Functionality, Robustness and Control of Nonlinear Network Dynamics: Modeling and Understanding the *C. elegans* Connectome

James Michael Kunert

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

Doctor of Philosophy

University of Washington

2016

Reading Committee:

J. Nathan Kutz, Chair

Marcel den Nijs

Paul Wiggins

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Functionality, Robustness and Control of Nonlinear Network Dynamics:
Modeling and Understanding the *C. elegans* Connectome

James Michael Kunert

Chair of the Supervisory Committee:
Professor J. Nathan Kutz
Applied Mathematics

Networks of many nonlinearly-coupled dynamical components are ubiquitous in the physical sciences, but often difficult to characterize. However, their dynamics are often low-dimensional, being dominated by a few functional, coherent patterns. We wish to understand: (1) How do nonlinear networks generate functional responses? (2) What role does the network's *structure* play in generating such responses? (3) To what extent are the network dynamics robust to network damage?

Towards these ends we model the *C. elegans* neuronal network, the connectivity of which is known. Chapter 2 constructs a full-Connectome dynamical model which can generate proxies for known behaviors (specifically demonstrating a proxy for forward motion). Chapter 3 explores the input space via interpretable bifurcation diagrams. The highly multistable dynamics give rise to long transient timescales (orders of magnitude longer than intrinsic nodal timescales). Chapter 4 models network injuries, which significantly distort dynamics. We develop a metric to quantify the injury level and help predict an injury's functional outcome. Chapter 5 uses Dynamic Mode Decomposition to relate connectivity to low-dimensional dynamical structure. In the process, we demonstrate consistency with proprioception-driven locomotion which is facilitated by network structure.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I wish to express my deep gratitude to my advisor, J. Nathan Kutz, for his role as a research advisor and mentor over the years. He not only made this research possible, but opened many opportunities to me professionally, and in his role as a mentor he helped me to develop into the researcher who I am today.

Similarly, I want to thank past mentors who, at earlier stages of my academic career, believed in me and granted me many opportunities to succeed. I would especially like to thank Stephen Gregory from the University of Oregon, along with Dee Winn and Steve Mackey from Umpqua Community College.

The research in this manuscript was only possible due to the guidance and input from my collaborators, and I would like to thank Eli Shlizerman, Josh Proctor, Steve Brunton and Pedro Maia for their advice and support. I am also indebted to the members of my Committee: Marcel den Nijs, Adrienne Fairhall, Jason Detwiler, Paul Wiggins and Bing Brunton.

Lastly, I want to acknowledge my friends and family, who have been so important throughout my graduate school career as sources of guidance and support. My partner Rachel and I, in particular, have relied on each other's support throughout each of our graduate programs. This work would not have been the same without the balance she's brought to my life.

DEDICATION

to my parents, Mike and Mary Kunert

Chapter 1

INTRODUCTION: THE *C. ELEGANS* CONNECTOME

1.1 Nonlinear Physical Networks: Exploring Functionality and Robustness

1.1.1 Nonlinear Networked Systems

Complex physical systems comprised of a network of nonlinear dynamical components of voltage activity are capable of producing robust functionality and/or low-dimensional patterns of coherent activity. By "low-dimensional", we mean that as the state of the system dynamically evolves, it is well-described by a superposition of a small number of activity patterns (or "modes"), much smaller than the apparent dimensionality of the system. Understanding the functional responses and control of these high-dimensional networked dynamical systems is of critical importance across the physical, engineering and biological sciences. These systems are often composed of a diverse set of dynamic elements whose connectivity are prescribed by sparse and/or dense connections that are local and/or long-range in nature. For many systems of interest, the diversity in connectivity and dynamics make it extremely challenging to characterize dynamics on a macroscopic network level.

In many such systems, even with large numbers of nonlinearly interacting nodes, meaningful input and output are dominated by low dimensional spatio-temporal patterns of activity [59, 69, 113]. The coherent swing instability in power grid networks [113], for instance, is an example of these phenomena which have been observed in experiments and computational studies, yet are difficult to characterize with theoretical techniques. Other examples of interacting dynamical systems that are well-known in physics, and that produce functional behavior or coherent patterns, include coupled oscillators (e.g. the Kuramoto oscillators), analog circuits, coupled lasers, many-particle systems, etc. Biophysical systems, whose interactions are often driven by chemical reactions, voltage activity, and/or ion exchange, produce

similar functionality and structured activity.

When a high-dimensional system exhibits such low-dimensional output structure, the underlying networked dynamics can be understood as trajectories which primarily evolve on low-dimensional attractors and/or induced transient trajectories between attractors. As a specific biophysical example, neuronal networks, which are typified by high-dimensional networks of neurons, display robust functional responses and behavioral assays that are encoded by such low-dimensional attractors or transient trajectories [61, 95, 93, 94, 72, 49, 105, 68].

1.1.2 Neuronal Networks

Neuro-sensory networks are ideal for characterizing the role of seemingly complex network interactions for producing robust functionality and can motivate bio-inspired engineering principles. Neuro-sensory integration, which attempts to understand the neural pathways from input stimuli to motor-neuron driven behavioral responses and low-dimensional movements, is one of the most challenging and open problems in the field of neuroscience today. The primary challenge lies in understanding how large networks of different classes of neurons (e.g. sensory-, inter- and motor-neurons which can be either inhibitory or excitatory) interact to produce the observed robust behavioral responses to stimuli. Ultimately, the biophysical processes produce a large, nonlinear network of electronic components that dynamically decode input stimulus and drive downstream neuronal function and behavior.

Of great interest in biological settings is the fact that such complex networks often produce robust and low-dimensional functional responses to dynamic inputs. Indeed, the structure of their large connectivity graph can determine how the system operates as a whole [129, 89]. Neuronal networks, in particular, may encode key behavioral responses with low-dimensional patterns of activity, or population codes, as they generate functionality [12, 23, 29, 74, 71, 32].

1.2 C. elegans