941), 756-760.
- Lewis, J.E., Kristan Jr., W.B. (1998a). A neuronal network for computing population vectors in the leech. *Nature* 391, 76-79.
- Lewis, J.E., Kristan Jr., W.B. (1998b). Quantitative analysis of a directed behavior in the medicinal leech: implications for organizing motor output. *J. Neurosci.* 18(4), 1571-1582.
- Lewis, J.E., Kristan Jr., W.B. (1998c). Representation of touch location by a population of leech sensory neurons. *J. Neurophysiol.* 80, 2584-2592.
- Lockery, S.R., Kristan Jr., W.B., (1990a) Distributed processing of sensory information in the leech. I. Input-output relations of the local bending reflex. *J. Neurosci.* 10(6):1811-1815.
- Lockery, S.R., Kristan Jr., W.B. (1990b). Distributed processing of sensory information in the leech. II. Identification of interneurons contributing to the local bending reflex. *J. Neurosci.* 10(6), 1816-1829.
- Louie, K., Glimcher, P.W. (2010) Separating value from choice: delay discounting activity in the lateral intraparietal area. *J. Neurosci.* 30(16), 5498-5507.
- Ly, C., Doiron, B. (2009). Divisive Gain Modulation with Dynamics Stimuli in Integrate-and-Fire Neurons. *PLoS Comput. Biol.* 5, 1-12.
- Lytton, W.W., Kristan, W.B. (1989). Localization of a leech inhibitory synapse by photo-ablation of individual dendrites. *Brain Res.* 504(1): 43-48.
- Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neurosci.* 9(11), 1432-1438.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neurosci.* 5, 793-807.
- Mason, A., Kristan, W.B., Jr. (1982). Neuronal excitation, inhibition and modulation of leech longitudinal muscle. *J. Comp. Physiol.* 146, 527-536.
- Mitchell, S.J., Silver, R.A. (2003). Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron* 38, 433-445.

- Ohshiro, T., Angelaki, D.E., DeAngelis, G.C. (2011). A normalization model of multisensory integration. *Nature Neurosci.* 14(6), 775-782.
- Olsen, S.R., Bhandawat, V., Wilson, R.I. (2010). Divisive normalization in olfactory population codes. *Neuron* 66, 287-299.
- Poirazi, P., Brannon, T., Mel, B. (200-3). Pyramidal neuron as a two-layer neural network. *Neuron* 37(6), 821-827.
- Pouget, A., Sejnowski, T.J. (1997). Spatial transformations in the parietal cortex using basis functions. *J. Cog. Neurosci.* 9(2), 222-237.
- Pouget, A., Dayan, P., Zemel, R.S. (2000). Information processing with population codes. *Nature Reviews Neurosci.* 1, 125-132.
- Pouget, A., Dayan, P., Zemel, R.S. (2003). Inference and computation with population codes. *Annu. Rev. Neurosci.* 26, 381-410.
- Prochazka, A. (1989) Sensorimotor gain control: a basic strategy of motor systems? *Prog. In Neurobiology* 33, 281-307.
- Ringach, D.L. (2010). Population coding under normalization. *Vision Res.* (50), 2223-2232.
- Solomon, S.G., Lee, B.B., Sun, H. (2006). Suppressive surrounds and contrast gain in magnocellular-pathway retinal ganglion cells of macaque. *J. Neurosci.* 26(34), 8715-8726.
- Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neurosci.*
- Sussillo, D., Abbott, L.F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544-557.
- Vu, E.T., Krasne, F.B. (1992). Evidence for a computational distinction between proximal and distal neuronal inhibition. *Science* 255(5052), 1710-1712.
- Wilson, N.R., Runyan, C.A., Wang, F.L., Sur, M. (2012). Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature* 488, 343-348.
- Zoccolan, D., Cox, D.D., DiCarlo, J.J. (2005). Multiple object response normalization in monkey inferotemporal cortex. *J. Neurosci.* 25(36), 8150-8164.

4 Experimental Methods

Leech Behavior and Anatomy

The leech is a segmented worm and based on different types of stimulations and contexts can produce many complicated behaviors (Muller et al. 1981). Figure 4.1 shows four behaviors that we can study experimentally and are produced by stimulating the animal in different locations along the anterior-posterior axis. Shortening typically results from the stimulation of the anterior of the animal (Palmer et al. 2014) and is a simultaneous contraction of the longitudinal muscles across the whole animal (Magni & Pellegrino, 1978; Shaw & Kristan, 1995; Kristan et al. 2005). Local-bending is a reflex that bends the animal's body away from the stimulation site, by contracting longitudinal muscles at the location of the stimulus and relaxing the muscles on the opposite side (Kristan, 1982; Lockery & Kristan, 1990a). More posterior stimulations can result in swimming or crawling (Briggman et al. 2005). Swimming is a dorsal-ventral sine-wave oscillation that the leech produces to propel itself through water. Crawling is a cyclical contraction-elongation process that the leech uses to move on surfaces.

The neuronal processes that underly these behaviors have been studied by comparing intact, semi-intact and isolated preparations (see Figure 4.2). A typical behavior is at first carefully studied in the intact animal – the kinematics of the muscles and position of the body can be objectively studied by attaching beads to the animals and tracking their locations. The leech can, remarkably, be partially dissected and still produce these behaviors. These semi-intact preparations allow some access to the nervous system where motor-neuron impulses can be monitored by recording from the nerves going to the muscles (Gray et al. 1938). The movements of the animal can then be related to these neuronal signals. For more complex physiology and VSD imaging, the entire leech's nervous system can be dissected out from the animal. Even without sensory feedback, the nervous system contains pattern generators that can produce the “fictive” versions of each of these behaviors. These are fictive behaviors because although we do not observe any movements (as there are no muscles attached to the nervous system), the motor-neurons produce similar output patterns as observed in the semi-intact preparations.




Figure 4.1: Many different behaviors can be elicited in the leech that are suitable for experimental study. Although there are many contextual factors that determine behavioral production, an important cue is the location of a stimulation along the anterior-posterior axis of the animal. Stimulation to the anterior end of the animal will produce a withdrawal response called shortening. More medial stimulations can produce local-bending and swimming, and stimulation to the posterior end of the animal can produce crawling.

The leech nervous system is made up of a head-brain, a tail-brain, and 21 mid-body ganglia (Kristan et al. 2005). It appears that each of the 21 ganglia are made up of a homologous set of about 400 neurons (Macagno, 1980). The head-brain and tail-brain are formed during development through the fusion of 4 and 7 ganglia, respectively (Kristan et al. 2005). These structures were specialized to control the anterior and posterior suckers, but also may have some executive-type functionality – for instance there are neurons in the head-brain that can turn on and off certain pattern generators and send axons through the length of the animal (Brodfhuehrer & Friesen, 1986; Esch et al. 2002; Crisp & Mesce, 2004; Puhl & Mesce, 2010; Mullins & Friesen, 2012). Ganglia 5 and 6 have additional neurons likely to handle the control of the reproductive organs that are innervated by these ganglia (Macagno et al. 1986). Although throughout the course of evolution these structures have become specialized, they were built upon this set of 400 neurons that we see repeated throughout the animal. Understanding this structure will thus be extremely important in understanding the leech’s full nervous system.

Figure 4.3 shows the ventral aspect of a single ganglion on the right. The 400 neurons




Figure 4.2: The relation between behavior and the leech nervous system is studied by examining intact, semi-intact and isolated nerve cord preparations. With the semi-intact preparation we can correlate the movements produced during different behaviors with the signals coming from the nervous system. These same signals can be seen in the isolated nerve cord and are called “fictive” behaviors. We can study the isolated nervous system in much more detail, utilizing intracellular recordings and voltage-sensitive dye imaging techniques.

that make up a ganglion are all monopolar and their somatas are spread across the ventral and dorsal faces of the ganglion. The processes and connections form between the 2 layers of cell bodies. Ganglia interact by sending axons anteriorly and posteriorly down the connective, which are the nerves that connect each ganglia. Of the 400 neurons that make up the ganglion, only about a third have been identified. The sensory neurons (shown in blue; Nicholls & Baylor, 1968), and many of the motor neurons (shown in red; Stuart, 1970; Ort et al. 1974; Thompson & Stent, 1976; Sawada et al. 1976; Norris and Calabrese, 1987) account for most of the identified cells. A handful of interneurons have been identified and studied (shown in green). The majority of the unidentified or partially characterized neurons are interneurons, which are hard to study due to their typically smaller size and lack of obvious characteristic features.

The neurons that we have characterized thus far in the leech have been well studied essentially because they are easily identifiable. These neurons have easily accessible features that give away their identity and can be used to find the same cells across animals for careful

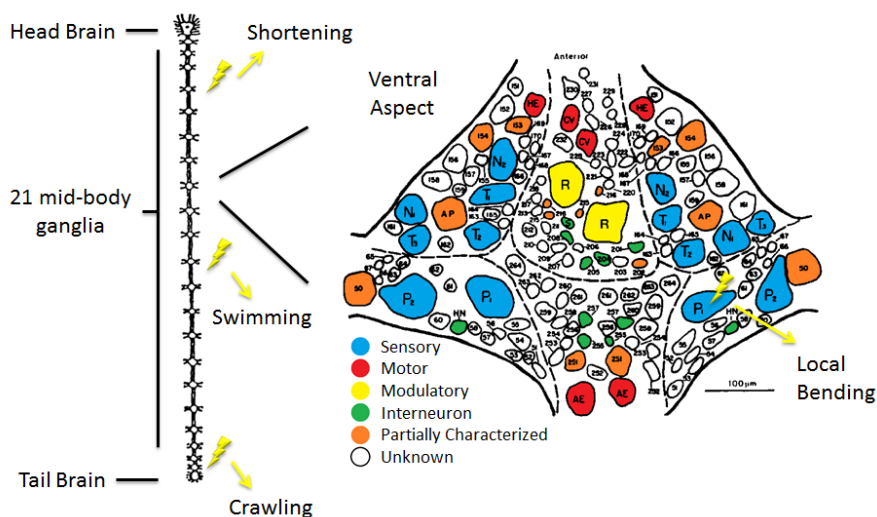


Figure 4.3: The leech nervous system is made up of a head-brain, a tail-brain and 21 mid-body ganglia (left). Each ganglia is composed of about 400 neurons whose somas are spread across the ventral and dorsal faces of the ganglion (right). In color are shown many of the previously identified neurons - mostly consisting of the sensory and motor neurons. Shown in white are neurons that have not been identified, whose numbers are underrepresented in this figure. The isolated nervous system can be stimulated to produce the various different behaviors by stimulating different nerves along the length of the nerve-cord. Local-bending can be initiated by stimulating the pressure-sensitive sensory neurons (P cells) with micro-electrodes.

study. Some neurons, such as the Retzius cells, can be identified simply from visual inspection – their soma location and size is enough information to recognize the cells across preparations. Only a few neurons can be identified from visual inspection alone, but often the soma size and position of a neuron will narrow down the possibilities (Weeks, 1982a; Weeks, 1982b). Many neurons, such as the sensory neurons and the S cell (Laverack, 1968; Mistick, 1974; Bagnoli et al. 1975; Mistick, 1978), can be identified by their unique action potentials, requiring only a micro-electrode recording for identification. However, most of the unidentified interneurons that make up the nervous system have very similar anatomical and physiological properties. Typically their action-potentials, as recorded from the soma, are only about 5 mV and have no clear defining features. These neurons require finding more complex properties in order to identify them, such as their functional roles during different behaviors.

Imaging is an attractive option to identify cells in a scalable fashion. In order to identify the other cells in the ganglion, it is necessary to extract functional information from each of them and compare the activity of different neurons to find distinguishing features. Doing this with classical electrophysiological techniques is problematic, because only a few neurons can

be studied at once and it would take an unfeasible number of experiments to collect the data necessary to identify cells. We developed a Voltage-sensitive dye (VSD) to image the activity of neurons during behavioral states.

Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires.

Abstract

Fluorescence imaging is an attractive method for monitoring neuronal activity. A key challenge for optically monitoring voltage is development of sensors that can give large and fast responses to changes in transmembrane potential. We now present fluorescent sensors that detect voltage changes in neurons by modulation of photo-induced electron transfer (PeT) from an electron donor through a synthetic molecular wire to a fluorophore. These dyes give bigger responses to voltage than electrochromic dyes, yet have much faster kinetics and much less added capacitance than existing sensors based on hydrophobic anions or voltage-sensitive ion channels. These features enable single-trial detection of synaptic and action potentials in cultured hippocampal neurons and intact leech ganglia. Voltage-dependent PeT should be amenable to much further optimization, but the existing probes are already valuable indicators of neuronal activity.

Introduction

Fluorescence imaging can map the electrical activity and communication of multiple spatially resolved neurons and thus complements traditional electrophysiological measurements (Scanziani & Hausser, 2009, Paterka et al. 2011). Ca^{2+} imaging is the most popular of such techniques, because the indicators are well-developed (Poenie et al. 1985; Tsein et al. 1985; Grynkiewicz et al. 1985; Minta et al. 1989), highly sensitive (Grynkiewicz et al. 1985; Minta et al. 1989), and genetically encodable (Nagai et al 2001; Miyawaki et al. 1997; Tian et al. 2009; Mank et al. 2008; Heim & Griesbeck, 2004; Palmer et al. 2006; Horikawa et al. 2010), enabling investigation of the spatial distribution of Ca^{2+} dynamics in structures as small as dendritic spines and as large as functional circuits. However, because neurons translate depolarizations into Ca^{2+} signals via a complex series of pumps, channels, and buffers, fluorescence imaging of Ca^{2+} transients cannot provide a complete picture of electrical activity in neurons. Observed Ca^{2+} spikes are temporally low-pass filtered from the initial depolarization and provide lim-

ited information regarding hyperpolarizations and subthreshold events. Direct measurement of transmembrane potential with fluorescent indicators would provide a more accurate account of the timing and location of neuronal activity. Despite the promise of fluorescent voltage-sensitive dyes (VSDs), previous classes of VSDs have each been hampered by some combination of insensitivity, slow kinetics (Cacciatore et al. 1999; González & Tsien, 1997; González & Tsien, 1995), heavy capacitative loading (Sjulson & Miesenböck, 2008; Akemann et al. 2009; Wang et al. 2010; Bradley et al. 2009; Chanda et al. 2005), lack of genetic targetability, or phototoxicity. Two of the more widely used classes of VSDs, electrochromic and FRET dyes, illustrate the problems associated with developing fast and sensitive fluorescent VSDs.

Electrochromic dyes respond to voltage through a direct interaction between the chromophore and the electric field (Scheme 4.4A). This Stark effect leads to small wavelength shifts in the absorption and emission spectrum. Because the electric field directly modulates the energy levels of the chromophore, the kinetics of voltage sensing occur on a timescale commensurate with absorption and emission, resulting in ultrafast (*fs* to *ps*) hypso- or bathochromic shifts many orders-of-magnitude faster than required to resolve fast spiking events and action potentials in neurons. This small wavelength shift dictates that the fluorescence signal can be best recorded at the edges of the spectrum, where intensity varies most steeply as a function of wavelength. The largest linear responses are $\sim 28\% \Delta F/F$ per 100 mV (Hubenar et al. 2003), although more typical values are $\sim 10\%$ per 100 mV (Grinvald, 1983; Fluhler et al. 1985). Photo-induced electron transfer (PeT)-based Ca^{2+} probes, such as fluo-3, give $\Delta F/F$ values of up to 150% for action potentials in cultured hippocampal neurons (Jacobs & Meyer, 1997). Therefore, although electrochromic dyes can keep pace with fast voltage oscillations in neurons, their insensitivity limits the systems in which these dyes can successfully report on voltage changes.

FRET-based voltage sensors use lipophilic anions that intercalate into the cellular membrane and distribute between the inner and outer leaflets depending upon the transmembrane potential (Scheme 4.4B). The Nernstian distribution is monitored by a second fluorophore immobilized on one side of the membrane, which undergoes FRET preferentially with the mobile anions on the same side of the membrane. Translocation of the lipophilic anion through the lipid bilayer governs the kinetics of voltage sensing, which can be in the millisecond range. Although these two-component systems can give large changes in intensity (5-34%) (Chanda et al. 2005) or ratio (80% per 100 mV) (González & Tsien, 1997), the slow translocation of mobile charges in the plasma membrane introduces a capacitative load and hampers the ability of the reporter to monitor fast changes.




Figure 4.4: Mechanisms of fluorescent voltage sensing. (A) Electrochromic VSDs sense voltage through the Stark effect, whereby the chromophore interacts directly with the electric field. Absorption of a photon significantly alters the excited state molecular dipole, which at hyperpolarizing potentials is stabilized (Left). At depolarizing potentials the charge shift inverted state is destabilized (Right). Changes in the energy levels of the chromophore result in small spectral shifts in the emission of the dye. (B) FRET-pair voltage sensors use lipophilic anions (red), which partition in a voltage-dependent fashion on the inner or outer leaflet of the membrane. Depolarization causes translocation of the anion, which can now quench the fluorescence of an immobilized fluorophore (green). (C) Molecular wire PeT VSDs depend upon the voltage-sensitive electron transfer from an electron-rich donor (orange) through a membrane spanning molecular wire (black) to a fluorescent reporter (green). At hyperpolarizing potentials, the electric field is aligned antiparallel to the direction of electron transfer, resulting in efficient PeT and quenched fluorescence (Left). Depolarization aligns the electric field in the direction of PeT, decreasing the rate of electron transfer and increasing fluorescence (Right).

To combine the best features of electrochromic and FRET-based VSDs, we have now tested a unique mechanism for voltage sensing, PeT through molecular wires. In these PeT sensors, a fluorescent reporter connects to an electron-rich quencher via a molecular wire, which minimizes the exponential distance dependence of intramolecular electron transfer (Davis et al. 1998) and allows efficient electron transfer over a major fraction of the thickness of the plasma membrane. At resting or hyperpolarized potentials, the transmembrane electric field promotes electron transfer from the quencher to the excited-state fluorophore through the molecular wire, quenching fluorescence (Scheme 4.4C). Depolarization reverses the electric field, hinders electron transfer, and brightens fluorescence (de Silva et al. 1995), just as Ca^{2+} binding dequenches indicators like fluo-3 (Adams, 2010). Electron transfer occurs within pico- to nanoseconds after photon absorption and returns to its initial state within a microsecond (Davis et al. 1998; de la Torre et al. 2005), slower than the electrochromic mechanism but essentially instantaneous on a biological timescale. Because electron transfer reverses quickly and is driven by photon absorption rather than membrane potential changes, capacitive loading should be negligible, as




Figure 4.5: Synthesis of VF probes

calculated in the Supplementary Appendix A. A full electronic charge traverses a Marcus-type thermal activation barrier to sense a large fraction of the membrane voltage, making voltage sensitivity high (Li, 2007). Quenching of the fluorescent reporter by the electron-rich donor modulates the fluorescence quantum yield independent of wavelength, permitting efficient use of photons for excitation and emission, allowing lower light levels or dye concentrations to be used. We report here the design, synthesis, and application of the VoltageFluor (VF) family of fluorescent sensors as molecular wire PeT-based probes for voltage imaging in neurons.

Results

Design and Synthesis of VF Sensors. Our initial voltage sensors incorporate dichloro-sulfofluorescein as a membrane-impermeant fluorophore, a p-phenylenevinylene (PPV) molecular wire, and N,N-dimethyl- or dibutylaniline as an electron-rich quencher (Scheme 4.5). VF-1.4.Cl comprises 2,7-dichlorosulfofluorescein connected via one vinylene unit to dibutylaniline (hence VF1.4.Cl). VF2.4.Cl adds a second PPV unit, and VF2.1.Cl features the same configuration, with methyl substituted in place of butyl groups.

Correct positioning of the fluorophore-wire donor within the membrane is vital to take advantage of the vectoral nature of the transmembrane electric field and electron transfer. First, the longitudinal axis of the molecular wire must be normal to the plane of the plasma membrane, to sample the full electric field. Second, dye molecules must all align in the same direction to avoid canceling out the electron transfer effect. Positioning the fluorophore at the extracellu-

lar leaflet of the membrane ensures fluorescence brightening upon depolarization; the opposite orientation of PeT would give fluorescence quenching upon depolarization.

The negatively charged sulfofluorescein will preclude dye internalization and force an orientation in which the fluorophore adsorbs to the outer leaflet of the plasma membrane, with the lipophilic molecular wire and alkyl aniline dangling into the lipid bilayer. As an intervening spacer, PPV molecular wires are an ideal choice because of their low attenuation values (Davis et al. 1998), synthetic tractability, and demonstrated ability to conduct current through lipid bilayers (Gerner et al. 2010). Anilines are common PeT donors and the di-alkyl groups should enhance uptake into the plasma membrane.

A modular synthetic design both allows for rapid generation of the voltage sensors and enables future derivatization (Scheme 4.5 and Supplementary Appendix A). Coupling of the molecular wire styrene unit 1, available in one step from 4-di-butylaminobenzaldehyde, with iodo-functionalized dichlorosulfofluorescein 2 via a Pd-catalyzed Heck reaction gives VF1.4.Cl in good yield. An analogous reaction with molecular wire 3, available in two steps from 1, gives VF2.4.Cl in 70% yield. A parallel reaction beginning from styrene 4 furnishes VF2.1.Cl in good yield. All dyes feature emission and excitation profiles typical of dichlorofluoresceins (VF1.4.Cl: $\lambda_{max} = 521nm$, $\epsilon = 93,000M^{-1}cm^{-1}$, $\lambda_{em} = 534nm$, $\Phi = 0.24$; VF2.4.Cl: $\lambda_{max} = 522nm$, $\epsilon = 97,000M^{-1}cm^{-1}$, $\lambda_{em} = 536nm$, $\Phi = 0.054$; VF2.1.Cl: $\lambda_{max} = 522nm$, $\epsilon = 98,000M^{-1}cm^{-1}$, $\lambda_{em} = 535nm$, $\Phi = 0.057$, 5 mM sodium phosphate, pH 9, 0.1% Triton X-100) (Supplementary Appendix A, Fig. S1). The dibutyl (VFx.4.Cl) dyes stain the cell membranes of HEK293 cells when loaded at a concentration of 2 μM for 15 min at 37° C in buffer with 0.1% DMSO as cosolvent (Fig. 4.4A, and Supplementary Appendix A, Fig. S2). VF2.1.Cl requires even lower dye concentrations (100 nM) and gives bright staining of HEK cell membranes, which is likely to be because of the greater aqueous solubility of VF2.1.Cl compared with VF2.4.Cl in aqueous solution (Supplementary Appendix A, Fig. S2). The membrane retention of the second generation dyes (VF2.x.Cl) is in contrast to di-4-ANEPPS, which at the same loading conditions, shows significant uptake into internal membranes. The bleach rates of the probes were tested in HEK cells and compared with di-4-ANEPPS. The bleach rates for VF1.4.Cl, VF2.4.Cl, and VF2.1.Cl at 7 W/cm² were measured to be $3.9 \pm 0.1 \times 10^{-2}s^{-1}$, $1.8 \pm 0.1 \times 10^{-2}s^{-1}$, and $8.0 \pm 0.1 \times 10^{-3}s^{-1}$, respectively. These results are two-, four-, and ninefold smaller than di-4-ANEPPS, which has a bleach rate, under identical illumination conditions, of $6.9 \pm 0.1 \times 10^{-2}s^{-1}$.

Characterization of the Voltage Response of VF Sensors. We characterized the volt-




Figure 4.6: Characterization of VF sensors in HEK cells. (A) Confocal image of HEK 293 cells stained with $2\mu\text{M}$ VF2.4.Cl. (Scale bar, $20\mu\text{m}$.) (B) Fractional changes in VF2.4.Cl fluorescence during a series of voltage steps to $+100$ or -100 from a holding potential of -60 mV (40-mV increments). (C) Fractional changes in VF2.4.Cl fluorescence from B plotted against membrane potential for voltage changes from a holding potential of -60 mV. Each datapoint represents three to four separate measurements. Error bars are SEM.

age sensitivity of all three indicators by making tight-seal whole-cell recordings of HEK293 cells stained with the VF sensors. Cells were voltage-clamped at -60 mV holding potential and sequentially stepped to depolarizing and hyperpolarizing potentials at 20-mV increments (Fig. 4.6B). For all three dyes, depolarizing steps produced fluorescence increases, whereas hyperpolarizing steps produced fluorescence decreases, in keeping with the proposed PeT mechanism. The fluorescence response is linear over the range of -100 mV to $+100$ mV (Fig. 4.6C), with voltage sensitivities $\Delta F/F$ per 100 mV of $20 \pm 1\%$ for VF1.4.Cl and $25 \pm 1\%$ for VF2.4.Cl. This statistically significant increase in voltage sensitivity ($P < 0.05$, two-tailed Student t test) is expected upon increasing the length of the molecular wire, and is 2.5- to 4-times more sensitive than di-4-ANEPPS, which, in our hands, gives sensitivities of between 6% and 10% $\Delta F/F$ per 100 mV. VF2.1.Cl shows fluorescence increases upon depolarization similar to VF2.4.Cl, with a voltage sensitivity of $27 \pm 1\%$ per 100 mV. This value is not significantly different from the sensitivity of VF2.4.Cl, suggesting that voltage sensitivity is largely determined by the length of the molecular wire and that the small increase in electron donating ability of the butyl relative to the methyl groups makes only a relatively small contribution to the increased voltage sensitivity of VF2.1.Cl. To investigate the speed of response of VF probes, we again made whole-cell recordings of cells stained with VF2.4.Cl, applied 100-mV depolarizing steps from a holding potential of -60 mV, and recorded both the electrophysiological signals and the optical signals. Fitting the electrophysiological and optical recordings gave identical time constants for both beginning and end of the pulse ($\tau_{ON,phys} = 139 \pm 0.2\mu\text{s}$, $\tau_{ON,optical} = 138 \pm 14\mu\text{s}$, $\tau_{OFF,phys} = 142 \pm 0.4\mu\text{s}$, $\tau_{OFF,optical} = 147 \pm 19\mu\text{s}$), showing that VF2.4.Cl and related sensors do not introduce any detectable lag in their fluorescence response to voltage, consistent with a PeT-based mechanism

for voltage sensing (Fig. 4.7A and B).

The fluorescence response of VF2.4.Cl to voltage changes is insensitive to the excitation wavelength, as is true for PeT-based probes, such as fluo-3 and Calcium Green-1. We assayed the wavelength dependence of VF2.4.Cl by changing the excitation wavelength in 5-nm steps and determined that the fluorescence response of VF2.4.Cl to a 100-mV depolarization from a holding potential of -60 mV varied only about 15% when testing wavelengths from 445 to 500 nm (Fig. 2C). In comparison, di-4-ANEPPS varies by nearly 100% over its excitation spectrum (Montana et al. 1989), and the PeT-based Ca^{2+} sensor, Calcium Green-1, varies by $\sim 20\%$ (Supplementary Appendix A, Fig. S3). These comparisons show that the wavelength independence of the voltage sensitivity is more consistent with PeT than a wavelength-shifting mechanism, such as electrochromism or solvatochromism (a wavelength shift because of alteration in local solvation).

PeT-based molecular wire sensors do not affect neuronal excitability by capacitive loading. We injected hyperpolarizing current into the Retzius cells of leech ganglia preparations and compared the time constants for these voltage steps in ganglia under different dye loading conditions. The ganglia were stained with dye at three times the working concentration [either VF2.1.Cl or diSBA-C4-(3) (Cacciatore et al. 1999)] and compared with unloaded cells. The presence of the translocating dye oxonol 413 substantially increases the capacitive load on the membrane, as measured by the increase in the RC time constant for the hyperpolarizing injection (Fig. 4.7D). On the other hand, ganglia loaded with VF2.1.Cl show no difference from control cells, demonstrating that molecular wire sensors place negligible capacitive load on the cell (Fig. 4.7D), confirming the predictions of the Supplementary Appendix A.

Detection of Action Potentials by VF2.4.Cl in Mammalian Neurons. To assess whether VF probes can detect action potentials in single trials, we used cultured rat hippocampal neurons. Bath application of $2\mu M$ VF2.4.Cl showed bright cell staining limited to the cell membranes of neurons and their support cells (Fig. 34.8). We then injected current into a neuron under whole-cell patch-clamp mode to trigger single action potentials and used a high-speed, back-illuminated EMCCD camera to track fast optical signals from VF2.4.Cl, enabling us to resolve action potentials in neurons in single sweeps (Fig. 4.8B). The optical trace matched the physiology trace and gave about a $20\% \Delta F/F$ increase in fluorescence and a 16:1 signal-to-noise ratio (SNR) in a single trial. The fact that VF2.4.Cl detected action potentials without spike-timed averaging suggests the possibility of measuring spontaneous action potentials in neurons at sites away from the recording pipette.




Figure 4.7: Characterization of the speed, wavelength sensitivity, and capacitance of the VF2 fluorescence response. (A) Rising edge of a 100-mV depolarizing step from -60 mV in HEK cells stained with VF2.4.Cl. (B) Falling edge of the same step. Black, solid trace is the integrated current measured electrophysiologically; red points are the optical recording. Time constants are calculated by fitting a monoexponential equation to each side of the step. Traces are the average of 100 sequential trials. (C) Voltage sensitivity vs. excitation wavelength. The normalized response of VF2.4.Cl to a 100 mV depolarization from -60 mV in HEK cells is plotted in red, and the excitation spectrum in HEK cells is the dotted black line. Error bars are SEM for $n = 3$ experiments. (D) Measurement of capacitive loading in leech Retzius cells. Traces show the normalized voltage decay following hyperpolarizing current injection into Retzius cell stained with $3\times$ VF2.1.Cl (red trace), $3\times$ oxonol 413 (black trace), or nothing (gray trace). (Inset) An expanded time scale revealing no difference between cells stained with VF2.1.Cl and control cells.




Figure 4.8: VF2 dyes resolve action potentials in neurons. (A) Rat hippocampal neurons stained with $2\mu\text{M}$ VF2.4.Cl for 15 min show strong membrane staining. (Scale bar, $20\mu\text{m}$.) (B) VF2.4.Cl can detect evoked action potentials in rat hippocampal neurons in single trials. The black trace is the recorded electrophysiology signal. Individual points represent the optical signal from VF2.4.Cl captured with a high speed EMCCD camera at a rate of 2kHz . (C) Optical imaging of spontaneous activity in leech Retzius cells using the dye VF2.1.Cl. Desheathed midbody leech ganglion stained with 200nM VF2.1.Cl for 15 min. Pixels within the region of interest (red circle around a single Retzius cell body) were averaged in each frame to produce the optical trace. (Scale bar, $25\mu\text{m}$.) (D) Simultaneous optical and electrophysiological recording of spontaneous activity in cell from C. The red trace is the hi-pass filtered VF2.1.Cl signal, sampled at 50Hz . The black trace is the electrophysiological recording, sampled at 10kHz . The optical trace shows near-perfect matching of the subthreshold membrane potential and a clear detectable signal indicating action potentials. Action potentials have variable amplitudes in the optical traces because of the relatively slow optical sampling rate (Supplementary Appendix A, Fig. S4).

Monitoring Spontaneous Activity in Leech Ganglia with VF2.1.Cl. A more stringent test of the usefulness of the PeT-based VSD is to determine whether it can accurately measure subthreshold activity in heterogeneous preparations. For this test we used leech ganglia, because their neurons have been well studied using both electrophysiological and other VSD recordings (Cacciatore et al. 1999, Briggman et al. 2005; Salzberg et al. 1977; Ross et al. 1987). We isolated a midbody ganglion and stained it with VF2.1.Cl for 15 min at 22° C (Fig. 4.8C). Insertion of a sharp electrode (25M Ω) into a Retzius cell to enabled recording of its spontaneous activity while simultaneously recording the fluorescence signals from the same cell. When sampling at a rate of 50Hz, the optical recording (Fig. 4.8D, red trace) faithfully followed the subthreshold fluctuations in the electrical recording (black trace). The optically recorded spikes are truncated as a result of undersampling the optical signal; sampling at a higher rate (722Hz) fully resolved the action potentials, but introduced a significant amount of sampling noise (Supplementary Appendix A, Fig. S4). Although the action potential was subsampled at just 50Hz, there is still a reliable transient in the optical trace that indicates the time when an action potential occurs, which is often what is needed.

The PeT-based VSDs show significant improvement in speed and accuracy compared with FRET-based VSDs previously used for leech recordings (Briggman et al. 2005; Taylor et al. 2003; Briggman & Kristan, 2006; Baca et al. 2008), which in turn had superseded electrochromic dyes (Cacciatore et al. 1999). The improvement in the recording of membrane potential fluctuations is not the result of a greater sensitivity (the $\Delta F/F$ for both the FRET and PeT dyes is about 10% per 100 mV in leech recordings), but to a greater SNR. The PeT-based VSD produces a much brighter signal, one that is well above the photon noise levels of the dye and the dark noise level of the camera. Increasing the concentration of the FRET-based dye does increase its SNR, but the consequent increase in the cell's capacitance (Fig. 4.7D) makes the dye useless for recording either action potentials or synaptic potentials. Tests of the toxicity and bleaching of the PeT-based VSD similar to those performed on the FRET-based dyes (Cacciatore et al. 1999) show that the PeT-based VSD has a slower rate of bleaching and is less toxic than the FRET-based dyes. Hence, considering all measures, the PeT-based VSD performs better than the FRET-based dyes (Briggman et al. 2010).

Discussion

Optimal VSDs would have large, fast responses to changes in voltage, place little or no capacitative load on the membrane, photobleach slowly with minimal photodynamic dam-

Table 4.1: Summary of VSD attributes

Attribute	Electrochromic	FRET	PeT
Nature of translocating charge	Electron	Libid soluble anion	Electron
Forward charge shift occurs when	Photon absorption	Membrane depolarization	Quenching
Reverse charge shift cocurs when	Photon emission or radiationless decay	Membrane repolarization	Electron-hole recombination
Fractional charge x Fraction of total voltage	~ 0.1	0.4 – 0.6	~ 0.5
Δ energy for 100mV ΔV	0.003 – 0.02eV	0.06eV	0.05eV
Comparison energy	Photon energy	kT	kT
Extended rigid fluorophore needed?	Yes	No	No
Use full ex/em band	No	Yes	Yes
Sensitivity $\Delta F/F$ per mV	Low	High	High
Speed	fs	ms – s	ms – μ s
Capacitative loading	None	Significant	None

age, and would be synthetically tractable for rational chemical modification and genetic targetability. We believe that the VF family of PeT-based probes surpass previous VSD classes by these criteria. The three PPV molecular wire, PeT-based molecules we tested (VF1.4.Cl, VF2.4.Cl, and VF2.1.Cl) exhibit good membrane staining and 20 – 27% $\Delta F/F$ per 100 – mV increases in fluorescence upon depolarization in HEK cells. These molecules possess the fast kinetics ($\tau_{ON/OFF} \ll 140\mu s$) and wavelength-independent voltage sensitivity consistent with a PeT mechanism for sensing voltage. Measurements of capacitance in leech neurons show that an insignificant amount of capacitance is placed on the membrane. The advantages of PeT-based dyes over both electrochromic and FRET-based methods for optical voltage sensing are described below and summarized in Table 4.1. A fourth technique, making use of genetically encoded voltage sensors, offers a promising method for optically monitoring voltage changes because the fluorescent proteins can be targeted to cells of interest, thereby increasing the SNR of the fluorescence response. In practice, however, fluorescent protein voltage sensors suffer from low sensitivity [0.5% (Ataka & Pieribone, 2002) to 10% $\Delta F/F$ per 100mV (Perron et al. 2009)], nonlinear responses (Siegel & Isacoff, 1997) and slow kinetics (tens to hundreds of milliseconds). Newer efforts have made use of proton translocation within bacterial rhodopsins (Kralj et al. 2011), but although these show large voltage sensitivities, the response time is still in the millisecond range, quantum efficiencies are very low, and their expression limited to prokaryotic systems. Voltage-driven translocation of ions through the membrane will generally add much more capacitance than electron translocation during transient excited states (Sjølson & Miesenböck, 2007).

VSDs using a PeT-based molecular wire approach should be highly sensitive. Because a full electronic charge travels through a substantial fraction of the transmembrane voltage (11 Å for VF1.4.Cl, 17 Å for VF2.4.Cl and VF2.1.Cl, or 37% and 57% of the 30 Å low-dielectric constant core of the plasma membrane) the change in driving force for PeT is large. For example, a 100-mV depolarization changes the PeT driving force by 0.05 eV (one electron \times half of 100-mV potential, or 0.05 V). Because PeT is a thermally controlled process, the value of 0.05 eV is large relative to the value of kT at 300 K (0.026 eV), yielding a large dynamic range between the rates of PeT at resting and depolarized potentials. FRET-based VSDs will have similar sensitivities; lipid-soluble mobile anions transverse distances calculated to be between 0.4 (González & Tsien, 1995) and 0.6 (Fernández et al. 1983) of the total membrane width, resulting in ΔG of ~ 0.05 eV for 100-mV depolarization, compared with a kT of 0.026 eV for the thermally activated process.

In contrast, electrochromic dyes have smaller ΔG values, 0.003 (Loew et al. 1978) to 0.02 (Kuhn & Fromherz, 2003) eV, and larger comparison energies. Because the interaction is a photochemically controlled process, the energy of the exciting photon is the comparison energy, which is 1.5-2 eV for dyes in the blue-to-green region of the spectrum. Therefore, PeT and FRET dyes have large changes in energy versus their comparison energy (0.05 eV vs. 0.026 eV), giving high sensitivities; electrochromic dyes have small changes compared with the excitation photon (0.003-0.02 eV vs. 2 eV), producing low voltage sensitivity.

The nature of the PeT mechanism also predicts that the kinetics of voltage sensing will be fast; forward electron transfer occurs on the nanosecond timescale as fluorescence is quenched, and back-electron transfer completes the cycle and occurs on a microsecond timescale or faster, meaning that the slow step, electron-hole recombination, finishes a full three orders-of-magnitude faster than an action potential. Electrochromic dyes display even faster kinetics, as forward charge shift occurs with absorbance, on the femtosecond time scale, and resolves itself upon emission of a photon, enabling these dyes to keep time with the fastest spiking neurons. FRET pair VSDs depend upon the migration of a lipophilic anion through an unstirred lipid bilayer and display kinetics in the millisecond-to-second time regime, limiting their application to monitoring only slow transients.

Because PeT shuttles an electron across the membrane and back on a microsecond or faster timescale, driven by photons rather than membrane potential changes, no capacitive loading should be observed. The same holds true for electrochromic dyes, which transfer electrons on even faster time scales. One disadvantage of electrochromic dyes is that they require

the entire voltage-sensing chromophores to be rigid to enable π orbital overlap, quantum yield, efficient charge transfer, and maximization of voltage sensitivity (Hubener et al. 2003). Such rigidity hinders synthesis and water solubility and may explain why electrochromic dyes are not improved by lengthening their chromophores. PeT probes do not require the entire molecular wire to be rigidly coplanar, and synthesis of longer wires is feasible.

PeT-sensing allows the entire emission spectrum to be used for monitoring voltage, because the quenching mechanism alters the Φ_{Fl} , decreasing the brightness of the dye, and does not shift the wavelength as do electrochromic methods. Because photons are not wasted, this allows lower intensity light to be used in experiments, reducing phototoxicity and increasing the duration of experimental procedures. The performance of electrochromic dyes has plateaued over four decades of development. Excitation at the far-red edge of the spectrum gives voltage sensitivities ranging from -35% to -52% $\Delta F/F$ per 100 mV; however, at the edge of the spectrum, the intensity is far below the peak and the voltage response becomes nonlinear (Kuhn et al. 2004).

Several limitations of the VF dyes remain to be addressed. VF derivatives are not yet genetically targetable. The sensors are readily taken up by the cell membranes of all tissue, increasing nonresponsive background fluorescence and decreasing the SNR. For heterogeneous preparations, such as intact leech ganglia and brain slices, this becomes an increasingly important issue, and one method to address this concern is through the genetic targeting of VSDs. VF sensors lend themselves to chemical derivatization, and efforts are underway to modify VF probes for targeting to genetically defined circuits of neurons.

Another drawback is that VF PeT sensors are not as sensitive to voltage as hoped. Our first derivatives show sensitivities ranging from 20-27% $\Delta F/F$ per 100 mV, and the most sensitive of existing electrochromic dyes exhibit -28% $\Delta F/F$ per 100-mV sensitivities in the linear range (Kuhn & Fromherz, 2003). Although it is encouraging that the first derivatives display sensitivities on a level approaching the most sensitive electrochromic dyes, we believe ample chemical space exists for improving the voltage sensitivity of molecular wire platforms. Because the voltage sensitivity is controlled by PeT, the efficiency of PeT can be rationally tuned (Ueno et al. 2004) by altering the electron affinities of the donor, wire, and acceptor to maximize the fluorescence turn-on in response to depolarizations. Additionally, extending the molecular wire to span an even greater distance across the plasma membrane should increase sensitivity as the transferred electron samples more of the electric field. The modular nature of the VF synthesis allows for rapid interchange of coupling partners to quickly assemble and assess the voltage

sensitivity of an array of compounds.

In summary, we present a unique method for monitoring voltage in neurons based on the voltage-sensitive PeT from an electron-rich donor to fluorescent reporter attached via a membrane-spanning molecular wire. The VF family of sensors have large, linear, turn-on fluorescence responses to depolarizing steps (20-27% $\Delta F/F$ per 100 mV), fast kinetics ($\tau \ll 140\mu s$), and negligible capacitive loading. VF2.4.Cl can detect and resolve evoked action potentials in primary culture hippocampal neurons, and VF2.1.Cl can monitor spontaneous spiking and synaptic potentials in leech Retzius cells with sensitivity and time-course essentially identical to the recorded electrophysiology signal. VF sensors provide a practical alternative to currently available VSDs, and future derivatives of the molecular wire platform will increase our ability to optically monitor the temporal and spatial dynamics of neuronal activity in defined circuits of neurons.

Methods

Imaging, electrophysiology, cell culture, leech imaging and electrophysiology, and data analysis methods are available in Supplementary Appendix A. Theoretical considerations of capacitive load are included in Supplementary Appendix A. Dyes were synthesized using standard synthetic procedures detailed in Supplementary Appendix A.

Acknowledgements

Part of this chapter was published in Miller, E.W., Lin, J.Y., Frady, E.P., Steinbach, P.A., Kristan, W.B., Tsien, R.Y. (2012). Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires. *Proc. Natl. Acad. Sci. USA.* 109(6): 2114-2119.

References

- Adams, S.R. (2010) How calcium indicators work. *Cold Spring Harbor Protocols.* 2010 pdb-top70.
- Akemann, W., Lundby, A., Mutoh, H., Knöpfel, T. (2009) Effect of voltage sensitive fluorescent proteins on neuronal excitability. *Biophys J.* 96:3959-3976.
- Ataka, K., Pieribone, V.A. (2002) A genetically targetable fluorescent probe of channel gating with rapid kinetics. *Biophys J.* 82:509-516.

- Baca, S.M., Marin-Burgin, A., Wagenaar, D.A., Kristan, W.B., Jr (2008) Widespread inhibition proportional to excitation controls the gain of a leech behavioral circuit. *Neuron*. 57:276-289.
- Bagnoli, P., Brunelli, M., Magni, F. (1975) The neuron of the fast conducting system in *Hirudo medicinalis*: identification and synaptic connections with primary afferent neurons. *Arch. Ital. Biol.* 113:21-43.
- Bradley, J., Luo, R., Otis, T.S., DiGregorio, D.A. (2009) Submillisecond optical reporting of membrane potential in situ using a neuronal tracer dye. *J. Neurosci.* 29:9197-9209.
- Briggman, K.L., Abarbanel, H.D., Kristan Jr., W.B. (2005) Optical imaging of neuronal populations during decision-making. *Science*. 307:896-901.
- Briggman, K.L., Kristan Jr., W.B. (2006) Imaging dedicated and multifunctional neural circuits generating distinct behaviors. *J. Neurosci.* 26:10925-10933.
- Briggman, K.L., Kristan Jr., W.B., González, J.E., Kleinfeld, D., Tsien, R.Y. (2010) Monitoring integrated activity of individual neurons using FRET-based voltage-sensitive dyes. In: Canepari M, Zecevic D, editors. *Membrane Potential Imaging in the Nervous System: Methods and Applications*. New York: Springer; pp. 61-70.
- Brodgheuer, P.D., Friesen, W.O. (1986) Initiation of swimming activity by trigger neurons in the leech subesophageal ganglion. I. Output pathways of Tr1 and Tr2. *J. Comp. Physiol.* 159:485-502.
- Cacciatore, T.W., Brodgheuer, P.D., González, J.E., Jiang, T., Adams, S.R., Tsien, R.Y., Kristan Jr., W.B., Kleinfeld, D. (1999) Identification of neural circuits by imaging coherent electrical activity with FRET-based dyes. *Neuron*. 23:449-459.
- Chanda, B., Blunck, R., Faria, L.C., Schweizer, F.E., Mody, I., Bezanilla, F. (2005) A hybrid approach to measuring electrical activity in genetically specified neurons. *Nat Neurosci.* 8:1619-1626.
- Crisp, K.M., Mesce, K.A. (2004) A cephalic projection neuron involved in locomotion is dye coupled to the dopaminergic neural network in the medicinal leech. *J. Exp. Biol.* 207: 4535-4542.
- Davis W.B., Svec W.A., Ratner M.A., Wasielewski M.R. (1998) Molecular-wire behaviour in p-phenylenevinylene oligomers. *Nature*. 396:60-63.
- Esch, T., Mesce, K.A., Kristan, W.B. (2002) Evidence for sequential decision making in the medicinal leech. *J Neurosci.* 22:11045-11054.
- Fernández J.M., Taylor R.E., Bezanilla F. (1983) Induced capacitance in the squid giant axon. Lipophilic ion displacement currents. *J Gen Physiol.* 82:331-346.

- Fluhler E., Burnham V.G., Loew L.M. (1985) Spectra, membrane binding, and potentiometric responses of new charge shift probes. *Biochemistry*. 24:5749-5755.
- Garner L.E., Park J., Dyar S.M., Chworos A., Sumner J.J., Bazan G.C. (2010) Modification of the optoelectronic properties of membranes via insertion of amphiphilic phenylenevinylene oligoelectrolytes. *J Am Chem Soc*. 132:10042-10052.
- González, J.E., Tsien, R.Y. (1997) Improved indicators of cell membrane potential that use fluorescence resonance energy transfer. *Chem Biol*. 4:269-277.
- González, J.E., Tsien, R.Y. (1995) Voltage sensing by fluorescence resonance energy transfer in single cells. *Biophys J*. 69:1272-1280.
- Gray, J., Lissmann, H.W., Pumphrey, R.J. (1938) The mechanism of locomotion in the leech *Hirudo medicinalis* Ray. *J. Exp. Biol*. 15:408-430
- Grinvald, A. (1983) Fluorescence monitoring of electrical responses from small neurons and their processes. *Biophys J*. 42:195-198.
- Grynkiewicz, G., Poenie, M., Tsien, R.Y. (1985) A new generation of Ca²⁺ indicators with greatly improved fluorescence properties. *J Biol Chem*. 260:3440-3450.
- Heim, N., Griesbeck, O. (2004) Genetically encoded indicators of cellular calcium dynamics based on troponin C and green fluorescent protein. *J Biol Chem*. 279:14280-14286.
- Horikawa, K., Yamada, Y., Matsuda, T., Kobayashi, K., Hashimoto, M., Matsu-ura, T., Miyawaki, A., Michikawa, T., Mikoshiba, K., Nagai, T. (2010) Spontaneous network activity visualized by ultrasensitive Ca(2+) indicators, yellow Cameleon-Nano. *Nat Methods*. 7:729-732.
- Hubener, G., Lambacher, A., Fromherz, P. (2003) Anellated hemicyanine dyes with large symmetrical solvatochromism of absorption and fluorescence. *J Phys Chem B*. 107:7896-7902.
- Jacobs, J.M., Meyer, T. (1997) Control of action potential-induced Ca²⁺ signaling in the soma of hippocampal neurons by Ca²⁺ release from intracellular stores. *J Neurosci*. 17:4129-4135.
- Kralj, J.M., Hochbaum, D.R., Douglass, A.D., Cohen, A.E. (2011) Electrical spiking in *Escherichia coli* probed with a fluorescent voltage-indicating protein. *Science*. 333:345-348.
- Kristan Jr., W.B. (1982). Sensory and motor neurons responsible for the local bending response in leeches. *J. Exp. Biol*. 96:161-180.
- Kristan Jr., W.B., Calabrese, R.L., Friesen, W.O. (2005) Neuronal control of leech behavior. *Prog. in Neurobiology*. 76:279-327.
- Kuhn B., Fromherz P. (2003) Anellated hemicyanine dyes in a neuron membrane: Molecular Stark effect and optical voltage recording. *J Phys Chem B*. 107:7903-7913.

- Kuhn, B., Fromherz, P., Denk, W. (2004) High sensitivity of Stark-shift voltage-sensing dyes by one- or two-photon excitation near the red spectral edge. *Biophys J.* 87:631-639.
- Laverack, M.S. (1968) Mechanoreceptors, photoreceptors and rapid conduction pathways in the leech, *Hirudo Medicinalis*. *J. Exp. Biol.* 50:129-140.
- Li, L.S. (2007) Fluorescence probes for membrane potentials based on mesoscopic electron transfer. *Nano Lett.* 7:2981-2986.
- Lockery, S.R., Kristan Jr., W.B., (1990a) Distributed processing of sensory information in the leech. I. Input-output relations of the local bending reflex. *J. Neurosci.* 10(6):1811-1815.
- Lockery, S.R., Kristan Jr., W.B. (1990b). Distributed processing of sensory information in the leech. II. Identification of interneurons contributing to the local bending reflex. *J. Neurosci.* 10(6):1816-1829.
- Loew, L.M., Bonneville, G.W., Surow J. (1978) Charge shift optical probes of membrane potential. *Theory. Biochemistry.* 7:4065-4071.
- Macagno, E.R. (1980) Number and distribution of neurons in leech segmental ganglia. *J. Comp. Neurol.* 15:283-302.
- Macagno, E.R., Peinado, A., Stewart, R.R. (1986) Segmental differentiation in the leech nervous system: Specific phenotypic changes associated with ectopic targets. *Proc. Natl. Acad. Sci. USA.* 83: 2746-2750.
- Magni, F., Pellegrino, M. (1978) Neural mechanism underlying the segmental and generalized cord shortening reflexes in the leech. *J. Comp. Physiol. A* 124:339-351.
- Mank, M, Santos, A.F., Drenth, S., Mrcic-Flogel T.D., Hover, S.B., Stein, V., Hendel, T., Reiff, D.F., Levelt, C., Borst, A., Bonhoeffer, T., Hübener, M., Griesbeck, O. (2008) A genetically encoded calcium indicator for chronic in vivo two-photon imaging. *Nat Methods.* 5:805-811.
- Minta, A., Kao, J.P., Tsien, R.Y. (1989) Fluorescent indicators for cytosolic calcium based on rhodamine and fluorescein chromophores. *J Biol Chem.* 264:8171-8178.
- Mistick, D. (1974) Rohde's fiber: a septate axon in the leech. *Brain Res.* 74: 342-348.
- Mistick, D. (1978) Neurones in the leech that facilitate an avoidance behavior following nearfield water disturbances. *J. Exp. Biol.* 75:1-23.
- Miyawaki, A., Llopis, H., Heim, R., McCaffery, J.M., Adams, J.A., Ikura, M., Tsien, R.Y. (1997) Fluorescent indicators for Ca²⁺ based on green fluorescent proteins and calmodulin. *Nature.* 388:882-887.

- Montana, V., Farkas, D.L., Loew, L.M. (1989) Dual-wavelength ratiometric fluorescence measurements of membrane potential. *Biochemistry*. 28:4536-4539.
- Muller, K.J., Nicholls, J.G., Stent, G.S. (1981) *Neurobiology of the Leech*. Cold Spring Harbor Laboratory, Cold Spring Harbor, New York.
- Mullins, O.J., Friesen, W.O. (2012) The brain matters: effects of descending signals on motor control. *J. Neurophysiol.* 107(10): 2730-2741.
- Nagai, T., Sawano, A., Park, E.S., Miyawaki, A. (2001) Circularly permuted green fluorescent proteins engineered to sense Ca²⁺ Proc. Natl. Acad. Sci. USA. 98:3197-3202.
- Nicholls, J.G., Baylor, D.A. (1968) Specific modalities and receptive fields of sensory neurons in CNS of the leech. *J. Neurophysiol.* 31:740-756.
- Norris, B.J., Calabrese, R.L., 1987. Identification of motor neurons that contain a FMRFamide-like peptide and the effects of FMRFamide on longitudinal muscle in the medicinal leech, *Hirudo medicinalis*. *J. Comp. Neurol.* 266, 95-111.
- Ort, C.A., Kristan Jr., W.B., Stent, G.S. (1974) Neuronal control of swimming in the medicinal leech. II. Identification and connections of motor neurons. *J. Comp. Physiol.* 94:121-154.
- Palmer, A.E., Giacomella, M., Kortemme, T., Hires, S.A., Lev-Ram, V., Baker, D., Tsien, R.Y. (2006) Ca²⁺ indicators based on computationally redesigned calmodulin-peptide pairs. *Chem Biol.* 13:521-530.
- Palmer, C.R., Barnett, M.N., Copado, S., Gardezy, F., Kristan Jr., W.B. (2014). Multiplexed modulation of behavioral choice. *J. Exp. Bio.* 217: 2963-2973.
- Perron, A., Mutoh, H., Akermann, W., Gautam, S.G., Dimitrov, D., Iwamoto, Y., Knöpfel, T. (2009) Second and third generation voltage-sensitive fluorescent proteins for monitoring membrane potential. *Front Mol Neurosci.* 2:5.
- Peterka, D.S., Takahashi, H., Yuste, R. (2011) Imaging voltage in neurons. *Neuron.* 69:9-21.
- Poenie, M., Alderton, J., Tsien, R.Y., Steinhardt, R.A. (1985) Changes of free calcium levels with stages of the cell division cycle. *Nature.* 315:147-149.
- Puhl, J.G., Mesce, K.A. (2010) Keeping it together: mechanisms of intersegmental coordination for a flexible locomotor behavior. *J. Neurosci.* 30(6): 2373-2383.
- Ross, W.N., Arechiga, H., Nicholls, J.G. (1987) Optical recording of calcium and voltage transients following impulses in cell bodies and processes of identified leech neurons in culture. *J Neurosci.* 7:3877-3887.
- Salzberg B.M., Grinvald A., Cohen L.B., Davila H.V., Ross W.N. (1977) Optical recording of

neuronal activity in an invertebrate central nervous system: Simultaneous monitoring of several neurons. *J Neurophysiol.* 40:1281-1291.

Sawada, M., Wilkinson, J.M., McAdoo, D.J., Coggeshall, R.E. (1976) The identification of two inhibitory cells in each segmental ganglion of the leech and studies on the ionic mechanism of the inhibitory junctional potential produced by these cells. *J. Neurobiol.* 7:435-445.

Scanziani M., Häusser M. (2009) Electrophysiology in the age of light. *Nature.* 461:930-939.

Shaw, B.K., Kristan Jr., W.B. (1995) The whole-body shortening reflex of the medicinal leech: motor pattern, sensory basis, and interneuronal pathways. *J. Comp. Physiol. A* 177:667-681.

Siegel M.S., Isacoff E.Y. (1997) A genetically encoded optical probe of membrane voltage. *Neuron.* 19:735-741.

de Silva, A.P., Gunaratne, H.Q.N., Habib-Jiwan, J.L., McCoy, C.P., Rice, T.E., Soumillion, J.P. (1995) New fluorescent model compounds for the study of photoinduced electron transfer: The influence of a molecular electric field in the excited state. *Angew Chem Int Ed Engl.* 34:1728-1731.

Sjulson, L., Miesenböck, G. (2007) Optical recording of action potentials and other discrete physiological events: A perspective from signal detection theory. *Physiology (Bethesda)* 22:47-55.

Sjulson, L., Miesenböck, G. (2008) Rational optimization and imaging in vivo of a genetically encoded optical voltage reporter. *J Neurosci.* 28:5582-5593.

Taylor, A.L., Cottrell, G.W., Kleinfeld, D., Kristan, W.B., Jr (2003) Imaging reveals synaptic targets of a swim-terminating neuron in the leech CNS. *J. Neurosci.* 23:11402-11410.

Thompson, W.J., Stent, G.S. (1976) Neuronal control of heartbeat in the medicinal leech. I. Generation of the vascular constriction rhythm by heart motor neurons. *J. Comp. Physiol.* 111:261-279.

Tian, L., Hires, S.A., Mao, T., Huber, D., Chiappe, M.E., Chalasani, S.H., Petreanu, L., Akerboom, J., McKinney, S.A., Schreiter, E.R., Bargmann, C.I., Jayaraman, V., Svoboda, K., Looger, L.L. (2009) Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat. Methods.* 6:875-881.

de la Torre, G., Giacalone, F., Segura, J.L., Martin, N., Guldi, D.M. (2005) Electronic communication through pi-conjugated wires in covalently linked porphyrin/C60 ensembles. *Chemistry.* 11:1267-1280.

Tsien, R.Y., Rink, T.J., Poenie, M. (1985) Measurement of cytosolic free Ca²⁺ in individual small cells using fluorescence microscopy with dual excitation wavelengths. *Cell Calcium.* 6:145-157.

Ueno, T., Urano, Y., Setsukinai, K., Takakusa, H., Kojima, H., Kikuchi, K., Ohkubo, K., Fukuzumi, S., Nagano, T. (2004) Rational principles for modulating fluorescence properties of fluorescein. *J. Am. Chem. Soc.* 126:14079-14085.

Wang, D., Zhang, Z., Chanda, B., Jackson, M.B. (2010) Improved probes for hybrid voltage sensor imaging. *Biophys. J.* 99:2355-2365.

Weeks, J.C. (1982a) Synaptic basis of swim initiation in the leech. I. Connections of a swim-initiating neuron cell 204 with motor neurons and pattern generating "oscillator" neurons. *J. Comp. Physiol.* 148:265-279.

Weeks, J.C. (1982b) Synaptic basis of swim initiation in the leech. II. A pattern-generating neuron cell 208, which mediates motor effects of swim-initiating neurons. *J. Comp. Physiol.* 148:265-279.

5 The Imaging Computational Microscope

Abstract

The Imaging Computational Microscope (ICM) is a suite of computational tools for automated analysis of functional imaging data that runs under the cross-platform MATLAB environment (The Mathworks, Inc.). ICM uses a semi-supervised framework, in which at every stage of analysis computers handle the routine work, which is then refined by user intervention. The main functionality of ICM is built upon automated extraction of component signals from imaging data, segmentation and clustering of component signals, and feature computation and visualization. Analysis of imaging data through ICM can also be done in real-time, enabling guided imaging experiments. ICM is built using MATLAB's object-oriented class design, which allows the tool to be used both as a graphic user interface (GUI) as well as in custom scripts. ICM is freely available under the GNU public license for non-commercial use and open source development, together with sample data, user tutorials and extensive documentation.

Introduction

Large-scale recordings of neural data, such as calcium imaging or multi-unit electrode arrays, are necessary to understand the function of intact nervous systems. However, most experimentalists still rely on manual analysis techniques to synthesize and make sense of these “big data” experiments. As the number of neurons that can be simultaneously recorded continues to increase, manual analysis techniques will not be able to scale with the data and automated techniques will be necessary.

The Imaging Computational Microscope (ICM) is a computational tool for analyzing functional neural imaging data, such as calcium or voltage-sensitive dye imaging. The main functionality of ICM is built upon automated extraction of component signals from imaging data using Principal and Independent Component Analysis (PCA, ICA respectively; Mukamel

et al. 2009; Hill et al. 2010), segmentation and clustering of component signals, and feature computation and visualization. These stages enable the experimentalist to start with raw imaging data and rapidly analyze large-scale data and produce high-level visualizations. The full set of computational analysis stages can be executed on the scale of minutes, which enables ICM to be used in real-time to guide experiments.

Methods and Results

Overview of Main Functions

The Imaging Computational Microscope consists of several stages of analysis, each of which benefits from user interaction and inspection. Each stage is controlled using a series of tab menus (Fig. 5.1 R4), and ICM updates its displays based on the current stage and menu that is open.

Data. In the data tab of the ICM, the user can browse the raw data, draw ROIs, perform motion correction, and enable concatenated-trial ICA (ctICA). In this stage, the Roi Editor (Fig. 5.1 R1) shows the raw imaging data, and ROIs can be manually drawn by clicking and dragging the mouse over the desired region of the image. The signals extracted by the ROI are plotted in the Data plot (Fig. 5.1 R6). These ROIs are stored as “ROI sets” and can be added or deleted using the ROI control list (Fig. 5.1 R3) as well as saved into and loaded from files. Multiple trials of data can be loaded into ICM and these can be browsed using the trial control list (Fig. 5.1 R2).

Pre-Processing. In the pre-processing tab, the user can down-sample and smooth imaging data (see R4 in Fig. 5.2). The ROI Editor will display the pre-processed data, and the extracted ROI signals of the pre-processed data will be plot in the Data plot (Fig. 5.1 R6).

PCA. In the PCA tab, the user can run Principal Components Analysis (PCA) on the imaging data, and browse the components. The principal components (PCs) are used to reduce the dimensionality of the imaging data and remove noise. When applied to imaging data, PCA produces a map (i.e. the spatial locations) and a source (i.e. the time-series) for each principal component. The PCs are typically combinations of many cellular signals or artifacts. Each PC map will be displayed as a frame in the ROI Editor, which can be browsed by clicking through the frames in the GUI. The corresponding source for each PC is plotted in the Component plot (Fig. 5.1 R7).

ICA. In the ICA tab, the user can run and visualize Independent Components Analysis (ICA) of the imaging data (Bell & Sejnowski, 1995). ICA is run on a subset of the PCs, which is

set by the user in the interface. The ICA algorithm also produces a map and a source for each of the independent components (ICs). The maps can be browsed as frames in the ROI Editor and the sources are plotted in the Component plot. The ICA algorithm has been shown to extract individual cellular signals, but it also extracts artifactual components. These must be viewed and sorted manually using the Remove box (see R4 in Fig. 5.4). Further post-processing of the extracted components can also be performed using the controls in the ICA Post-Processing panel.

Segmentation. In the Segmentation tab, the user can segment the ICs and generate ROIs. The ROI Editor displays the results of the segmentation for each IC. The ROIs generated from segmentation are locked to their IC and can be used to quickly browse through all of the ICs.

Visualization. In the Visualization tab, ICM displays color-maps based on features computed from the IC sources. The user can manipulate the visualization settings using the controls in the Visualization Settings panel. The visualizations can be crafted using two built-in visualization algorithms, which are controlled in the Visualization panel (5.1 R8). The algorithms assign each IC source 2-3 coefficients, based on some aspect of the component's activity. These coefficients are used to derive colors and create a colored activity map. Outside functions can also be used to create visualizations by assigning a color to each IC and setting these colors using the function `set_viz_colors`.

We will illustrate the use and utility of ICM in more detail with two example walk-throughs of imaging data analysis.

PCA-ICA extraction of Calcium Imaging Data

ICM can load data from a .tiff or .mat file using the ICM menu in the interface (Fig. 5.1), or data can be loaded into ICM through matlab functions `set_data` or `add_trial`. Data is loaded as a trial, and trials are managed using the trial management list (Fig. 5.1 R2).

The data can be browsed using the ROI Editor (Fig. 5.1 R1), which allows the user to draw manual ROIs and look through the frames of the imaging data. ROIs can be drawn by simply clicking and dragging the mouse on the image in the ROI editor. The signals from manually drawn ROIs are plotted in the Data plot (Fig 5.1 R6). Multiple sets of both manually and automatically generated ROIs can be managed using the ROI manager (Fig 5.1 R3).

Each stage of the analysis is set and controlled with a tab menu in the stage tab panel (Fig 5.1 R4). ICM displays different information depending on which stage of the analysis is

being viewed. The current stage of the analysis is displayed beneath the stage tab panel (Fig 5.1 R5), which also indicates when ICM is busy computing for each stage.

In the next step, the preprocessing tab is opened, and the smooth window and down sampling values are set (Fig. 5.2). The smooth window averages all pixels within a moving $M \times N \times T$ window (rows, columns, time). The down sampling only keeps every $M \times N \times T$ pixels, producing a smaller imaging data set. For this example, the original image data is $512 \times 512 \times 1600$, but this is too large for the PCA-ICA analysis to be run on a desktop computer. We used a $4 \times 4 \times 2$ smooth window and $4 \times 4 \times 2$ down sampling to reduce the image size to $128 \times 128 \times 800$. When the “Pre Process” button is pressed the original data is pre-processed and the pre-processed data is displayed in the ROI Editor. The original data can still be viewed in the data tab, and the pre-processing can be removed by clicking the “Reset” button in the Data tab.

Next, the principal components are computed. Every pixel in the preprocessed imaging data set is arranged in a matrix, where the rows are each pixels, and the columns are the values of the pixels over time. When the user clicks the “Run PCA” button (Fig. 5.3), the principal components of this matrix are then computed, which decomposes the imaging data into several components, each of which has a “source” and a “score”. Sources are the time series of the extracted components, and the scores indicate the coefficient of the source for each pixel. The scores are rearranged back into an image to produce a “map”, which shows the spatial locations from which the sources are produced. The principal components are typically combinations of cellular signals, and do not reveal individual cells.

The principal components can be browsed in the ROI editor with the PCA tab open. Each frame in the ROI Editor shows the map of a different principal component, and the source of each component is plotted in the Component Plot. Several more example principal components can be seen in Figure 5.6A.

To reveal individual cells, the independent components are then computed from a subset of the principal components. Typically the top N principal components are used, where N is slightly larger than the number of neurons being recorded from. This can be set using the “PCs” edit box in the ICA tab (Fig. 5.4). The ICA algorithm will attempt to find the same number of independent components as principal components included, and so there should be at least as many PCs used as cells. Typically, more PCs are needed than cells, because many independent components are extracted that correspond to motion, background, bleaching, or other artifacts.

Once the PCs are chosen and “Run ICA” is clicked, ICM computes the independent components using the fastica algorithm by default. Further changes to the fastica settings can

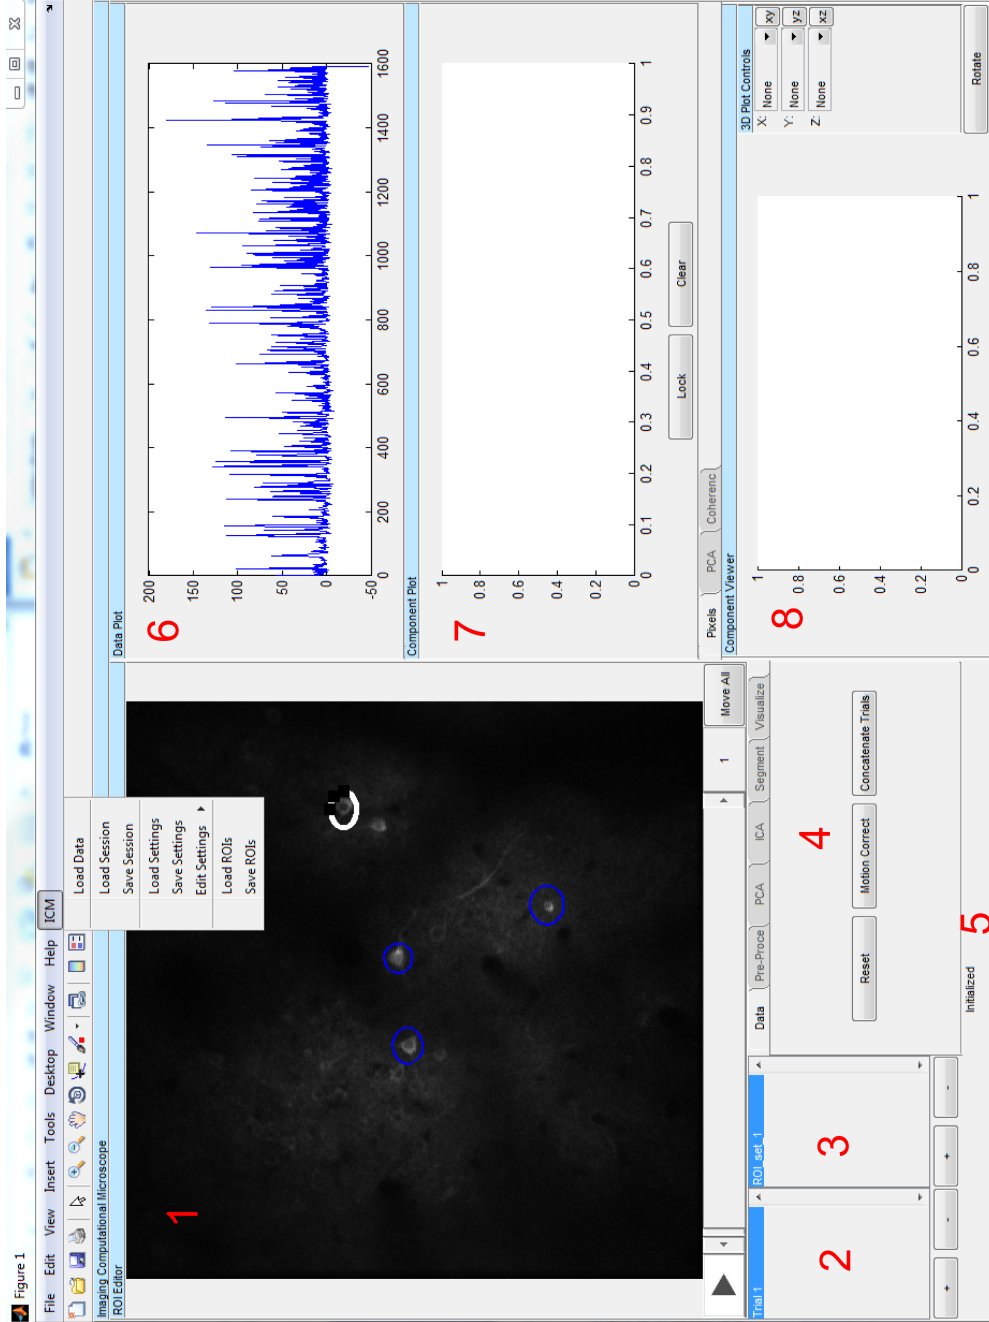


Figure 5.1: ICM screenshot with labeled regions.

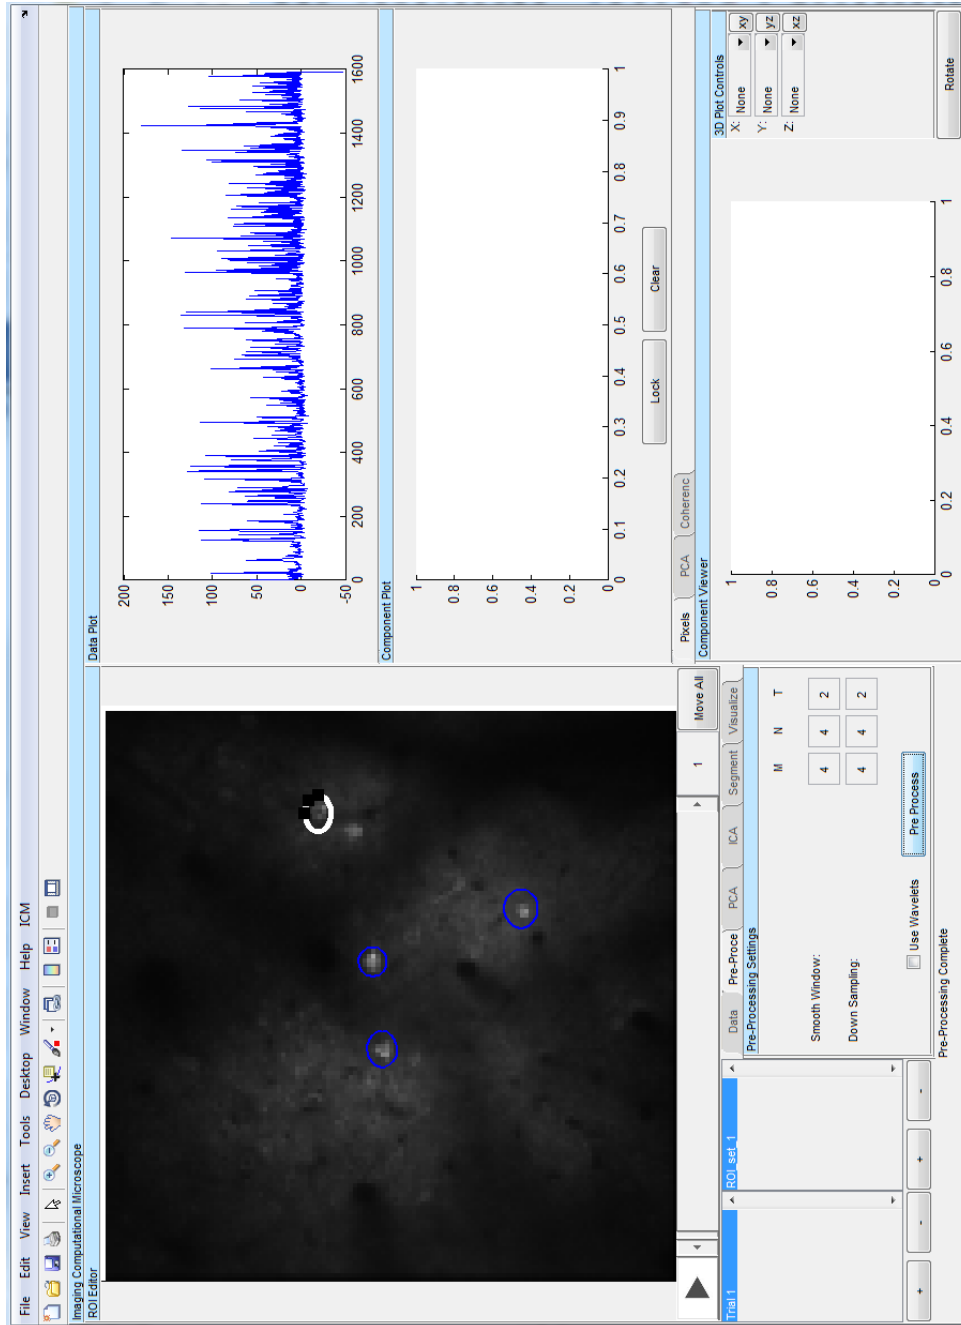


Figure 5.2: ICM screenshot in pre-processing stage.

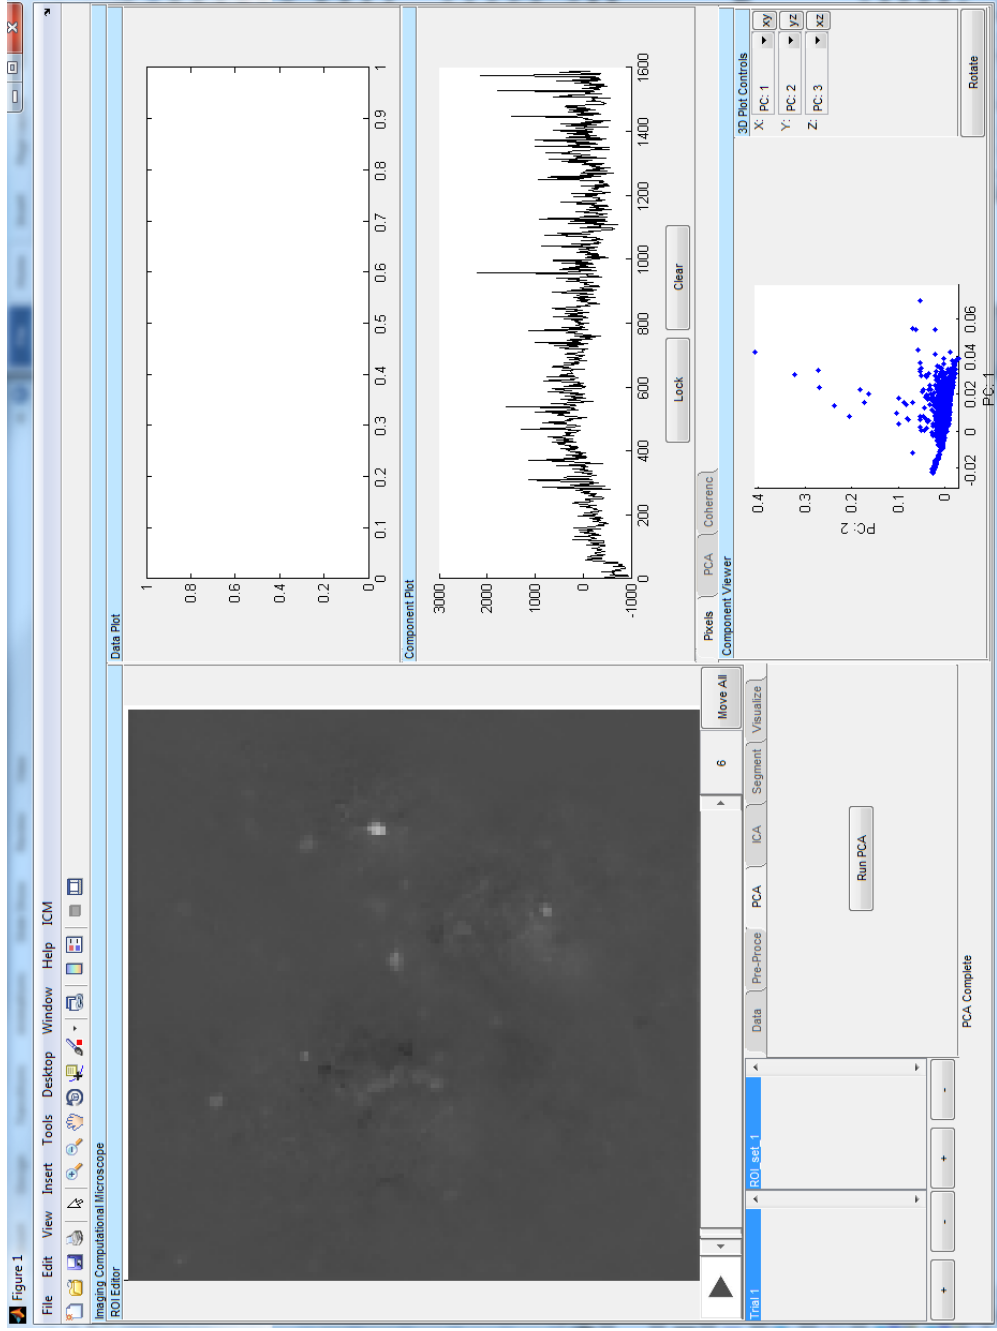


Figure 5.3: ICM screenshot in PCA stage.

be made through a settings struct that can be changed programatically (see Documentation), as well as choosing different ICA algorithms (e.g. infomax (Bell & Sejnowski, 1995) and spatio-temporal ICA (Mukamel et al. 2009)). Like PCA, ICA also produces sources and scores, where the sources correspond to the independent signals and the scores describe which pixels are contributing to the source. The scores are rearranged back into an image to produce a map, and the maps show the spatial locations of the independent components.

Several example independent components are shown in Figure 5.6B and C. The ICA algorithm pulls out both components that are cellular signals (Fig. 5.6B) as well as components that are artifacts (Fig. 5.6C). These must be sorted manually using the interface, and components which are artifactual can be removed from further analysis by adding them to the “Remove” edit box. Further post-processing can be performed on the ICs using the tools in the “ICA Post-Processing” panel.

Regions-of-interest can then be automatically generated from the ICA maps. To compute the regions of interest, ICM uses a threshold to segment the IC maps. In the Segment tab (Fig. 5.5), the threshold level and amount of down-sampling are set for the segmentation algorithm. When the “Segment ICs” button is pressed, binary masks are created for each IC, which are displayed in the ROI Editor when the Segment tab is open. Each contiguous region of the binary mask is then matched with the best fitting oval to produce the ROI. The ROIs produced are saved in the ROI manager (Fig. 5.1 R3).

ROIs produced automatically remain linked with their corresponding IC and this allows the user to browse the ICs by clicking on the corresponding ROI. The ROIs also show the user the spatial localization and extent of all of the ICs.

Each IC can potentially have multiple ROIs, because there can be multiple spatially isolated regions for a single IC. In this example data, several cell soma’s are slightly misaligned with the image plane, but their dendritic branches are still recorded (see IC 26 and 30 at bottom of Fig. 5.6B). Because of this, the segmentation algorithm breaks up these cells into multiple ROIs (Fig. 5.6B, right).

A major advantage of the PCA-ICA extraction is that this algorithm does not depend on spatial-localization to extract component signals. Calcium signals from cells or axons that do not have a single localized spatial region would be virtually impossible to extract if only ROIs were used. The ICA component decomposition does not depend on spatial localization, which allows for clear signals to be extracted from out of focus cells or long axons (see bottom two examples in Fig. 5.6B). This suggests that many cells may be missed entirely when relying

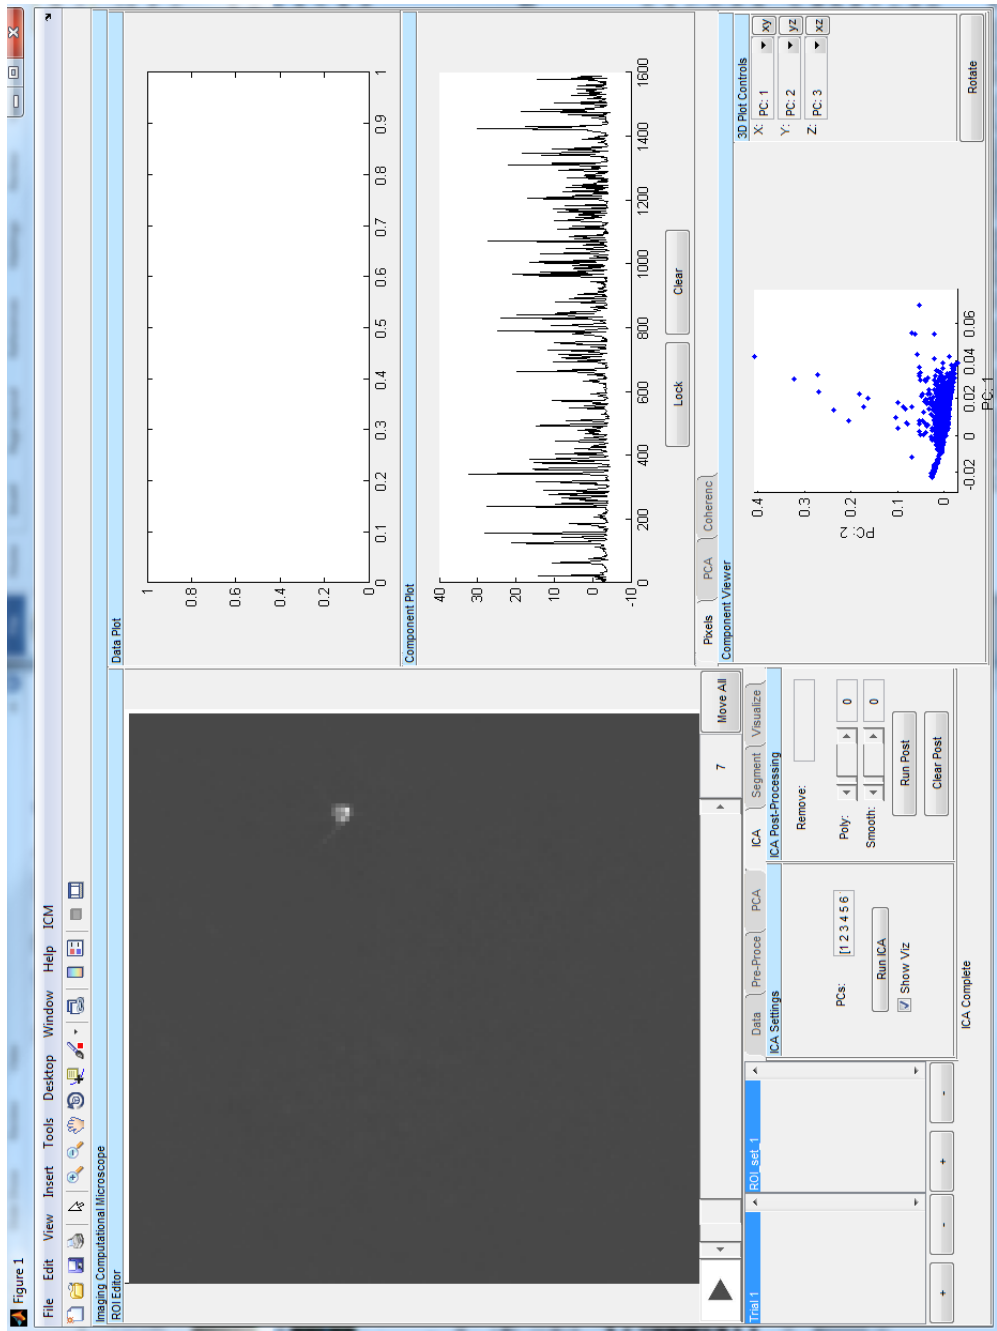


Figure 5.4: ICM screenshot in ICA stage.

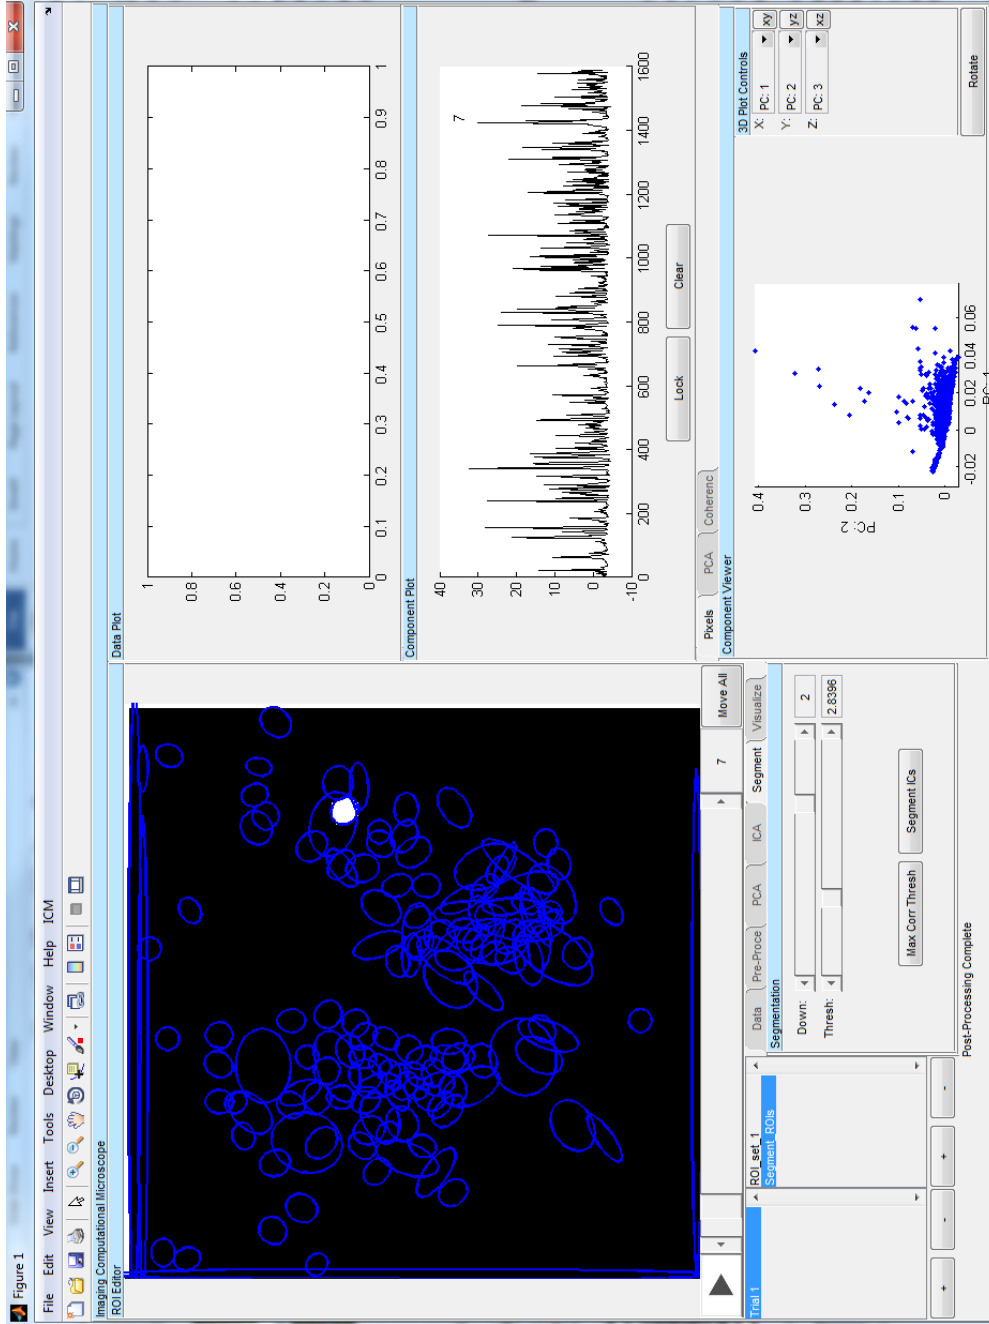


Figure 5.5: ICM screenshot in segment stage.

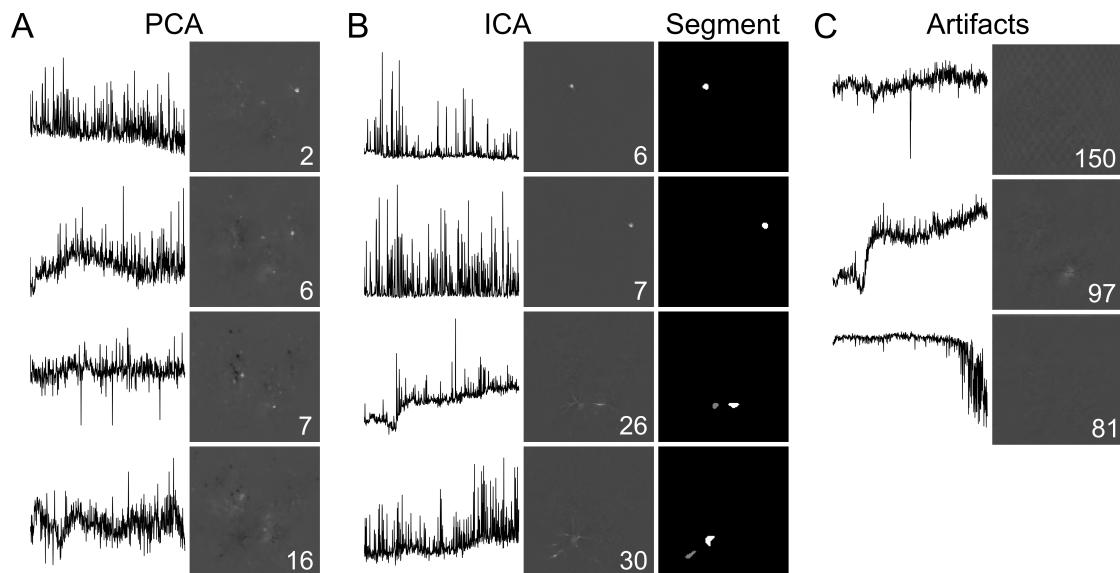


Figure 5.6: Example components. (A) Four example principal components are shown. The PCs are typically mixtures of cells and artifacts. Sources (left) show calcium spikes, anti-spikes and background fluctuations, and maps (right) show a mixture of cells and background. (B) Four example independent component from cellular sources are shown. The bottom two components show the dendrites of two neurons that are out of the image plane (middle), but the Calcium spikes of these cells (left) can still be captured using the PCA-ICA extraction. The results of segmentation for each IC is shown on the right. (C) Three example independent components that are artifacts. The top shows corrugation due to image sampling artifact, the middle is background fluorescence, and the bottom is an edge artifact.

on ROI methods and that the ICA algorithm can get much higher signal-to-noise ratio under certain imaging conditions. Further, ICA can separate components that have overlapping spatial locations. For example, IC 150 (Fig. 5.6C) shows clear corrugations due to artifacts of the imaging acquisition system. These striations cover the entire image, but ICA can separate this artifact from the cellular components because the pixels share statistical patterns caused by the artifact. It would be impossible for ROI methods to separate spatially overlapping components.

Finally, visualizations are created in the visualization stage. Two simple visualization systems are built into ICM, and here we will illustrate the principal component visualization of the IC sources. The PCA component viewer can be opened in the visualization control tabs (Fig 5.1. R8), and this allows the user to select three principal components to visualize. Each independent component has coefficients in the principal component space, and the user selects which three dimensions to view. The component maps are then colored based on the coefficients of the ICs given the dimensions chosen, and these are overlaid on the image data to create a visualization of activity. The settings of the visualization can be manipulated using the controls

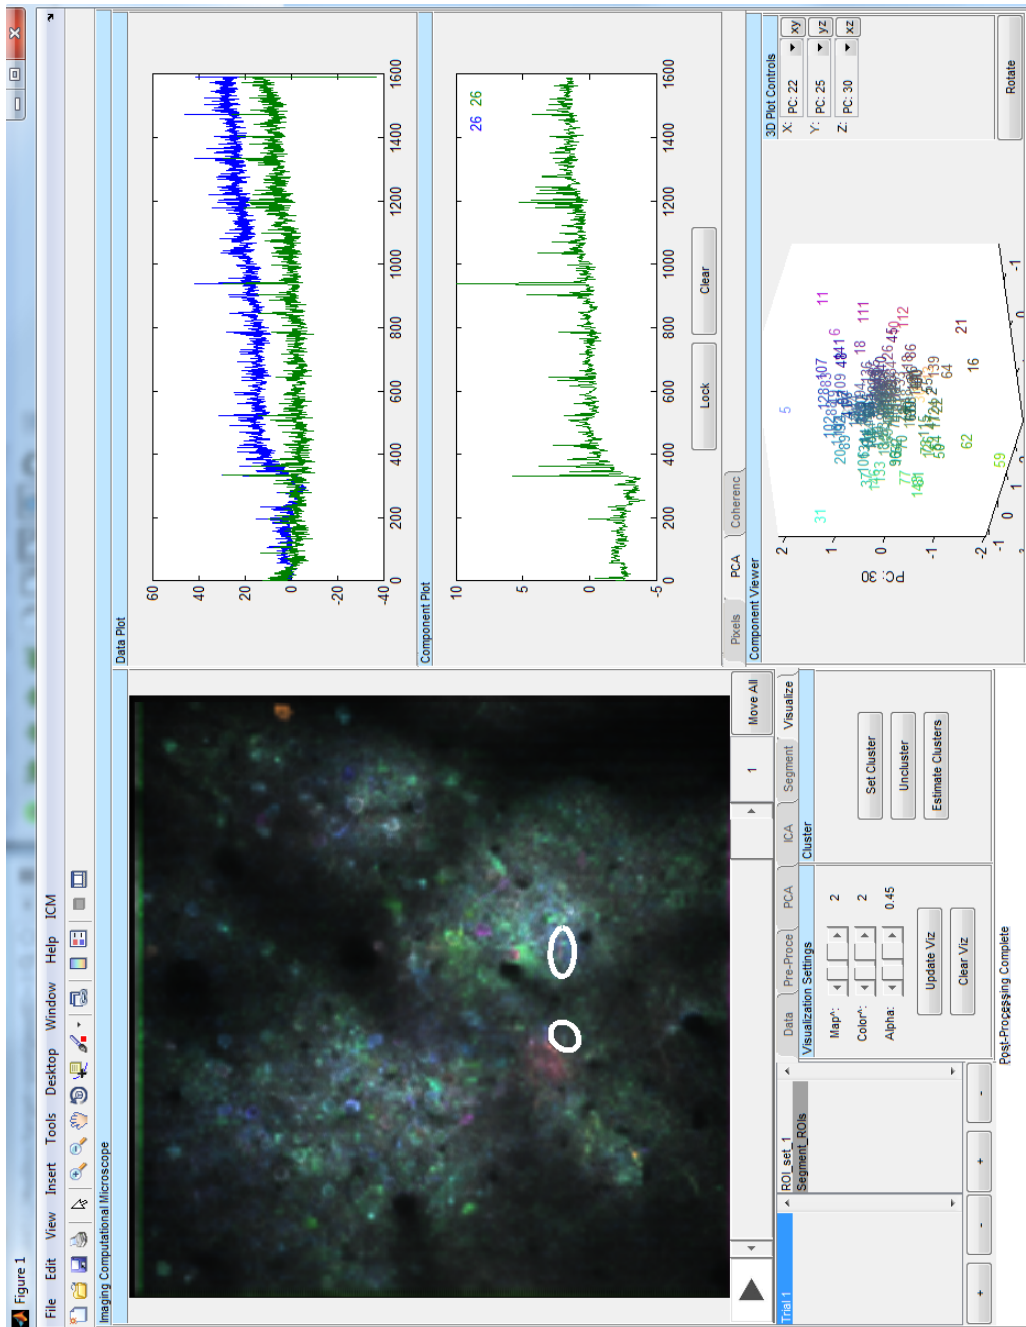


Figure 5.7: ICM screenshot in visualization stage.

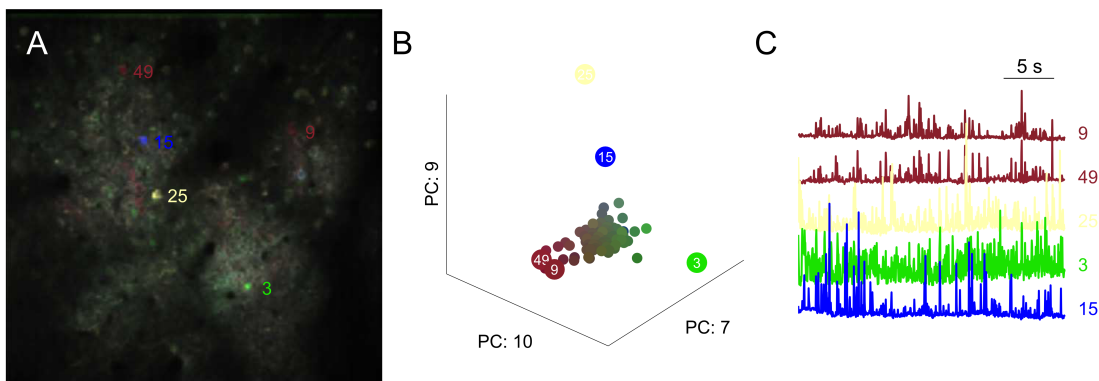


Figure 5.8: ICM visualization reveals patterns in large-scale imaging data. (A) A visualization of cellular activity is created by ICM by coloring IC maps based on the PC coefficients. Five cells are indicated with numbers as examples. Cells 9 and 49 are nearly the same color. (B) Three principal components are chosen in the GUI and the coordinates of each IC are plotted in three dimensions. The colors for each IC are based on these coordinates. Five examples are labeled. (C) The sources of the five example components are shown. Components 9 and 49 have very similar spiking activity, which is reflected by their similar color. The other cells have different spiking activity and are colored differently.

in the “Visualization Settings” panel (Fig. 5.7).

Visualizations are extremely useful for quickly assessing activity patterns in imaging data. In Figure 5.8A, we illustrate the PCA visualization and highlight some example components of interest. To create this visualization, PC 7, 10, and 9 were selected in the PC component viewer (Fig. 5.1 R8) and ICM uses the coefficients from these three coordinates to create the visualization (Fig. 5.8B). This particular example shows a handful of cells with different colors, and notably IC 9 and 49 are very close to the same color. This is the result of the fact that cells 9 and 49 have very correlated spiking activity (Fig. 5.8C), which places them in nearly the same location in PCA space (Fig. 5.8B).

Multi-trial extraction of Voltage-Sensitive Dye Imaging Data

Often many different trials are performed while imaging under different experimental conditions, and it is important to extract the same cells under these different conditions for analysis. The best solution for this situation is concatenated-trial ICA, where the multiple imaging trials are aligned and concatenated into a single imaging data set. This extracts the same cell as the same component from each trial, ensuring that cross-trial comparisons are valid.

In this section, we illustrate the programmatic use of ICM with an example of concatenated trial ICA (ctICA). This walk-through follows the script `walkthrough2.m` in Supplementary

Appendix B. This will illustrate scripted use of ICM, and reveal some of the more advanced features of the software. We will follow along the Matlab script and also point out how actions in the script can be done in the GUI.

To enable ctICA, several data files would first be loaded through the GUI menu or using the function `add_trials` (Supplementary Appendix B Section 2). These trials would be imaging data collected from the same region of the brain and the changes in position of the image field or its z-depth must be minimal.

To perform ctICA, ICM concatenates multiple trials as if they were a single imaging acquisition. For this to work, however, the pixels must be consistently aligned with the sources of the signals. ICM uses an image registration algorithm (Evangelidis & Psarakis, 2008) to align the acquisitions across trials. To enable ctICA, the “Concatenate Trials” button is pressed in the Data tab when ICM is in the “Initialized stage”. This calls the function `align_trials`, which performs the image registration on all of the trials currently loaded and puts ICM into ctICA mode (Supplementary Appendix B Section 3).

Once ICM is in ctICA mode, the process of extracting the independent components (ICs) is essentially the same as the single trial extraction. Pre-processing, PCA, and ICA are run using the controls in each of their tabs. These VSD recordings are much less sensitive than Calcium imaging recordings, and to maximize the signal we set the pre-processing smooth window to $6 \times 6 \times 1$, which is about the size of the smallest cell. This smoothing makes each pixel record the average of the 6×6 pixels around it, increasing the signal-to-noise ratio of the voltage signals, which will improve the extraction of the ICs (Supplementary Appendix B Section 4).

PCA is then run on the concatenated data as if they were a single acquisition. Each PC will have a map and a source, but the source is now a time-series that extends through all three concatenated trials. The map is the same for all three trials. After PCA is run, the maps and sources of each trial can be viewed in ICM. The single concatenated source is split back up to individual trial sources, and these can each be seen by looking at the PCA tab in the different trials using the trial selection list (Fig. 5.1 R2). In Figure 5.9A, four example PCs are shown along with the three sources from the individual trials (Supplementary Appendix B Section 5).

ICA is then run using the PCs defined in the settings struct. The number of PCs to use in ctICA should be even larger than in individual trials, because more artifactual components may be present due to distortions when concatenating the trials. Again, the ICA maps are the same across all three trials, but the sources will be unique for each trial. This means that the locations from which the sources come from are the same (i.e. the component is the same cell across




Figure 5.9: Examples of ctICA components. (A) Four example principal components from ctICA analysis. Three sources (right) are shown corresponding to signals from the three concatenated trials. The component number is indicated at the bottom right of the component maps (left). (B) Four example independent components. The three sources (right) shows activity from an individual cell across three different trials. The component number is indicated at the bottom right of the maps (left).

trials), but the cell signals are different (because the cell's activity is different in different trials). These maps and sources for each trial can be viewed using the trial selection list (Supplementary Appendix B Section 6).

Segmentation is run in the same fashion as the individual trials. Since the segmentation is done on the ICA maps and these are the same across concatenated trials, there will be the same segments for all trials (Supplementary Appendix B Section 7). Segmentation is needed when there are cellular signals that are highly correlated, because the ICA algorithm can lump strongly correlated signals into a single component. In Figure 5.10A an IC map is shown for a component that captures the strongly correlated activity of two electrically coupled bilateral pairs. These cells can easily be segmented because they are spatially separated (Fig. 5.10A, right).

The ICA algorithm can also split a single cell source into multiple components (Fig. 5.10C). This sometimes occurs with very large cells. The user can manually set components into

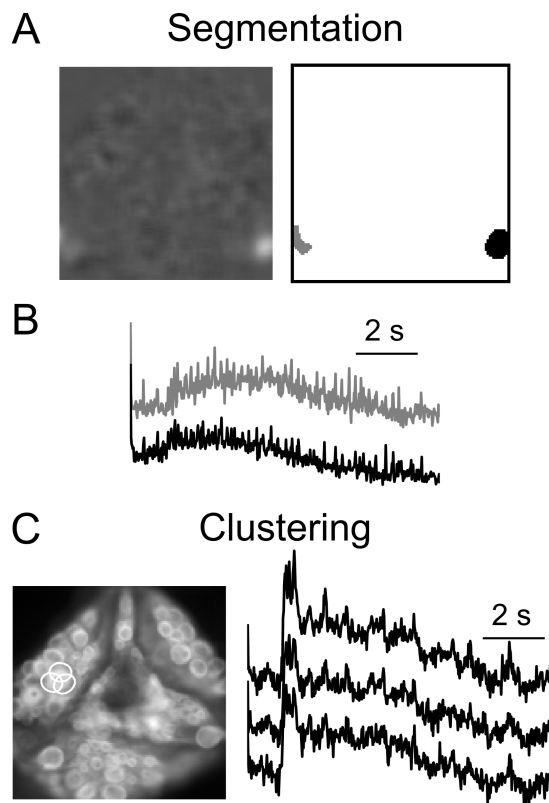


Figure 5.10: Segmentation and Clustering of ICs (A) An example IC where two correlated cells are extracted as the same component (left) can be spatially segmented with a threshold (right). (B) The ROI signals from the gray and black segments derived in panel A show these two cell’s spiking activity is highly correlated. (C) An example of three extracted ICs that correspond to a single cell. These can be manually clustered using the GUI.

a cluster by selecting the ROIs of the ICs and pressing “Set Cluster” in the Cluster control panel in the Visualize tab (see Fig. 5.7).

Finally, visualizations can be made for each trial using the concatenated results. In the first trial, a cell was stimulated with electrodes at 0.5 Hz while VSD activity was simultaneously recorded. ICM’s coherence tools are useful for creating a visualization of rhythmic signals captured by the VSD imaging. To create this visualization, first the IC of the stimulated cell was selected by browsing the ICs using the GUI (Fig. 5.11A top trace), which shows the spiking of the stimulated cell. This was set as the “base” in the coherence tool, which then calculates the coherence between the base and the other VSD signals (Bokil et al. 2010; chronux.org). The frequency of the coherence was then set to the stimulus frequency, and the coherence magnitude and phase at that frequency for every IC is plotted as a polar plot (Fig. 5.11B). The ICs with a

coherence magnitude greater than the significance level (dashed red line in Fig. 5.11B) are given a color based on their phase. This color is used to color the IC maps and create a visualization of the activity (Fig. 5.11C) (Supplementary Appendix B Section 8).

Note that because the fastica algorithm uses a random initialization, it will not always return the exact same components or order of components. In the script, we include some advanced commands to initialize the mixing matrix, which should then lead to the exact same components being computed. If the script is not producing the correct visualizations, then it is possible that the component order of the ICs is incorrect and the settings of the visualizations are also incorrect. It is typical to have to browse the ICs after they are computed and search for the cellular signals of interest by hand.

In the second trial, the leech is performing the swim behavior, which is a cyclical oscillatory behavior and many neurons can be seen oscillating in the VSD recordings. This can again be visualized using ICM's coherence tools. A known cell that shows strong oscillations is chosen manually as the base (Fig. 5.11D top trace), and the dominant frequency of the oscillation is chosen as the coherence frequency (Fig. 5.11E). In this case, the base cell oscillation is defined as the 180° phase, and so the phases of all the other components are rotated. This is used to create a visualization of the oscillations of the neurons (Fig. 5.11F) (Supplementary Appendix B Section 9).

In the third trial, we use an outside function to create the visualization colors and import these colors into ICM. This trial is the shortening behavior and we use factor analysis to find low-dimensional descriptors of the IC sources (Fig. 5.11G). Factor analysis fits pre-defined curves to each IC source and returns the coefficients of these fits (Fig. 5.11H). A color for each IC is then computed based on these coefficients and these are imported into ICM using the function `set_viz_colors`. ICM then creates the visualization with the given colors (Fig. 5.11I; Supplementary Appendix B Section 10).

Performance of ICA vs. ROI

Regions-of-interest (ROIs) are classically used to extract signals from imaging data. An intrinsic issue with this method is that often the signals of interest will overlap and a single pixel can be receiving signal from multiple sources. This type of scenario is impossible to compensate with ROI methods, and some signal must be sacrificed or lost. An important advantage of the ICA algorithm is that it can use the higher-order statistics of the pixels to extract component signals and does not depend on spatial localization. This means that components that overlap on

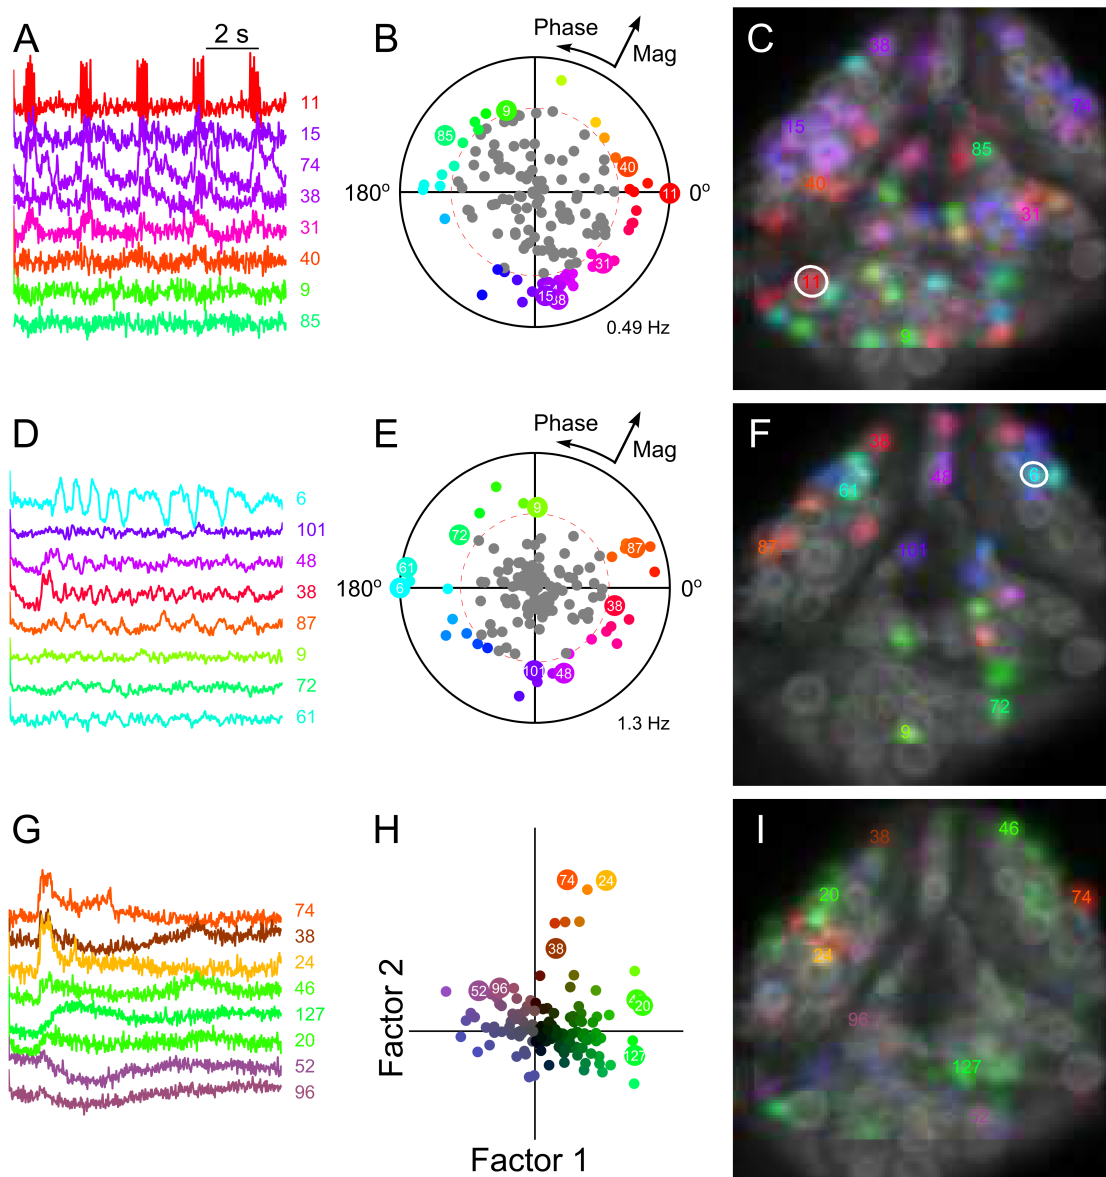


Figure 5.11: Visualizations of multiple trials with ctICA. (A) A sensory P cell was stimulated (top trace) and several VSD traces of example cells are shown which respond to the P cell stimulus. (B) The coherence of the VSD traces was computed against the P cell trace. Cells which had a significant coherence (red dashed line) were assigned a color based on the phase of their oscillation. (C) The IC maps are colored and an activity map is created to visualize neural responses. The white circle indicates the stimulated cell. (D-F) Same as A-C, except the coherence magnitude and phase of an oscillatory motor pattern, swimming, are calculated. (G) Several VSD traces of the shortening behavior. (H) The shortening responses are decomposed into two factors and the coefficients for each IC are plotted. These coefficients determine the color assigned to each IC. (I) An activity map is generated using the colors derived in panel H.

the same pixels can be separated by ICA, and that the ICA algorithm could potentially produce more accurate signals than ROIs.

The ICA algorithm has been previously used to extract optical signals, and reliable cell signals can be extracted (Mukamel et al. 2009; Hill et al. 2010). However, the objective quality of this procedure has never been measured nor compared to other techniques for extracting optical signals. This is because for most imaging data, the imaged signals are based on Calcium, and there is no underlying ground-truth measure of Calcium. However, for VSD imaging, we do have an objective ground-truth measure: the intracellular potential, and can compare optical signals to intracellular recordings by performing simultaneous VSD imaging and electrophysiological recordings.

To assess the relative advantages of ICA based signal extraction, we compared the quality of signals extracted by ICA to those derived by the classical region-of-interest (ROI) technique. We made simultaneous intracellular and optical VSD recordings (Miller et al. 2012) of several cells in the leech ganglion under several different conditions (Fig. 5.12A). We came up with an objective measure of performance by computing the “signal-to-noise ratio” (SNR) of the extracted optical signal with the electrophysiological signal. The most desired result of VSD imaging would be that the optical signals perfectly mimic the intracellular recordings. However, the sampling frequency is different between the intracellular and optical recordings, and optical signals are corrupted by bleaching artifact. To correct for these expected downsides of imaging, we down-sampled the intracellular recording (Fig. 5.12B) to the optical frequency. The down-sampling averages all of the intracellular recording samples that occur during the exposure for each frame of the imaging data, mimicking the way the optical signal is sampled. This results in the Ideal Optical signal (Fig. 5.12C), which would be the best optical signal we could possibly record given our imaging sampling methods.

To remove the bleaching artifact, we use a polynomial to fit the Raw Optical trace to the Ideal Optical signal (Fig. 5.12D). Figure 5.12E shows the Fit Optical signal overlaid onto the Ideal Optical signal. The SNR is computed as the standard deviation of the Ideal Optical signal divided by the standard deviation of the difference between the Ideal and the Fit (Fig. 5.12F).

$$SNR = \frac{\sigma(Ideal)}{\sigma(Ideal - Fit)} \quad (5.1)$$

Motion artifact is another big component of errors in optical signals. ICM includes a built-in motion-correction algorithm (Evangelidis & Psarakis, 2008). This simply registers each frame of the movie to the previous frame using an affine transform (i.e. scale, skew, rotation and

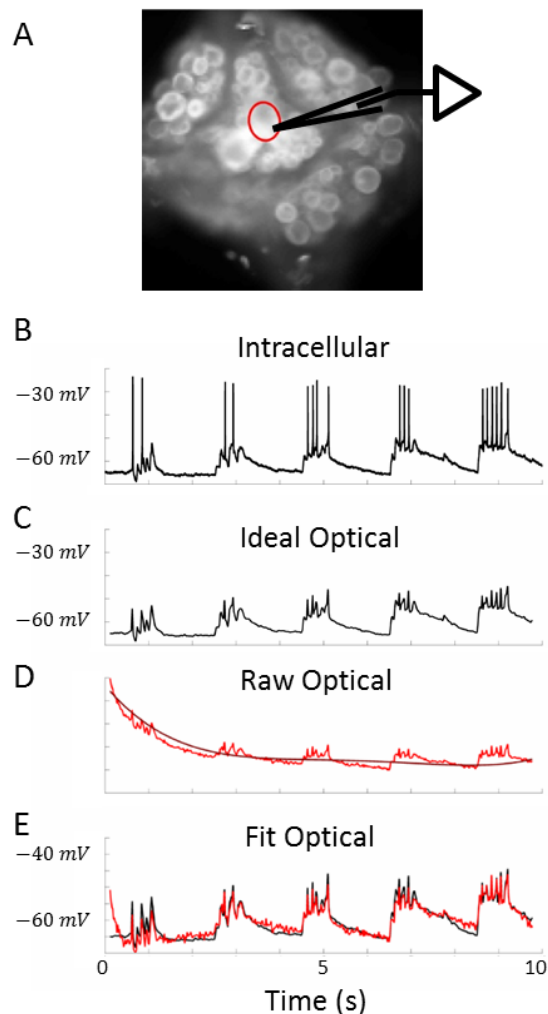


Figure 5.12: Calculation of Signal-to-Noise Ratio (SNR). (A) Simultaneous optical and electrophysiological recordings of several cells under several different conditions (B) Raw intracellular voltage trace. (C) The Voltage trace is down-sampled to the optical frequency based on a method that mimics the optical sampling mechanism. (D) The raw optical trace is shown in red and a polynomial is used to fit the optical trace to the ephys trace (dark red). (E) The fit optical trace is overlaid on the ideal optical trace. (F) The SNR is calculated as the standard deviation of the Ideal trace divided by the standard deviation of the difference between Ideal and Fit (i.e. the residuals).

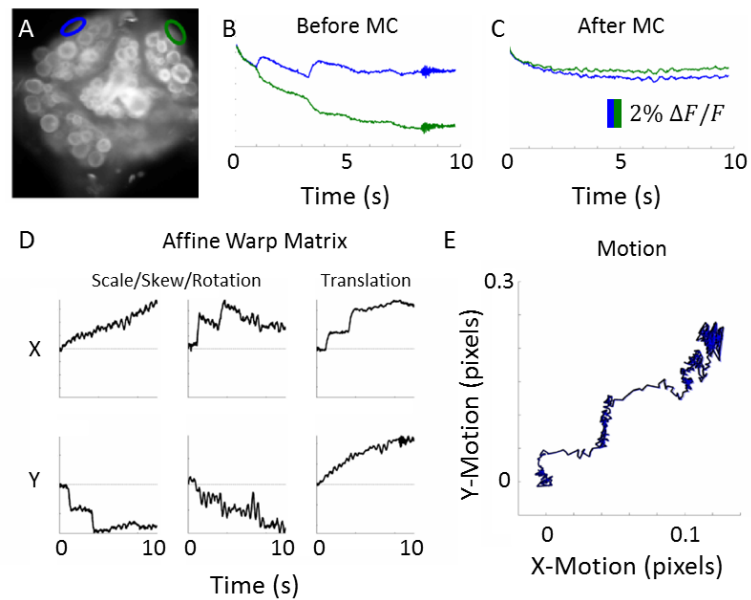


Figure 5.13: Sub-pixel motion correction with ECC. (A) Raw image of ganglion stained with VSD with two ROIs at the edge of the ganglion. (B) Sub-pixel motion artifact is apparent and large compared with VSD signals. (C) Motion artifact can be removed with ECC image registration algorithm. (D) The registration algorithm performs an affine transformation of the pixels. Each panel corresponds to the warp matrix values over time. (E) The resulting motion artifact that is removed. The separation of blue and black highlights the impact of the skew and rotation values of the affine transformation.

translation). This algorithm can even remove sub-pixel motion artifact, which cannot be seen by eye. Figure 5.13 illustrates sub-pixel motion correction removed by the algorithm. The blue and green ROIs shown in Figure 5.13 are at the edge of the ganglion, and they reveal the subtle motion artifact (Fig. 5.13B) as two slight tugs towards the blue ROI. This motion can be fairly well removed using the motion correction algorithm (Fig. 5.13C), which computes the Affine Warp matrix for each frame to remove the motion (Fig 5.13D, E).

We assessed the performance of ICA and ROI extraction techniques by comparing the SNR of several simultaneous recordings. We also compared the two techniques with and without sub-pixel motion correction (no recordings with significant motion were used). Fifteen recordings were made during several different conditions: spontaneous activity, chemical pre-synaptic stimulation, electrical pre-synaptic stimulation, and the swimming behavior.

Figure 5.14 shows some examples of the different extraction techniques under two different conditions (chemical pre-synaptic stimulation, and swimming Fig. 5.14C, F, respectively). The ROI used for extraction was an oval drawn by hand over the cell being recorded from in-

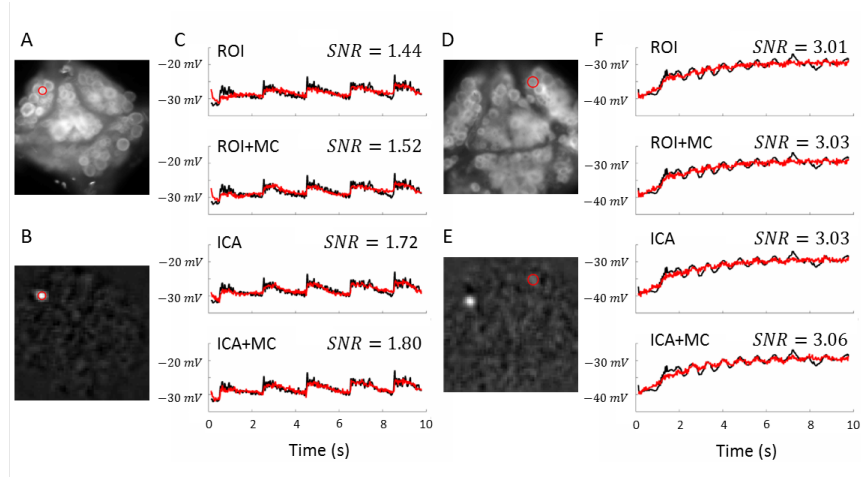


Figure 5.14: Comparison of ROI and ICA based signal extraction. (A) Raw image with red ROI shown during pre-synaptic stimulation. The optical signal is the average of all pixels within the ROI for each frame. (B) ICA pulls out a component that is manually selected coming from the same cell. The ROI is shown on top to compare localization of component and hand-drawn ROI. (C) The Fit trace is overlaid on the Ideal trace for 4 conditions: ROI only, ROI with Motion Correction (ROI+MC), ICA only, and ICA with Motion Correction (ICA+MC). (D) Raw shown during swim behavior. (E) ICA component. This component shows weight in the bi-lateral pair of cells, as these cells are highly correlated. (F) SNR is compared under the 4 conditions. The 5 mV oscillation are much more clear in the ICA+MC case, even though the SNR increase is fairly small.

tracellularly. The ICA component corresponding to the cell was selected using ICM. In some cases, cellular signals are correlated enough to come out as a single component, and occasionally a single IC corresponds to multiple cells (i.e. Fig. 5.14E).

For each of the 15 recordings, we compared the SNR across the 4 different methods. Because the SNR of an individual cell varies from cell to cell, we measured the difference in SNR between the different conditions. Compared to the ROI method, ICA improves SNR for almost every recording. Motion correction also improves SNR, and ICA+MC leads to the best SNR results (Fig. 5.15).

Discussion

Here we have demonstrated the Imaging Computational Microscope as a tool for automated analysis of large-scale imaging data. This tool can be used to rapidly extract component signals from imaging data and create visualization for real-time experimental feedback. The PCA-ICA extraction has been shown to be a promising technique for analyzing imaging data,

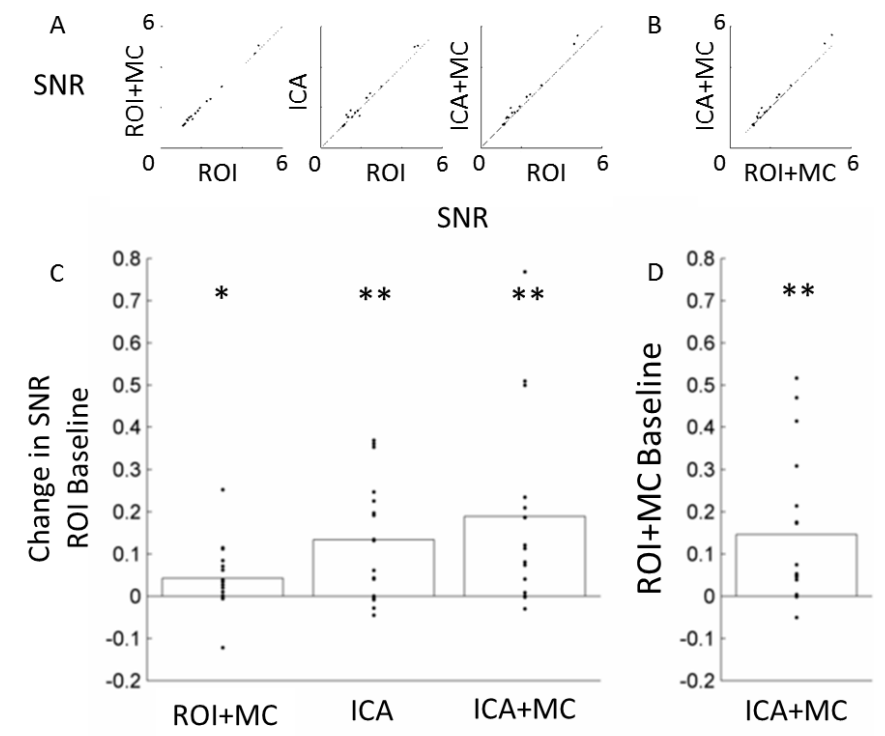


Figure 5.15: SNR comparison of different methods. (A) The SNR for each method is compared against the ROI method. (B) ICA+MC is compared with ROI+MC as the baseline. (C) The difference in SNR for each trace is plotted as black dots, and the average is plotted as the bar graph. Each method significantly increases SNR (ROI+MC: $p = 0.0358$, ICA: $p = 0.0013$, ICA+MC: $p=0.0024$, t-test). (D) Same as C but with ICA+MC compared to ROI+MC ($p=0.004$).

and we have created ICM to aid in the rapid use of these algorithms. Further, ICM incorporates user feedback and visualization to overcome some of the short-comings of the PCA-ICA extraction. ICM also includes the ability to perform concatenated trial ICA, which is an important technique to analyze multiple trials of imaging data with consistent components.

We showed that cellular signals extracted with PCA-ICA have improved SNR compared with ROI based methods by measuring the difference between optically recorded and intracellular voltage signals. This improvement is likely because many spatially overlapping components are recorded by each individual pixel, such as signals from neighboring cells or background fluorescence, and ROIs cannot separate these components. ICA, however, can separate spatially overlapping components, which allows artifacts or overlapping cellular signals to be removed from the true source signal.

ICA stability

There are several important theoretical assumptions underlying the use of ICA to extract components from imaging data. The first is that there must be enough data points for the independence of the underlying sources to be accurately computed. This means that there should be many more frames (samples) than there are cells (sources). The second is that ICA assumes that the sources are stationary throughout the data acquisition. To compensate for slight distortions in the spatial location of the image plane, ICM has image registration and motion-correction algorithms that use the visual features of the imaging data to align the sources with the pixels. ICA fundamentally also requires that the signals of interest do not have Gaussian statistics, because ICA cannot assess independence from Gaussian probability distributions.

The stability of ICA depends on the dimensionality of the data, and in general the higher the dimensionality, the more data needed to extract independent components. Imaging data is extremely high-dimensional, because the dimensionality is proportional to the number of pixels, which is easily larger than 10K for most imaging applications. For some applications, such as EEG (Makeig et al. 1996; Delorme & Makeig, 2004), ICA can be directly applied to the data because the dimensionality of the data is in the dozens. The primary purpose of PCA is to reduce the dimensionality so that the ICA algorithm can work with datasets with very large dimensionality. This technique can be extended to other modalities of data collection with large dimensionality, such as multi-electrode arrays (Jäckel et al. 2012) and 3-dimensional acquisition from 2-photon, light-sheet (Ahrens et al. 2013; Freeman et al. 2014), or fMRI.

Difference between ICA algorithms

ICM includes multiple ICA algorithms that can be chosen from. However, these algorithms typically produce similar results, and it is unclear which, if any, algorithm has an advantage. The computation of ICA is an optimization problem, and generally the different algorithms use slightly different methods of estimating the optimization function. The default algorithm, *fastica* (Hyvarinen & Oja, 1997), is the fastest ICA extraction and produces consistent results. These results are on the face indistinguishable from the results of other ICA algorithms, but there may be certain advantages in some circumstances. There are many settings that can be tuned to alter the extraction, but usually the most important setting is the choice of PCs to use for ICA. We recommend that the PCs used for ICA be experimented with and the results can be judged by manually viewing the ICs.

Typically, the first N PCs are used for ICA, but occasionally the imaging signal can have

noise or artifacts present in the first few principal components. It can sometimes be beneficial to remove the early PCs before performing ICA, especially if they visually appear as pure artifact. However, this demands caution, because the PCs are usually mixtures of components and removing an artifactual PC may also remove some underlying signal.

Multi-Trial ICA

Performing the ICA analysis across multiple trials has some issues because of the random nature of the ICA decomposition. Running the analysis, even on the same data, does not mean that the same components will be extracted in the same order. This is because the initialization of the mixing matrix is random, and the order of the components is somewhat dependent on the initialization of the mixing matrix. We experimented with an alternative solution for multi-trial ICA where we initialized the mixing matrix in one trial using the ICA results from a different trial. This can be performed with ICM using the function `init_spatial_guess`. This essentially amounts to providing a guess of spatial locations for expected cellular signals, and the ICA algorithm does find components in similar spatial locations. However, because many cells were tightly packed together, this algorithm often returned neighboring cells as the same component and generally was less consistent than ctICA. Further, this becomes problematic when cells are quiet or show little signal in one or more of the trials. If a cell is not active, then ICA has a much harder time of extracting the signal, because there is less information in the statistics of the pixels. This means that ICA can completely miss a cell from one trial and this will distort the component relationships that are returned in multi-trial analysis.

It is essential for ctICA to have the sources and the pixels aligned across trials, and to help with this we use an image registration algorithm to align the trials. This registration algorithm can only perform an affine transformation on the entire image, and cannot correct for motion of sub-regions of the image or differences in focal depth or z-motion. These types of artifacts will hinder the results of ctICA, but it can still give some good results as long as most pixels over a source remain over the source. More sophisticated image-alignment techniques, especially those with locally deformable registration, can further improve the results of ctICA. Additionally, some of the artifacts from mis-alignment can be compensated through blurring the data in the pre-processing stage. Further, the alignment artifacts are sometimes extracted as a component from ctICA, and these components can simply be ignored.

Memory requirements

The algorithms and data involved in ICM's analysis require very large amounts of memory. Even just a single imaging data set can be gigabytes in size. It is recommended that the software is run with at least 16 GB of memory on the computer, but even this much can be quickly used up by these algorithms.

Acknowledgements

This chapter is presented as a manuscript that is in preparation as a publication as Frady, E.P., Kristan Jr., W.B. "The Imaging Computational Microscope."

References

- Ahrens, M.B., Orger, M.B., Robson, D.N., Li, J.M., Keller, P.J. (2013) Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods* 10: 413-420.
- Bell, A.J., Sejnowski, T.J. (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.* 7: 1129-59.
- Bokil, H., Andrews, P., Kulkarni, J.E., Mehta, S., Mitra, P.P. (2010). Chronux: A platform for analyzing neural signals. *Journal of Neuroscience Methods* 192: 146-151.
- Delorme, A., Makeig, S. (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods* 134: 9-21.
- Evangelidis, G.D., Psarakis, E.Z. (2008) Parametric Image Alignment Using Enhanced Correlation Coefficient Maximization. *IEEE Trans. on PAMI [Pattern Analysis and Machine Intelligence]* 30(10):1858-1865.
- Freeman, J., Vladimirov, N., Kawashima, T., Mu, Y., Sofroniew, N.J., Bennet, D.V., Rosen, J., Yang, C.T., Looger, L.L., Ahrens, M.B. (2014) Mapping brain activity at scale with cluster computing. *Nature Methods* 11(9): 941-949.
- Hill, E.S., Moore-Kochlacs, C., Vasireddi, S.K., Sejnowski, T.J., Frost, W.N. (2010). Validation of Independent Component Analysis for Rapid Spike Sorting of Optical Recording Data. *J Neurophysiol* 104: 3721-3731.
- Hyvarinen, A., Oja, E. (1997) A fast fixed-point algorithm for independent component analysis. *Neural Comput.* 9, 1483-1492.
- Jäckel, D., Frey, U., Fiscella, M., Franke, F., Hierlemann, A. (2012). Applicability of independent component analysis on high-density microelectrode array recordings. *J Neurophysiol* 108:

334-338.

Lee, T.W., Girolami, M., Sejnowski, T.J. (1999) Independent component analysis using an extended infomax algorithm for mixed subgaussian and supergaussian sources. *Neural Comput.* 11: 417-41.

Lee, T.W., Girolami, M., Bell, A.J., Sejnowski, T.J. (2000) A Unifying Information-theoretic Framework for Independent Component Analysis. *Comput. Math. Appl.* 31: 1-21

Makeig, S., Bell, A.J., Jung, T.P., Sejnowski, T.J. (1996). Independent component analysis of electroencephalographic data. In Touretzky, D., Mozer, M., Hasselmo, M. editors. *Adv. Neural Inf. Process. Syst.* 8: 145-51.

Miller, E.W., Lin, J.Y., Frady, E.P., Steinbach, P.A., Kristan Jr., W.B., Tsien, R.Y. (2012) Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires. *PNAS* 109(6): 2114-2119.

Mitra, P.P., Bokil, H. (2007) *Observed Brain Dynamics*, Oxford University Press, USA.

Mukamel, E. A., Nimmerjahn, A., Schnitzer, M. J. (2009). Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data. *Neuron.* 63(6):747-760.

6 Scalable semi-supervised cell identification reveals canonical swim and preparatory networks.

Abstract

Using a voltage-sensitive dye (VSD), we imaged the multi-functional responses of hundreds of leech neurons during multiple behaviors to identify and functionally map homologous cells using a semi-supervised machine-learning framework. From these data we defined canonical networks, cell-by-cell in a scalable fashion, so that dozens of cells were identified and studied in real time. We identified 21 neurons as part of a canonical swim network as well as nine cells involved in a partially overlapping preparatory network. We showed that the preparatory network rapidly depolarized before the onsets of swimming and shortening, and that the same population of cells was activated during both behaviors. This network further shows an earlier response to stimuli than cells not in the preparatory network, indicating that it is part of a dedicated rapid-response network. This network is likely to be mediated by the S cell, and we use scalable identification and VSD imaging to identify cells and probe for potential synaptic partners to the S cell. We verify that the S cell is indeed pre-synaptic to several cells in the preparatory network using real-time cell identification and computationally guided electrophysiology.

Introduction

The identification of individual neurons and neuron types is a challenging problem that is especially hard when dealing with large-scale neural recordings. Cells and cell types are usually classified using prominent features, or cells can be identified with genetically targeted markers (Bernard et al. 2009), but these techniques can identify only a limited number of cells or cell types. Our visual systems are excellent at using multiple features to identify objects, and because our brains are so good at recognizing such patterns, the task of identifying a particular

cell is often taken for granted. However, this identification task is not a trivial problem for a computer. Further, when we record from large numbers of neurons, identifying an individual neuron becomes the task of identifying many neurons simultaneously. This is a much more challenging problem, and cannot be solved by solely relying on human pattern recognition. Simply identifying neurons and neuron types in large-scale data presents a complicated machine learning challenge, one that is essential to overcome for large-scale recording efforts such as those being proposed by the BRAIN Initiative (Alivisatos et al. 2012).

The neurons in the leech nervous system are individually identifiable by their homologous anatomy, electrophysiological properties, and relationships to behaviors from ganglion-to-ganglion and from animal-to-animal (Muller et al. 1981). However, the majority of the approximately 400 neurons within a single ganglion (Macagno, 1980) remain unidentified, primarily because most neurons require examination of their synaptic connections to other neurons or their relation to one or more behavioral responses to be identified. The somata of homologous neurons are located in roughly similar ganglionic positions, but this can vary enough that location alone can only narrow the identification to one in 5-10 neurons (Weeks, 1982; Friesen, 1989). This need to use physiological parameters to identify a single cell of interest triggers the question: how can we scale up the identification of cells to large data sets?

In the leech nervous system, advances in voltage-sensitive dyes (VSD; Miller et al. 2012) have enabled an unprecedented resolution of imaging neural activity during behavioral states. We can record voltage from over 80% of the neurons on the ventral surface with exceptional signal-to-noise ratio and temporal resolution. We used VSD imaging to simultaneously record from many leech neurons during multiple fictive behaviors and developed a semi-supervised machine-learning framework for identifying neurons across animals based on their VSD activity patterns. We used combinations of functional and anatomical features to abstract the imaging data set into a “rich medium-dimensional feature space”, and then used machine-learning techniques to piece together a canonical map of neural activity. Current machine-learning techniques still fall short of the capabilities of human pattern recognition. To bridge the gap of computational sophistication and to create a scalable architecture, we turned to semi-supervised techniques – merging algorithms, machine learning, visualization, and human feedback at each stage of the framework.

We used these techniques to identify canonical swim and preparatory networks by stitching together data collected across eight different animals. We characterized the responses of dozens of previously unidentified cells during each of these behaviors, and we showed that the

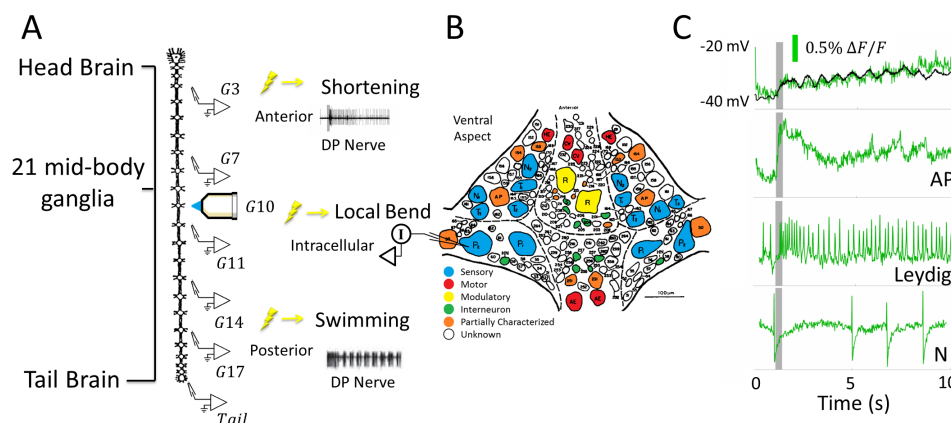


Figure 6.1: Voltage-sensitive dye imaging during multiple behaviors (A) The isolated nerve cord of the leech was dissected and ganglion 10 was prepared for VSD imaging. Anterior or posterior nerves were be stimulated with an extracellular suction-electrode to elicit shortening or swimming, respectively. (B) Schematic of leech ganglion. Cells that have been previously identified are shown in color. Local bending is activated by targeted stimulation of a P cell. (C) Voltage-Sensitive Dye imaging of cells. A simultaneous intracellular (black) and optical (green) recording is shown (top), which illustrates clear signals of oscillations less than 5mV. Three other example traces show optically recorded action-potentials.

preparatory network is rapidly activated by stimuli that produce several different behavioral outputs. One previously identified neuron, the S cell (Laverack, 1968; Mistick, 1974; Bagnoli et al. 1975; Mistick, 1978), is a component of the preparatory network, and earlier work suggests that the S cell is part of a “giant fiber” system in the leech. We tested the possibility that the S cell plays a role in mediating the preparatory response by testing whether the S cell was pre-synaptically connected to other preparatory neurons. We first identified dozens of cells in the swim and preparatory networks, enabling us to perform S cell experiments with many cells identified within the context of these networks. Using simultaneous VSD imaging and intracellular recordings from S cells, we showed that the S cell is functionally coupled to several neurons in the preparatory network. We verified several of these synaptic connections using targeted electrophysiology aided by rapid identification of cells.

Methods and Results

The isolated leech nervous system can elicit different fictive behaviors depending on where it is stimulated (Briggman et al. 2005; Briggman & Kristan, 2006). We used VSD imaging (Miller et al. 2012) to optically record from hundreds of cells in a single midbody ganglion

(ganglion 10) during 3 behaviors: shortening, swimming and local bending (Fig. 6.1A). In an intact leech, swimming and shortening are elicited by stimulating the animal near the posterior and anterior, respectively (Palmer et al., 2014). Fictive shortening was activated by stimulating a nerve (DP) from either ganglion 3 or 7, and swimming was activated by stimulating 11, 14, 17, or the tail brain. Local bending was activated by intracellular stimulation of a sensory P cell (Fig. 6.1B; Kristan, 1982; Lockery & Kristan, 1990a,b; Lewis & Kristan, 1998a,b,c). Of the nearly 400 neurons in a leech ganglion, only about a third have been previously identified (Fig. 6.1B, colored cells; Muller et al. 1981). About 80% of cells on the ventral surface can be imaged with VSDs, but due to the variability of the ganglion across animals as well as variability in dissections and VSD application, a random subset of these cells are imaged in each experiment. We recorded from 8 different animals (A-H) to sample the cells in the ganglion multiple times. We imaged a ganglion during the onset of each behavior one to four times at 50 Hz for 10 seconds, which allowed us to monitor synaptic potentials, action potentials, and oscillations of the membrane potential in individual cells in single trials (Fig. 6.1C). We refer to an individual recording as a “trial”, and the set of recordings from a single animal as an “experiment” or “animal”.

Typically, imaging data is decomposed into cellular responses using manually defined regions of interest (ROIs). New computational techniques using principal and independent component analysis (PCA, ICA; Mukamel et al. 2009; Hill et al. 2010) have been shown to be capable of automatically extracting cellular signals directly from imaging data. The ICA-based algorithms use the higher order statistics (i.e. skew and kurtosis) of the pixels in the imaging data to extract components, but the algorithm does not distinguish cellular signals from artifacts. In addition, these algorithms require that the number of components must be specified a priori. If too few components are specified, multiple neurons may be lumped together as a single component; if too many components are specified, signals from an individual neuron might be classified as two or more components. In simultaneous recordings from many neurons, both problems can occur in the same data set, and must be resolved manually. To handle these shortcomings, we built a user-interface called the “Imaging Computational Microscope” (ICM; Chapter 5), which automatically extracts signals from imaging data. This tool allows the user to quickly visualize the results of the PCA-ICA extraction (Fig. 6.2A, B), alter the parameters of the extraction, segment correlated signals (Fig. 6.2C), automatically generate ROIs (Fig. 6.2D), cluster decomposed individual sources (Fig. 6.2E), remove artifactual components (e.g., Fig.6.2B, bottom two panels), and create “activity maps” for visualization (Fig. 6.3).

We analyzed multiple trials of different behaviors by performing concatenated trial ICA (ctICA). In this procedure, we used an image registration algorithm (Evangelidis & Psarakis, 2008) to align multiple trials as if they were obtained from a single imaging acquisition (i.e. a single time-series). It is important to align the trials such that the same pixels are recording from the same neurons, because the extraction algorithm is based on the statistics of the pixels and implicitly assumes that nothing moves relative to the pixels. Changes in z-depth, or deformations of sub-regions of the ganglion were not accounted for in the image registration, and any misalignment of pixels with neurons can cause errors in the algorithm; such z-depth changes and sub-region deformations were very subtle in our data and did not interfere with the component extraction (these small problems were further mitigated by slightly blurring the images). We performed PCA on the pixels from the concatenated imaging data, and then performed ICA on the top 100-150 principal components (PCs). The number of PCs was matched to the number of cells expected in the imaging data set. This procedure is built into ICM and takes only minutes to run on a high-end desktop computer.

Both PCA and ICA are linear component decompositions and when applied to imaging data produce several components, each of which has a “source” and a “map” (Fig. 6.2A and B). Sources are the time series of the extracted components, and the maps are the spatial locations. The ICA algorithm will extract both artifacts and cell signals as components; the artifacts can be quickly sorted by looking through the sources and maps with the GUI. Components that are chosen as signals are called “neurons”. Because in ctICA the trials are concatenated into a single acquisition, the map of each component is the same for all the concatenated trials, and the source is a single time-series that extends through all trials. The sources are broken up for each trial to give individual “traces”, each of which is the activity of a component during a single trial (in Fig. 6.2A and B, two example traces are shown for each component from different trials).

Each leech ganglion, to a first approximation, consist of the same homologous set of cells across animals (Muller et al., 1981; Kristan et al. 2005), but slight anatomical variability makes identifying most neurons non-trivial. To characterize the consistent homologies across animals, we searched for neurons with similar anatomical and functional properties across multiple animals. By using a range of functional and anatomical features, we made a tag for each neuron. To compensate for that fact that not every neuron is imaged in a single ganglion, we used eight different animals. We designed computational algorithms based on coherence and factor analysis (Supplementary Appendix C) to extract low-dimensional functional features characterizing the main aspects of the shortening, local bending and swimming behaviors. We used

these features as both low-dimensional descriptors of neural activity and anatomy, as well as the characteristics used to identify homologous cells across animals.

The shortening response was elicited by stimulating anterior DP nerves (ganglion 3 or 7) and motor neuron spikes were monitored by recording extracellularly from other DP nerves (Fig. 6.3A, top trace). During shortening (Fig. 6.3A, colored traces), neurons showed a variety of responses, but they were well characterized by two factors (Fig. 6.3B) that were derived from neural responses during all shortening trials across animals (Supplementary Appendix C). Factor 1 describes prolonged depolarizing (positive Factor 1) or hyperpolarizing (negative Factor 1) responses that last as long as the motor neuronal burst that constitutes the fictive shortening. Factor 2 depicts a rapid onset response, and is generally positive indicating that many neurons rapidly depolarize after each stimulus. Factor coefficients (Fig. 6.3C) were computed by fitting each trace with the factors using a linear regression (Supplementary Appendix C). The components are colored based on their two factor coefficients – positive Factor 1 increases the green channel, negative Factor 1 increases the blue channel, and Factor 2 is shown in the red channel. These two coefficients describe 70-80% (check) of the variability of the traces, and are a low-dimensional way of representing the shortening response. An “activity map” was generated by coloring the maps (i.e. the neuronal somata) of each component based on the colors derived from the traces, and then overlaying every component’s map on an image of the ganglion (Fig. 6.3D). This visualization shows the activity of 103 cells and is useful for rapid high-level analysis and visual inspection of large-scale neural activity. The maps are labeled by experiment (A-H) and stimulus location (ganglion 1-21) on the top right and bottom right, respectively, in Fig. 6.3 and all subsequent figures.

To elicit local bending, we activated one of four mechanosensory P cells with 300 ms long bursts of 5-6 action potentials repeated every 2 seconds. This stimulus results in the repetitive initiation of the local-bend response, which was monitored in the DP nerve recordings (Fig. 6.3E, top trace). The traces from the same example neurons as shown for shortening (Fig. 6.3A) indicate that neurons with similar responses in shortening do not necessarily produce similar responses during local bending, illustrating the multi-functional nature of many of these neurons. To characterize this behavior compactly, we calculated the coherence of each VSD trace with the intracellular stimulus burst (Fig. 6.3F). Coherence analysis (Mitra & Bokil, 2007; Bokil et al. 2010) calculates a correlation magnitude and phase in Fourier space for each component, which we plotted in polar coordinates at the stimulus frequency (Fig. 6.3F). In these plots, neurons that are excited by the burst of P-cell spikes are indicated by coordinates in phase (0°) with the

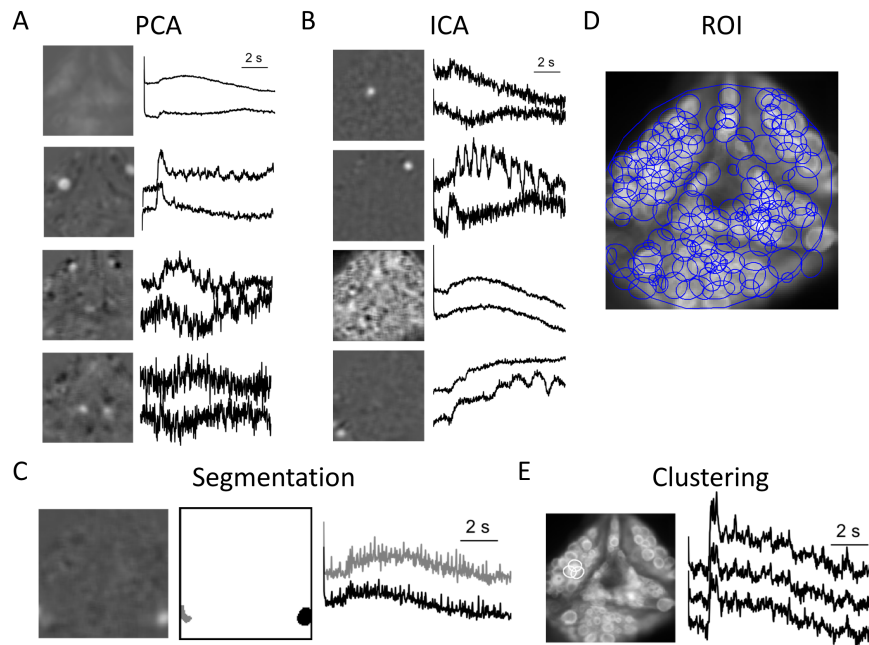








Figure 6.2: Semi-supervised signal extraction with concatenated-trial ICA (ctICA). (A) PCA is performed on concatenated trials. Four example principal components are shown. Each principal component has a PC mask (left) and multiple traces (right). Two example traces of each PC are shown from different trials. (B) ICA is performed on the top 120-150 principal components. The IC masks (left) and traces (right) are shown for four examples. The top two components are cellular signals. The bottom two components are examples of background neuropil response and motion artifact. These components are discarded. (C) The independent component masks (left) are then segmented using a threshold (middle). An example is shown where two cells are part of the same independent component, but can be separated through spatial segmentation. These cells are strongly electrically coupled to each other and ICA does not separate them because they have similar activity (right). (D) Regions-of-interest are automatically generated from the segmentation of the components. (E) Three components that are from a single source cell are shown. The algorithm occasionally splits single sources (especially larger cells), and components with similar locations and correlated traces can be clustered using the GUI.

Figure 6.3: Low-dimensional feature extraction from multiple behaviors. (A) Shortening is activated by stimulating an anterior DP nerve (gray bar). Other DP nerves show motor neuron spikes and are used to read behavioral output (top). Several example VSD traces from identified cells are shown in color. (B) Cell traces are decomposed into two factors, a slow factor that follows the behavior (Factor 1, green) and a fast factor (Factor 2, red). (C) The factors are fit to each trace and the coefficients of the factor fits are plotted. The components are given colors based on these coefficients. The coefficients of the examples from A are shown as large circles with white labels. (D) Activity map of shortening. Each component's map is colored based on the shortening factor coefficients and overlaid onto an image of the ganglion. The example cells are circled. The top right letter, G, indicates that this data is from Animal G, and the bottom left number, 3, indicates which ganglion was stimulated. (E) A burst of 6 action-potentials is given every 2-seconds to sensory P cells (gray bars) to repetitively elicit the local bend response. The motor neuron spikes of the behavior can be seen in the DP nerve (top). The VSD traces of the same example cells are shown in color. (F) The coherence of the sources was calculated against the stimulus at 0.5 Hz (the stimulus frequency), and the coherence phase and magnitude for each trace are plotted as a polar plot. The components are colored based on their phase and magnitude. (G) Activity map of local bending. Colors are based on the phase and magnitude of the coherence analysis. The same example cells are circled. White circle indicates the stimulated P cell. (H) Swimming is activated by stimulating a posterior DP nerve (gray bar). The swim motor pattern is monitored through DP nerve recordings (top). The VSD traces of the same example components are plotted. (I) The coherence of the traces were calculated against the DE-3 motor-neuron spike output at the swim frequency (large spikes in DP nerve recordings). The phase and magnitude are plotted as a polar plot and determine component colors. (J) Activity map of swimming. Same example cells are circled.




Shortening



Local Bending



Swimming



stimulus (pink, Fig. 6.3F) and inhibition is indicated by coordinates that are 180° out of phase (cyan, Fig. 6.3F). We then colored the components based on the phase and magnitude of the coherence analysis and visualized them as an activity map (Fig. 6.3G). The P cell stimulated is shown by white circle in the activity map, and the experiment is designated by the letter (A-H) in the top right.


The swimming motor program was initiated by stimulating posterior DP nerves (ganglion 11, 14, 17, Tail Brain) and monitored as recurring bursts of extracellular spikes at 1-2 Hz by a motor neuron (DE-3) in a DP nerve (Fig. 6.3H, top). Optical traces from the same sample neurons show that the membrane potentials of many cells oscillated at different phases relative to the spikes in the DP nerve during each cycle of swimming. Coherence analysis correlating the VSD traces with the DE-3 spike bursts (Fig. 6.3I) characterized the neuronal membrane potentials that oscillate with the swim motor output and at what relative phase. Based on these coordinates, the components are given different colors representing their phases, which are used to color the somata in an activity map (Fig. 6.3J).

For functional features to be useful in identifying homologous cells across animals, they must have consistent activity patterns across multiple bouts of a behavior (i.e. “trials”) and bouts across animals (i.e. “experiments”). To visualize this variability and to assess the consistency of the features, we analyzed each neuron during two trials of the same behavior. We computed the shortening factors for each component during two trials of shortening across four different experiments (Fig. 6.4A; Each experiment is indicated by a letter at the top-right). For each of the four experiments, we show two activity maps generated from the component responses from two different trials. The shortening factor coefficients from these two trials are plotted together on a single axis, and the coefficients from each individual component are connected as a line segment. A short line indicates small variability during different trials, making these features potentially useful as distinguishing characteristics to identify the cell. To the extent that such a line does not overlap with other lines, this feature could provide robust identification of a neuron’s identity. There were a variety of neurons that consistently depolarized (positive Factor 1), hyperpolarized (negative Factor 1) and/or showed rapid onsets (positive Factor 2). Although the shortening feature and the location of the neuronal somata suggested the identity of a few neurons, most neurons required additional features to be uniquely identified.


Multiple trials of local bending showed consistent and robust responses from many neurons, as indicated by short, isolated lines in plots of their calculated coherence (Fig. 6.4B). Stimulating either a ventral or dorsal P cell (indicated by white circles in each activity map) re-

Figure 6.4: Analysis of within animal and across animal functional variability. (A) The shortening feature is extracted from multiple trials of shortening within the same animal. The same cell's response during the two trials of shortening is plotted as a line segment, where the ends of the line segment are the shortening factor coefficients from each of the trials. The activity maps of the two trials are shown next to the coefficient plot. Four different animals are shown for comparative purposes. Well isolated and short line segments indicates that the feature is useful for identifying that cell. Further the common patterns of the behavior can be seen across animals, indicating that the shortening response (especially factor 2) is a good indicator of cell identity. (B) The same is done for the local-bend feature. A clear set of cells can be isolated that are colored pink and green/cyan. In these examples different P-cells were stimulated, showing that many cells respond consistently to stimulation of either P cell. (C) The same is done for the swim oscillation feature. Many of the cells show consistent phases of oscillations during swimming, making them easily identifiable.


A Shortening



B Local Bending



C Swimming



liably activated a consistent population of neurons. The local bend feature (defined in Fig. 6.3E, F) was most sensitive to cells that closely follow the sensory stimulus (Fig. 6.4B). This feature picks out two clusters of cells – pink cells which are functionally excited by the P cell stimulus, and cyan/green cells which are inhibited. This analysis does not determine whether the activity patterns resulted from direct or indirect synaptic connections from the P cells.

The swimming features (defined in Fig. 6.3H,I) defined neurons that have oscillations phase-locked to the swim motor output (Fig. 6.4C). This feature was so consistent and robust that many neurons can be identified by a single swim activity map alone, because the phase of many neurons' oscillations in conjunction with its somatic position provide a unique identifying tag. This identifiability is reflected in the phase plot (Fig. 6.4I) by the many short, non-overlapping lines that are spread out through different phases of the swim cycle.

Although some of the neurons can be identified by a single behavior, most require combinations of behaviors to provide enough unique information. To aid in identifying more neurons, we combined the multiple functional features with anatomical features (location and size of somata) to craft a canonical “rich, medium-dimensional” feature space that uses all the functional and anatomical features to make comparisons across animals. The set of features described creates a 10-dimensional feature space: two dimensions for the x,y position of the soma, two dimensions for soma size, and two dimensions for each of the functional features describing the 3 behaviors. The space is rich because the features were extracted from different behaviors and this information is close to independent. The volume of n-dimensional space increases exponentially with the dimensionality, so a 10-dimensional feature space can have a huge volume, but this volume is much smaller than that describing the individual neurons (which would be in the hundreds). We term this space “medium-dimensional” because it is more than the readily displayed 3-dimensional space, but much less than the 100-plus dimensional space of the neural responses. Although the components appear crowded in the graphs in Figure 6.4, when they are plotted in 10-dimensional space, many neurons can be readily separated and identified individually.

To explore the feature space and identify cells across animals, we developed semi-supervised machine-learning algorithms that estimate the invariances and relative importance of the features within the feature space. The anatomical and functional features for each component acts as a 10-dimensional address, and the feature space is converted into a “distance space” by computing the distance between the addresses of each pair of components across animals. The learning algorithms are based on analyzing and warping the distances between the compo-




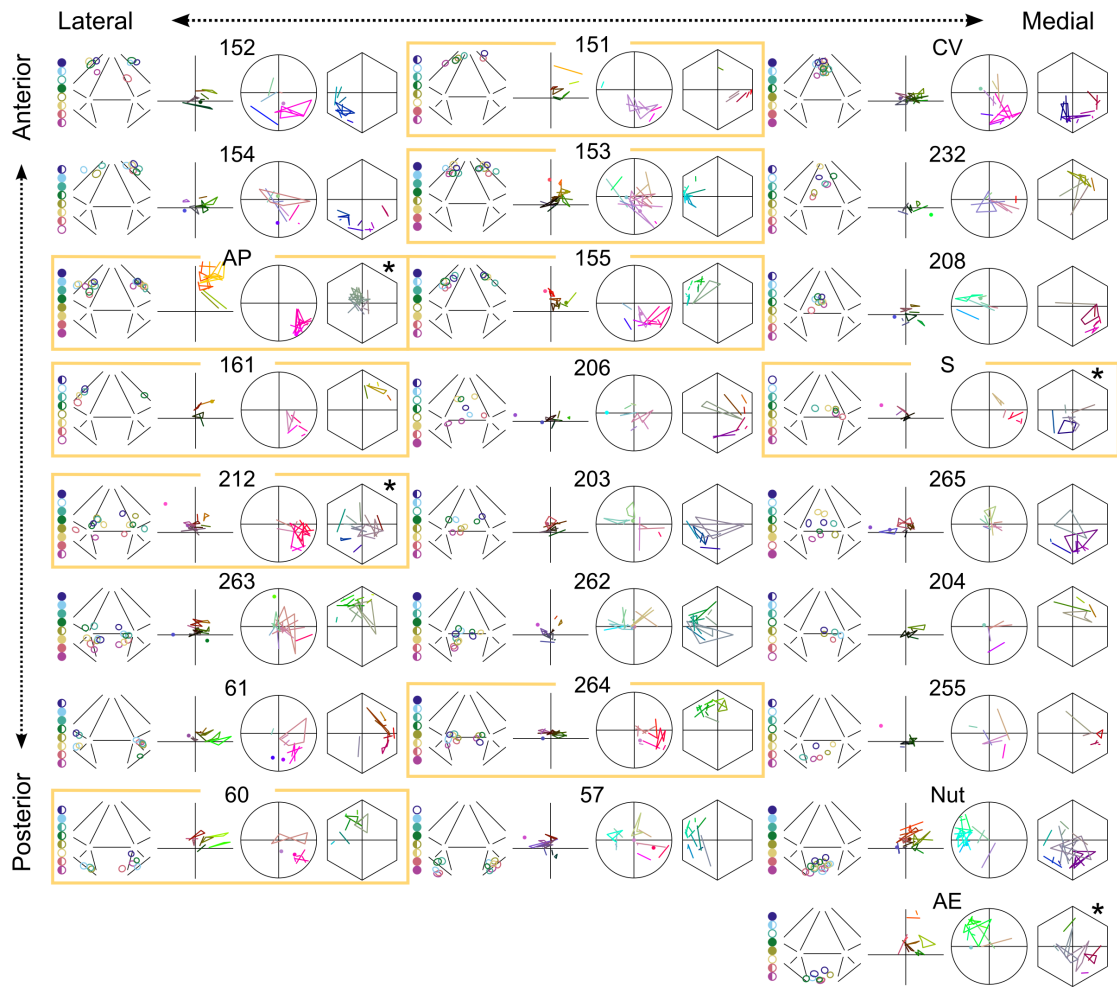
Figure 6.5: Canonical learning of and visualization of medium-dimensional feature space. (A and B) Three example activity maps of the three behaviors are shown from two example experiments (H and C) for visual comparison (top: shortening, middle: local bending, bottom: swimming). (C) Three cells are selected (green: 155, blue: 208, red: 264) by the user indicated by the three colored ROIs. Swim activity map is shown in background. (D) For each of the three selected cells in Animal H (panel C), ICM colors components from Animal C based on how close each component is to the selected component in the warped distance space. The three selections are visualized through the different R,G,B color channels. ICM uses these distances to compute the most likely homologs across animals, and the homolog that is predicted is shown as an ROI of the corresponding color. (E and F) The matches across animals can then be given labels and the components that are common matches for multiple animals are accumulated into a label category. The red colored components indicate components that were put into an identified cell category and have been given their category name or number, and the black components were components that were not identified and have arbitrary numbers.

nents across animals. This warp is encoded by values in the “warp matrix”, which applies an affine transform to all the components’ addresses in a single animal (Supplementarary Appendix). To determine how to warp the feature space, “matches” are chosen across animals and used as landmarks. Based on these matches, the distance space of one animal is warped such that the distance between matches decreases and the distance between non-matches increases. This is directly analogous to using anatomical landmarks and warping a 3D fMRI brain scan into a template brain, except that we define landmarks using both anatomical and functional features and we warp a 10-dimensional space, rather than a 3-dimensional space. This warping allows each ganglion to be translated, rotated, and skewed to be more similar to one another. This procedure compensates for the variations that are due to the positioning or stretching of the ganglion across experiments. Further, the functional features can be similarly rotated and scaled to compensate for variability in the behavior. Phase differences during swimming, for instance, may be caused by natural variability across animals, but variability can also be caused by experimental differences, such as differences in the location of the extracellular suction electrode on the DP nerve. The warping can compensate for this type of functional variability in the same fashion as spatial warping accounts for anatomical variability.

The warp matrix begins as an identity matrix, and the initial matches can be automatically chosen based on the unwarped feature space using the Hungarian Algorithm (Munkres, 1957; Supplemental 2). The automatic identification can accurately match a few cells that have distinguishing features, but this can be greatly improved by incorporating user feedback. Using a GUI, the user can manually set matches (obtained by visually inspecting data like Fig. 6.5A, B), which are used to learn the aspects of the features that are most suitable for identifying cells. Visualizing the warped distance space can aid the user in finding new matches. The interface allows the user to select a component in one animal (Fig. 6.5C, colored circles) and then the interface displays the distance to components from another animal (Fig. 6.5D, color intensity) as well as the match chosen by the match-selection algorithms (i.e Hungarian; Fig. 6.5D, colored circles). The user can test for other possible matches by changing the match manually, which is then used to refine the warped distance space. The use of multiple matches improves the warping, because feature invariances are learned by combining information from many matches across multiple animals.

Using this multiple-iteration, semi-supervised approach, we have identified two canonical networks for swimming and for preparatory activity that have consistent functional and anatomical patterns across animals (Fig. 6.6). The 45 neurons (21 are bilaterally paired and

Figure 6.6: Table of cells in canonical swim oscillator and preparatory network. Through homologous feature matching, we identified 21 canonical swim oscillators, and 9 preparatory neurons. The cells are given numbers based on their locations and best match to other cells in the literature. For each entry, the experiment indicator (8 small colored circles, left), the ganglion position (ganglion), the shortening factors (cross), the local-bend coherence (circle) and the swim coherence (hexagon) are shown. The experiment indicator describes the data sets in which the cell is identified, a full-circle means that both bilateral pairs of cells are present in the ganglion. A semi-circle means only one bilateral pair was seen or the cell was not paired, and an empty circle means the cell was not identified in that experiment. The colors of the experiment indicators are the same colors as the ganglion position ROIs. The position ROIs (ganglion) show the position and size features for each identified cell, and the across animal variability can be seen by the distribution of ROIs. The shortening factors (cross), local-bend coherence (circle) and swim coherence (hexagon) summarize the activity of each cell during the behaviors (see Fig. 6.3) and show the within and across animal variability. For each of these plots, a single connected line or polygon corresponds to a single component across multiple trials. Each corner of the polygon are the feature coefficients from a single trial of a behavior. Each individual polygon is from a different animal. The cells boxed in yellow are part of the preparatory network, and cells that do not show a significant swim oscillation have an asterisk above the swim coherence hexagon. The entries in the table are organized based on their anatomical position.



3 are unpaired) in these two networks are partially overlapping: 5 are active only during the preparatory phase, 28 are active only during swimming, and 12 are active during both behaviors. The entries for these identified cells are organized based on their general anatomical locations, and are given names based on an established ganglionic map (Muller et al., 1981). The entries boxed in yellow are cells involved in the preparatory network, and all entries except AP, 212, AE and S (Fig. 6.6, indicated by asterisk) show consistent oscillations during swimming.

For each identified cell, its unwarped 10-dimensional feature space is an entry in the table, each of which summarizes the anatomical and functional features of an identified cell across all trials in all experiments (Fig. 6.6). Each entry in the table is made up of five elements: the experiment indicator (vertical circles), the position indicator (ganglion), the shortening factors (cross), the local-bending coherence (circle), and the swimming coherence (hexagon). The experiment indicators are the eight vertical circles to the left, and represent in which experiments the cell was identified. These circles correspond to experiments A-H ordered from top to bottom. Most neurons in the ganglion are bilaterally paired; a solid circle indicates that both members of the pair were identified in a particular experiment, and a hemi-circle indicates that only one of them was identified or that the cells are not paired (cells S, 204, and 208). An open circle means that the cell was not identified in that experiment. The position indicator shows the location and size of the ROI of the identified component drawn on a schematic ventral view of the ganglion. In addition, the ROI is colored based on its experiment of origin (A-H; the color code for the ROIs are the same as the experiment indicators). Both the left and right ROIs of an identified bilateral pair are shown together in the position indicator.

The next 3 elements in each neuron's description (Fig. 6.6) show its functional features derived from shortening (cross), local-bending (circle), and swimming (hexagon) (explained in Fig. 6.3). Data from a single experiment is represented by the vertices of a connected polygon. In some experiments only a single trial of a behavior was captured; in these cases, a dot represents the coefficients from that trial. If two trials were captured, they are represented as the start and end of a line segment. Likewise, data from three trials are represented by a triangle, four trials by a rectangle, etc, with each corner of the polygon representing feature coefficients from one of the trials. Because each polygon (or dot or line segment) represents an individual experiment, the within-animal variability is represented as the variance of each polygon's vertices, and the variability across animals is the variance between polygons.

Fig. 6.7 is a projection onto the standard leech ganglion (Muller et al., 1981) of all the neurons that we have detected in this study that oscillate with swimming (left diagram) and that

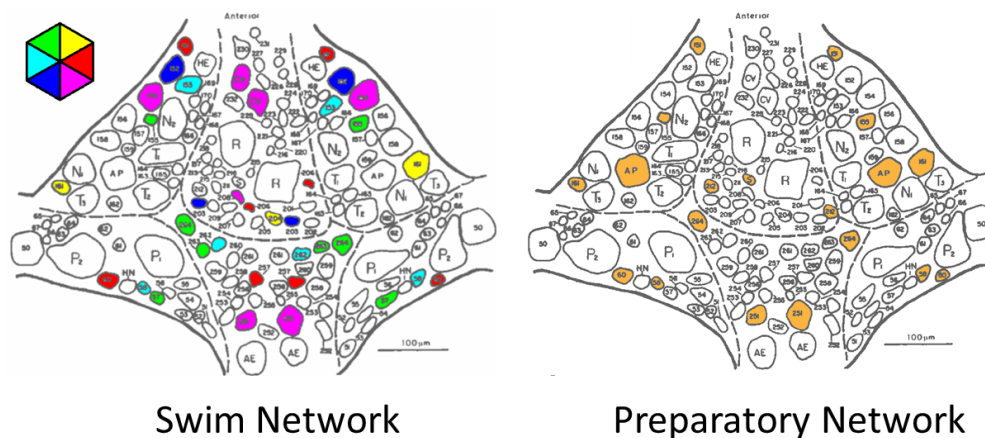


Figure 6.7: Map of canonical swim and preparatory networks. The identified cells that oscillate with swimming are grouped into one of six phases and colored by phase (left). The cells in the preparatory network are shown on the right.

belong to the preparatory network (right diagram). We attempted to match our identified neurons with those previously characterized. Our data is in good agreement with previous work, but it is not necessarily clear that all cells identified here truly correspond to cells with the same labels in the literature (Supplemental 3). Our identifications are in general agreement between VSD-produced phase maps reported previously (Briggman & Kristan, 2006), although the previous study used an older generation of VSD which had a much slower time-constant, which limited the accuracy of the phase information.

The extent and nature of the variability in the neuronal responses is best seen in some representative neurons that are most easily identified. The AP and the Nut cell are quite easy to distinguish simply by eye, because they are typically larger than their neighboring cells and in characteristic locations. The AP cell is located in the anterior lateral packets and is the largest cell in the table (ganglion, AP Entry, Fig. 6.6), which means that it has a high signal-to-noise ratio (SNR) because more pixels are collecting signal. The AP cell showed a consistent rapid depolarization during shortening (cross, AP Entry, Fig. 6.6), strong depolarizing responses to all P cell input (circle, AP Entry, Fig. 6.6), and no consistent oscillation during swimming (hexagon, AP Entry, Fig. 6.6). The Nut cell is somewhat smaller than AP, but is readily identified by its size and posterior medial position (ganglion, Nut Entry, Fig. 6.6). The Nut cell variability is much higher than the AP cell, but this amount of variability is more typical of the rest of the neurons identified. The Nut cell usually shows a rapid depolarization during shortening, but has some variability (cross, Nut Entry, Fig. 6.6). The Nut cell was usually inhibited by P cell stimulus

(circle, Nut Entry, Fig. 6.6), and typically oscillated with peak depolarization in the purple phase during swimming (hexagon, Nut Entry, Fig. 6.6), but occasionally this oscillation was not seen – possibly because of true functional variability or because sometimes the oscillation was too subtle for detection in the VSD signal.

Cell 208 is an unpaired cell that was found in every animal (experiment indicators, 208 entry, Fig. 6.6). One reason that this cell was found consistently is because it is located in the middle of the ganglion, making it much more likely that it would be in the imaging field of view and on the right side of the ganglion (ganglion, 208 entry, Fig. 6.6). It did not show consistent patterns during shortening (cross, 208 entry, Fig. 6.6), but it did receive consistent inhibition from P cell stimulation (circle, 208 entry, Fig. 6.6) as well as showing consistent oscillation during swimming centered on the red phase (hexagon, 208 entry, Fig. 6.6).

The identified neurons that proved most difficult to detect in this data set were cells 204, S, 206 and 203. These cells are very small, and so have low signal-to-noise ratio. We were initially surprised that the S cell was not obvious in our dataset, because it has large overshooting action potentials. However, the S cell action-potentials are unusually brief, lasting only a millisecond. By comparison, the Leydig cell action-potential lasts 30 milliseconds and can be easily seen in the VSD recordings (Fig. 6.1C, bottom trace;). Because they proved to be difficult to detect, cells 204, S, 206 and 203 are the least stringent in this data set, but enough information was extracted that upon close visual inspection these cells can be more confidently identified.

Cells 263 and 264 are nearly indistinguishable from each other, because their features were similar in all 10 dimensions – they were in nearly the same position, they both had a small Factor 2 during shortening, and they both oscillated in the green phase during swimming (compare entries 263 and 264 in Fig. 6.6). Differences in the local bend response were used as the distinguishing dimension between these two cells. However, 263 appeared to receive small input from some P cell inputs, and so it is possible that some of these instances are cross-classified.

The shortening feature shows a large amount of variability within and across animals, but some cells had consistent patterns during shortening. The AP cell, for instance, showed a clear, rapid depolarization during the shortening behavior, along with several other cells in the preparatory network (discussed below). A few cells, such as 57, 265, 262, showed fairly consistent hyperpolarizing responses during shortening (cross, Fig. 6.6).

The local bend feature was measured by activating sensory P cell. However, there are four different P cells and local bend trials were all lumped together, even though different P

cells are activated within and across animals. Some cells, such as 208, Nut, and AE, had purely inhibitory responses to all P cell stimulations, while other cells, such as AP, 212, 264, 155, had only excitatory responses. Other cells, such as 57, 203, 152, and 153 appeared to have both excitatory (pink phase) and inhibitory (cyan phase) responses.

The swim oscillations of neurons were the most robust feature and many of the swim oscillators can be identified using only position and phase of the oscillation. We performed several electrophysiological verification studies to validate the cells identified through VSD recordings and to show that these cells can be identified and targeted in novel animals (Fig. 6.8). We first imaged VSD signals during swimming and created a swim activity map in real-time (Fig. 6.8A, C, E). We used this activity map to target specific swim oscillators with intracellular electrodes and recorded their membrane potentials during a bout of swimming (Fig. 6.8B, D, F). We performed the same coherence analysis on the intracellular recordings and all 6 of the cells targeted have the same phase oscillation as predicted by the VSDs. Cell 153 is the easiest of the swim oscillators to identify, because it has a large 10-20 mV oscillation that is consistently in the anti-phase of the DP swim burst (hexagon, 153 entry, Fig. 6.6; Fig. 6.8F). Other cells, like 152 (Fig. 6.8B) and 203 (Fig. 6.8D), have oscillations of only a few millivolts.

We noticed that many of the neurons shared a common set of responses to stimulation: rapid onset excitation during shortening (positive shortening Factor 2, cross, Fig. 6.6), and excitation by P cell input (pink phase, circle, Fig. 6.6). These feature values mean that these cells were activated both by local sensory input, and by the initial input from distant stimuli that elicited the shortening response. We inspected these responses more closely, and measured the response latencies, identified as a sudden change in the VSD traces in response to shortening or swimming stimuli. We saw that many of these neurons showed rapid-onset responses to all stimuli, regardless of which DP nerve was stimulated or which behavior resulted. This means that many of the cells show a rapid sensory-driven response profile from both local and distal stimuli that is independent of the behavioral response. This suggested to us that these cells may be involved in preparing the animal for a rapid execution of a behavior and we deemed these cells as the canonical “preparatory network” (yellow boxed entries, Fig. 6.6).

For this data set, we stimulated different DP nerves and plotted the responses to stimuli at different distances from the stimulus site (Fig. 6.9A,B). The ganglion imaged was always ganglion 10, so the traces were sorted based on the distance of the stimulus from ganglion 10, either anteriorly (ganglia 3 and 7) or posteriorly (ganglia 11, 14, 17, and Tail Brain). As the stimulus distance increases away from ganglion 10, the response latency increases. We measured

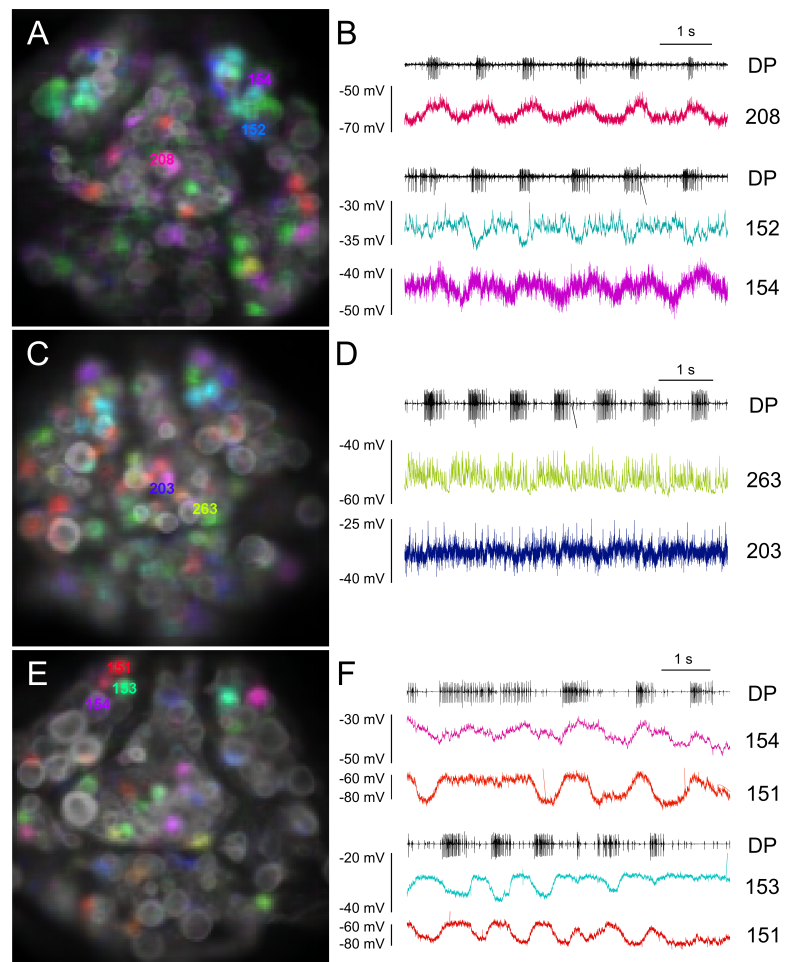


Figure 6.8: Electrophysiological verification of targeted swim oscillator neurons. (A) An activity map was generated from a VSD imaging trial of swimming and three cells (208, 152, 154) were identified and targeted for electrophysiological recordings. (B) Swimming was activated while intracellular recordings were made to targeted cells. The coherence of the intracellular recordings was computed and the traces are colored by the coherence phase in the same fashion as the VSD recordings. (C-F) The same verification experiment was carried out in two more animals targeting different swim oscillators. The electrophysiological recordings of these sample cells are in agreement with the phases predicted by the VSD mapping.

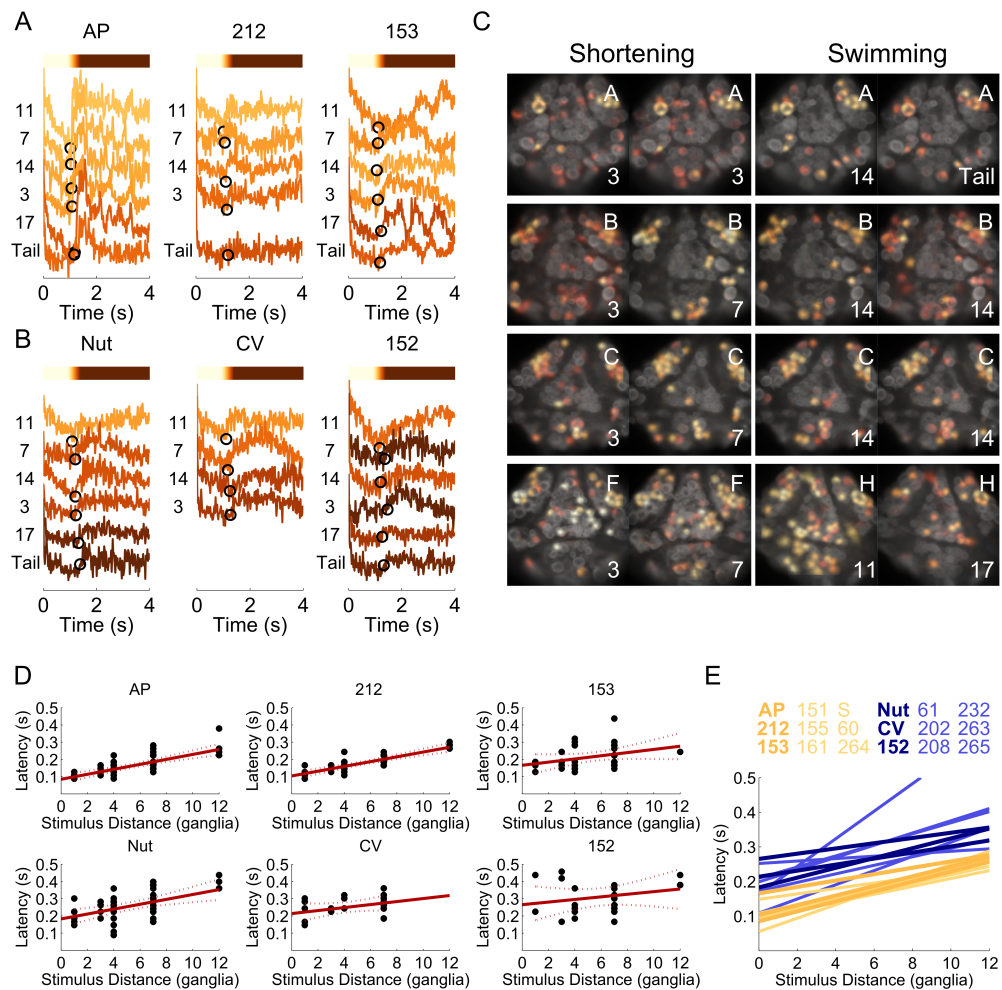


Figure 6.9: The preparatory network shows rapid responses to the stimulus regardless of behavior. (A and B) Example VSD traces from three identified cells within the preparatory network (A) and three that are not (B). Each trace is sorted based on the distance the stimulus was from ganglion 10. The ganglion stimulated is shown to the left of each trace. The response latency is shown as a black circle, and this determines the color of the trace (color bar legend, above traces). (C) Activity maps of the response delays for each cell for several examples of shortening and swimming. Each component is colored by the response delay from one trial. Examples are shown for both shortening and swimming and for different stimulated ganglia. The letters at the top right indicate which experiment, and the numbers to the bottom right indicate which ganglion was stimulated. (D) The response latency of the example cells from all experiments are plotted against the stimulus distance from the imaged ganglion. The red line is a linear fit to the response latency. (E) The linear fits of the response latency for all 9 cells in the preparatory network are plotted in yellow. The fits of the three examples cells that are not part of the preparatory network are plotted in blue.

the responses of the preparatory network neurons as well as some identified cells that are not in the preparatory network but are activated during shortening and/or swimming behaviors. The three examples of cells within the preparatory network (Fig. 6.9A) and three that are not (Fig. 6.9B) are colored based on their response delays (black circles) according to the color bar above each plot. From these examples, it can be seen that the preparatory network examples had a shorter response latency than the other neurons.

We visualized the preparatory network by making activity maps of the response latencies during both shortening and swimming behaviors (Fig. 6.9C). Cells with response latencies less than 300 ms are colored based on their response latency using the same color code as in Figs. 6.9A and 6.9B. The activity maps show several example trials with different behaviors and different stimulated ganglia, in which stimulation of ganglia 3 and 7 elicited shortening, and stimulation of 11, 14, 17, and the Tail Brain elicited swimming. The activity maps within the same animal showed remarkable consistency regardless of the behavior elicited or the stimulus location, suggesting that the preparatory network is rapidly activated by all stimuli that produce any response.

To quantify these impressions of the preparatory network timing, we pooled all data across experiments for each cell and computed linear fits of the response latencies relative to the ganglion stimulated (Fig. 6.9D). In Figure 6.9E, we plot the fit line of the response latency vs distance for every cell in the preparatory network as well as the three other example cells. This analysis showed that all 9 of the cells in the preparatory network shared this consistent and rapid response. There was one cell which showed a rapid onset to shortening as well as input from P cells (cell 61, Fig. 6.6), possibly indicating that this cell was part of the preparatory network. However, on closer inspection this cell did not have as rapid responses as the other cells in the preparatory network, and so we excluded it. Cells 151 and 153 did show these rapid responses to all sensory stimuli, however unlike other cells in the preparatory network these cells occasionally showed hyperpolarizing responses to P cell stimulation.

The S cell is a prime candidate for being the mediator of the rapid responses of the preparatory network. In each ganglion there is an unpaired S cell, which forms strong electrical connections to the S cells in the neighboring ganglia. This connectivity pattern creates a chain of electrically coupled cells that extends through the entire animal. The S cell network is reminiscent of a “giant fiber system” or a fast conducting system (Mistick, 1978), but its causal role in any behavior is unclear (Sahley et al., 1994) and activation of the S cell network does not elicit any particular behavior (Shaw & Kristan, 1999). However, this rapid network may be useful in

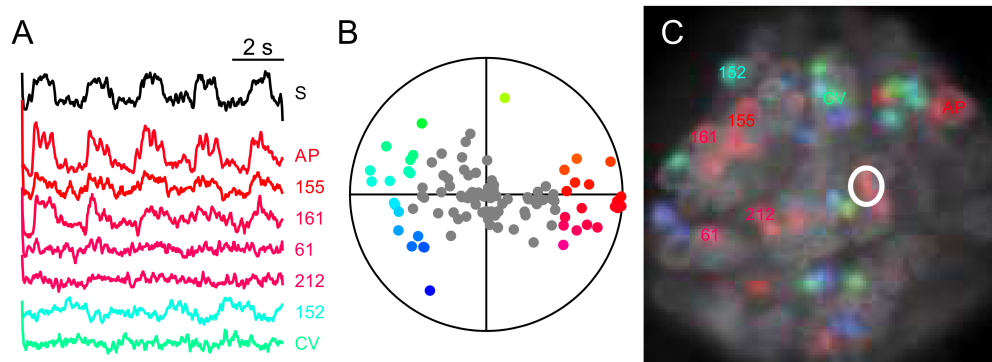


Figure 6.10: The S cell is functionally coupled to the preparatory network. (A) The S cell was stimulated with a 0.5 Hz square wave of current (black is optical trace of S cell). Several traces from identified cells are shown as examples. AP, 155, 161, and 212 are part of the preparatory network. (B) The coherence of all components was computed against the S cell, indicating cells that are functionally excited (red dots) and inhibited (cyan dots) by S cell stimulation. (C) S coupling activity map. Several cells are first identified based on their activity during swimming and shortening (colored text) and the S cell is targeted for stimulation (white circle). Components are colored based on their coherence in panel B.

putting the nervous system and the muscles in a ready state so that one of multiple behaviors can be rapidly executed. From our data, the S cell was an early-activated neuron, so we tested whether other cells in the preparatory network received excitation from the S cell.

To perform this test, we first needed to identify the cells in the preparatory network in each ganglion. Using the framework presented, we activated swimming, shortening and local-bending to identify these neurons within minutes. We then recorded from the S cell with an intracellular electrode and excited it with pulses of depolarizing current while imaging the other neurons in the ganglion (Fig. 6.10A). Many neurons responded to the S cell stimulus, and we used coherence analysis to extract the phases of the traces for all of the cells recorded (Fig. 6.10B). This analysis showed that many neurons are functionally excited (red phase) and inhibited (cyan phase) by S cell activation. These colors were used to make an activity map (Fig. 6.10C). We show in Figure 6.10 examples of some of the identified cells that respond to S cell stimulation. Notably, several identified cells in the preparatory network (AP, 155, 161, 212) proved to be excited by S cell activity as revealed by VSD imaging.

In order to validate the connectivity of the S cell predicted by the voltage-sensitive dyes, we performed several follow-up electrophysiological experiments. Fortunately, the S cell has a unique action potential that allows it to be identified readily using electrophysiology. We first tested for the S cell to AP cell connection because both neurons are readily identified in single

ganglia, without need of VSD recordings. Injecting pulses of current into the S cell to activate action-potentials showed a clear one-for-one EPSP for each S cell spike (Fig. 6.11D).

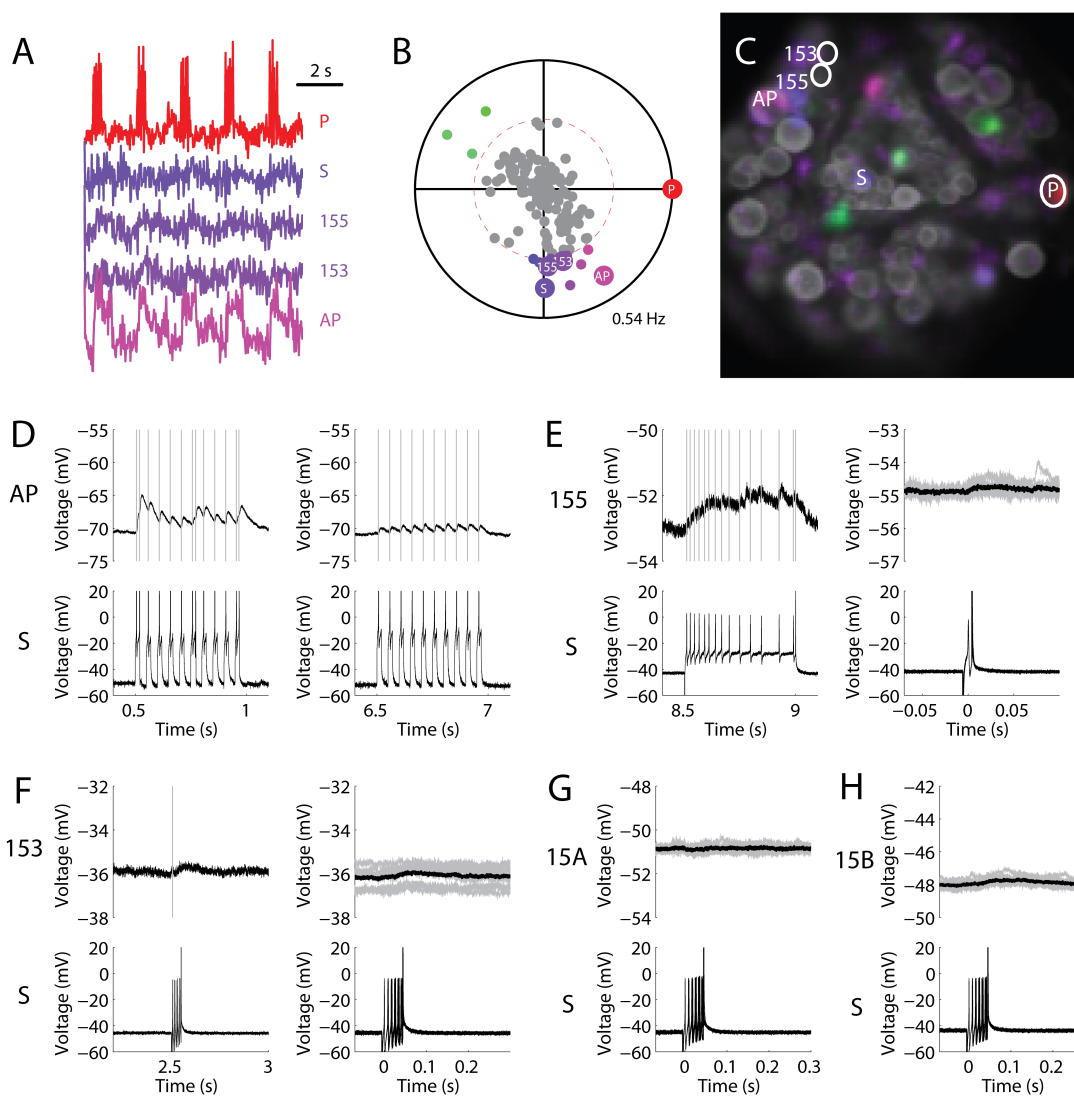
In a second experiment, we activated a P cell and monitored voltage-sensitive dye activity (Fig. 6.11A-C) to identify cells 155 and 153 (Fig. 6.11A-C), aided by the real-time generation of an activity map (Fig. 6.11C). S cell spikes produced clear excitatory effects on both neurons, but these synaptic potentials were much smaller than the EPSPs made by the S cell onto the AP cell, and the one-for-one nature of the S cell inputs onto cells 153 and 155 was less clear (Fig. 6.11E, F). We also targeted two other neighboring cells, which were not identified, and these cells showed no significant response to S cell activation (Fig. 6.11F, G).

Discussion

We used voltage-sensitive dye imaging and electrophysiology to identify and map the activity of dozens of neurons in a scalable fashion (Fig. 6.1), using information extracted from the multi-functional nature of leech neuronal activity during different behaviors (Fig. 6.3; Briggman & Kristan, 2006). To do this, we developed computational tools to automatically extract signals from imaging data (Fig. 6.2), to rapidly create activity maps for visualization (Fig. 6.4), and to build a rich medium-dimensional feature space that summarizes large-scale data and can be used to stitch experiments together into a standard form (Fig. 6.5). We used this information to identify dozens of cells and the canonical swimming and preparatory networks (Fig. 6.6 and 6.7).

The neurons we identify in this paper do not include many of the neurons that have been previously identified and characterized (we mostly ignored the colored cells in Figure 6.1B, although we include previously identified cells if they were part of the swim or preparatory network). The ignored neurons included most of the sensory neurons and a few neurons that have unique features. Many of these cells can be identified simply by eye – such as the Retzius and P cells, because their anatomical features are distinguishing (both cells are much larger than other cells that are nearby). In fact, we removed the somata of Retzius cells (the largest somata on the ventral surface), because they sometimes cover the somata of other, smaller neurons that we wanted to image. From previous studies (Loer & Kristan, mid-1980's), most stimuli that produce motor neuron activity patterns also produce vigorous excitation of Rz cells. Many other cells can be identified by unique action-potentials that can be picked up by the VSD recordings – for instance the N cell's long after-hyperpolarization is an obvious unique feature, and the Leydig cell spikes are also easily seen and highly characteristic (Fig. 6.1C). Another cell, HE,

Figure 6.11: Electrophysiological validation of predicted S cell connections. (A) In order to identify preparatory network cells, a P cell is activated with a burst of spikes every 2 seconds and the ganglion is monitored with VSDs. Preparatory cells S, 155, 153, and AP can be identified due to their size, position and the fact that they receive P cell input. (B) The coherence of the components was calculated to quickly identify P cell followers. (C) The activity map highlights cells that respond to the P cell stimuli, allowing rapid identification of preparatory cells of interest. This map was used to then target cells with electrophysiology. (D) The S cell and AP cell were targeted with sharp electrodes, the AP cell was slightly hyperpolarized to below spike threshold. Two example traces show one-for-one spiking to large 2-3 mV EPSPs in the post-synaptic AP cell. (E) Cell 155 was then targeted for a paired physiology experiment. Exciting the S cell with current showed a clear excitatory response in the post-synaptic 155, but the one-for-one spiking is unclear. Further, the strength of the synapse is significantly smaller (only about 0.5 mV). (F) Another preparatory network cell, 153. A small excitatory response can be seen. (G and H) Two other cells nearby to 153 and 155 were targeted, these cells were unidentified. There is very little to no response in the other neighboring cells.



which is involved in controlling the heart, oscillates to pump blood and this oscillation can be picked up in the VSD recordings and used to identify the cell. Even though these cells can be identified, we did not include them because their identifying features are obvious and these cells have been previously well characterized.

Due to variability across preparations, both because of variations in cell body locations and dye loading, VSD data from a single experiment includes information only from a subset of all of the neurons in the ganglion. Cells at the lateral edges of the ganglion are likely to be sampled less often, because they can often appear on the opposite face of the ganglion. The smaller cells in the ganglion are also undersampled, because smaller cells have lower SNR in the optical signals. So even though small cells may be in the field of view, their activity can be missed because of low SNR. This also makes it harder in general to identify small cells as less information can be extracted from their optical signals.

The computational tools we developed relied extensively on user feedback and refinement. At each of the three stages of the framework, the experimentalist interacts with computational algorithms. Visualization of the results of the PCA-ICA decomposition is extremely useful to understand imaging data, and visualization tools are used in a wide variety of scientific contexts (Walter et al. 2010). Being able to manipulate settings and remove certain components greatly increased the SNR and pulled out signal in the midst of artifacts. Adjusting the parameters is also necessary, and it is useful to visualize the ICA analysis with different sets of parameters and numbers of components. Many algorithms could also be added to further improve the ICA results and speed, such as automatic identification of artifacts versus cells. For instance, cells could be automatically sorted from artifacts by looking for spatially localized signals of appropriate size.

Once feature detectors are designed, visualizations of activity maps can be produced from data in a few minutes. With experience, visualizations become an essential tool for guided experiments. A few simple visualizations often proved the most efficient way of identifying a particular cell of interest. Virtually all of the swim oscillators can be identified based on their swim phase alone, because the phases of the swim oscillators along with soma position provides enough information to identify these cells.

Often, the goal of machine-learning algorithms is to transform high-dimensional data into a rich medium-dimensional feature space. Here we have crafted such a rich feature space by using simple algorithms to extract features from three distinct behaviors. The desired goal of machine-learning algorithms is to transform cell traces into a feature space that is only a linear

transform away from the final output. When the final output of the system is a linear transform away, then the feature space represents the structure of the information in an efficient manner related to the output. Efficiently reducing high-dimensional data into a medium-dimensional feature space is indirectly related to extracting the underlying causal structure (Schmidhuber, 2014). Machine learning can be used to help find this structure, but we took a slightly different approach. Rather than developing sophisticated algorithms to build our feature space, we relied on simple features from a distinct set of behaviors. Theories of population coding suggest that neurons represent information in their joint activities, and that neurons use multi-dimensional receptive fields to represent multi-dimensional information. This means that neurons are often multi-functional, and these multi-functional responses are seen in the leech. Agreeing with population coding theory, it seems that receptive fields of neurons are often spread out across the parameter-space. Such broad tuning maximizes the robustness of the neural code. This broadness proved to be useful for our efforts to map the activity of the ganglion, as spread-out receptive fields mean that neurons will be more distinct from one another and can be identified. We did not need to build sophisticated machine-learning algorithms (our approach was rather simple), because simple algorithms extracting information from a diverse set of conditions/contexts provides a lot of information.

A medium-dimensional space means that the number of features describing the data is relatively small compared to the amount of data. In our case, we recorded from about 100 cells from eight different animals giving us 800 data points that we describe in a 10 dimensional feature space. This relative amount is what makes a feature space “medium-dimensional”. If the feature space is “high-dimensional”, meaning that is near the same dimensionality of the data, there is a risk of over-fitting because the feature space can contain both the common structure as well as the noise that is particular to the data. If the feature space is too low-dimensional, then there is not enough information to represent the important structure of the data. (A general rule-of-thumb is that the feature space be on the order of the square-root of the dimensionality of the data (Schrodinger).

The preparatory network was identified through its consistent responses to sensory stimuli. This network stood out because these cells were tightly locked to the sensory input. Because we were providing the sensory stimulus, we knew its exact timing and sensory correlated responses were easily detected. Both the shortening and local-bending features provide dependable information on cells that are reliably driven by sensory input. The preparatory network was surprising because it was activated by every type of sensory stimuli, whether medial or lateral P

cell activation during local bend, or anterior or posterior DP nerve activation during shortening and swimming. We investigated the responses of these cells closely (Figure 6.9), and discovered that the same network was rapidly activated prior to both shortening and swimming. The rapid and consistent sensory nature of these cells have lead us to believe that this is the “Do something” (Friesen & Kristan, 2007) or preparatory network, which we hypothesize readies the nervous system and musculature for a more rapid induction of a behavioral response. The S cell has been carefully studied because of its role in an apparent “Giant Fiber System” in the leech (Shaw & Kristan, 1999), and we hypothesized that the S cell was the mediator of the preparatory network responses. Previous work has shown that it is a challenge to halt or prevent the S cell spike, and showing the S cell was necessary for the preparatory response by blocking its spikes is currently beyond our technical ability. However, with our scalable identification framework, we were able to identify the S cell and the other cells in the preparatory network and then test sufficiency with dozens of cells identified in context (Fig. 6.10). We were able to show that the S cell is functionally coupled to at least 4 cells in the preparatory network (AP, 155, 161, 212) and verified these and other connections using computationally guided electrophysiology (Fig. 6.11). The AP cell (which is the most obvious cell in this network) has been studied in intense detail (Wessell et al. 1999a,b), but its function remains unknown. Further study into how the AP cell may be involved in a preparatory behavior is warranted.

The techniques we have introduced here can also be considered in computer science terms by considering the complexity of “brain activity mapping” experiments. Complex systems will require a large synthesis of information before they can be understood. This means that information must be extracted from the system, and the experimental methods are the “algorithm” used to extract information from the system. For instance, if one wanted to know the polarity of every synaptic connection in the leech ganglion using only paired intracellular recordings, then one would have to perform a paired recording between every pair of cells. This algorithm has a complexity of $O(N^2)$ for solving this problem, because N^2 experiments would have to be performed to know the connectivity of N neurons. Using optical techniques where essentially all of the cells can be recorded from simultaneously, one would only have to perform a coupling experiment for each of the N cells and the algorithm’s complexity drops to $O(N)$. Note the real payoff isn’t quite so grand, because the resolution of the VSD imaging is lower than the electrophysiology, which means more VSD experiments will be required to get the same amount of information as a single intracellular recording. But even if you had to do 10 VSD experiments for every intracellular recording, that is small in comparison to the scale of the data – even just

the 100 cells as in our recordings is large in comparison to 10. The difference vanishes as data from experiments get larger in scale and the dimensionality of the data dominates the complexity of the algorithm.

Acknowledgments

This chapter is presented as a manuscript that is in preparation as a publication as Frady, E.P., Kristan Jr., W.B. “Scalable semi-supervised cell identification reveals canonical swim and preparatory networks.”

References

Alivisatos, AP. Chun, M. Church, GM. Greenspan, RJ. Roukes, ML. Yuste, R. (2012) The Brain Activity Map Project and the Challenge of Functional Connectomics. *Neuron* 74:970-974.

Bagnoli, P., Brunelli, M., Magni, F. (1975) The neuron of the fast conducting system in *Hirudo medicinalis*: identification and synaptic connections with primary afferent neurons. *Arch. Ital. Biol.* 113:21-43.

Bernard, A., Sorensen, S.A., Lein, E.S. (2009) Shifting the paradigm: new approaches for characterizing and classifying neurons. *Current Opinion in Neurobiology* 19:530-536. DOI: 10.1016/j.conb.2009.09.010

Bokil, H., Andrews, P., Kulkarni, J.E., Mehta, S., Mitra, P.P. (2010). Chronux: A platform for analyzing neural signals. *Journal of Neuroscience Methods* 192: 146-151.

Boyd, Y., Veksler, O. and Zabih, R. Fast Approximate Energy Minimization via Graph Cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2001).

Briggman, K. L., Abarbanel, H. D. I. and Kristan, W. B. Optical Imaging of Neuronal Populations during Decision-Making. *Science* 11 (2005).

Briggman, K.L., Kristan, W.B. (2006) Imaging Dedicated and Multifunctional Neural Circuits Generating Distinct Behaviors. *J. Neurosci.* 26(42):10925-10933.

Brodhauer, P.D., McCormick, K., Tapyrik, L., Albano, A.M., Graybeal, C. (2008) Activation of two forms of locomotion by a previously identified trigger interneuron for swimming in the medicinal leech. *Invert. Neurosci.* 8:31-39.

Bunke, H. Recent Developments in Graph Matching, *International Conference on Pattern Recognition* (2000).

Evangelidis, G.D., Psarakis, E.Z. (2008) Parametric Image Alignment Using Enhanced Correlation Coefficient Maximization. *IEEE Trans. on PAMI [Pattern Analysis and Machine Intelli-*

gence] 30(10):1858-1865.

Friesen, W.O. (1989) Neuronal control of leech swimming movements I. Inhibitory interactions between motor neurons. *J. Comp. Physiol. A.* 166:195-203.

Friesen, W.O. (1989) Neuronal control of leech swimming movements II. Motor neuron feedback to oscillator cells 115 and 28. *J. Comp. Physiol. A* 166:205-215.

Friesen, W.O., Kristan, W.B. (2007) Leech locomotion: swimming, crawling, and decisions. *Current Opinion in Neurobiology* 17:704:711.

Hill, E.S., Moore-Kochlacs, C., Vasireddi, S.K., Sejnowski, T.J., Frost, W.N. (2010). Validation of Independent Component Analysis for Rapid Spike Sorting of Optical Recording Data. *J Neurophysiol* 104: 3721-3731.

Kristan Jr., W.B., Calabrese, R.L., Friesen, W.O. (2005) Neuronal control of leech behavior. *Progress in Neurobiology* 76: 279-327.

Laverack, M.S. (1968) Mechanoreceptors, photoreceptors and rapid conduction pathways in the leech, *Hirudo Medicinalis*. *J. Exp. Biol.* 50:129-140.

Lewis, J.E., Kristan Jr., W.B. (1998a). A neuronal network for computing population vectors in the leech. *Nature* 391, 76-79.

Lewis, J.E., Kristan Jr., W.B. (1998b). Quantitative analysis of a directed behavior in the medicinal leech: implications for organizing motor output. *J. Neurosci.* 18(4), 1571-1582.

Lewis, J.E., Kristan Jr., W.B. (1998c). Representation of touch location by a population of leech sensory neurons. *J. Neurophysiol.* 80, 2584-2592.

Lockery, S.R., Kristan Jr., W.B., (1990a) Distributed processing of sensory information in the leech. I. Input-output relations of the local bending reflex. *J. Neurosci.* 10(6):1811-1815.

Lockery, S.R., Kristan, W.B. (1990b) Distributed Processing of Sensory Information in the Leech. II. Identification of Interneurons Contributing to the Local Bending Reflex. *J. Neurosci.* 10(6): 1816-1829.

Lovasz, L. Random walks on graphs: a survey. *Combinatorics: Paul Erdos is Eighty 2* (1993).

Macagno, E.R. (1980) Number and distribution of neurons in leech segmental ganglia. *J. Comp. Neurol.* 15:283-302.

McKay, B. D. Practical Graph Isomorphism. *Congressus Numerantium* (1981).

Miller, E.W., Lin, J.Y., Frady, E.P., Steinbach, P.A., Kristan Jr., W.B., Tsien, R.Y. (2012) Optically monitoring voltage in neurons by photo-induced electron transfer through molecular wires.

PNAS 109(6): 2114-2119.

Minka, T. Divergence Measures and Message Passing. Microsoft Research Tech Report (2005).

Mistick, D. (1974) Rohde's fiber: a septate axon in the leech. *Brain Res.* 74: 342-348.

Mistick, D. (1978) Neurones in the leech that facilitate an avoidance behavior following nearfield water disturbances. *J. Exp. Biol.* 75:1-23.

Mitra, P.P., Bokil, H. (2007) *Observed Brain Dynamics*, Oxford University Press, USA.

Mukamel, E. A., Nimmerjahn, A., Schnitzer, M. J. (2009). Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data. *Neuron.* 63(6):747-760.

Muller KJ, Nicholls JG, Stent GS (1981) *Neurobiology of the Leech* (Cold Spring Harbor Laboratory, New York).

Munkres, J. Algorithms for the assignment and transportation problems. *J. Soc. of Industrial and Applied Mathematics* 5 (1957).

Nusbaum, M.P., Friesen, W.O., Kristan Jr., W.B., Pearce, R.A. (1987) Neural mechanisms generating the leech swimming rhythm: Swim-initiator neurons excite the network of swim oscillator neurons. *J. Comp. Physiol A.* 161: 355-366.

Nusbaum, M.P. (1986) Synaptic basis of swim initiation in the leech III. Synaptic effects of serotonin-containing interneurons (cells 21 and 61) on swim CPG neurones (cells 18 and 208). *J. exp. Biol.* 122:303-321.

Palmer, C.R., Barnett, M.N., Copado, S., Gardezy, F., Kristan Jr., W.B. (2014). Multiplexed modulation of behavioral choice. *J. Exp. Bio.* 217: 2963-2973.

Sahley, C.L., Modney, B.K., Boulis, N.M., Muller, K.J. (1994) The S cell: an interneuron essential for sensitization and full dishabituation of leech shortening. *J. Neurosci.* 14:6715-6721.

Schmidhuber, J. (2014) *Deep Learning in Neural Networks: An Overview*. Preprint IDSIA-03-14 / arXiv:1404.7828 [cs.NE]

Shaw, B.K., Kristan Jr., W.B. (1999) Relative Roles of the S Cell Network and Parallel Interneuronal Pathways in the Whole-Body Shortening Reflex of the Medicinal Leech. *J. Neurophysiol.* 82:1114-1123.

Taylor, A.L., Cottrel, G.W., Kleinfeld, D., Kristan Jr., W.B. (2003) Imaging Reveals Synaptic Targets of a Swim-Terminating Neuron in the Leech CNS. *J. Neurosci.* 23(36):11402-11410.

Walter, T. Shattuck, DW. Baldock, R. Bastin, ME. Carpenter, AE. Duce, S. Ellenberg, J. Fraser, A. Hamilton, N. Pieper, S. Ragan, MA. Schneider, JE. Tomancak, P. Heriche, JK. (2010) Visu-

alization of image data from cells to organisms. *Nature Methods* 7(3)

Weeks, J.C., Kristan Jr., W.B. (1978) Initiation, maintenance and modulation of swimming in the medicinal leech by the activity of a single neurone. *J. exp. Biol.* 77: 71-88.

Weeks, J.C. (1982) Segmental specialization of a leech swim-initiating interneuron, cell 205. *J. Neurosci.* 2(7): 972-985.

Wessel, R., Kristan Jr., W.B., Kleinfeld, D. (1999a) Dendritic Ca^{2+} -activated K^{+} conductances regulate electrical signal propagation in an invertebrate neuron. *J. Neurosci.* 19:8319-8326.

Wessel, R., Kristan Jr., W.B., Kleinfeld, D. (1999b) Supralinear summation of synaptic inputs by an invertebrate neuron: dendritic gain is mediated by an inward rectifier K^{+} channel. *J. Neurosci.* 19:5875-5888.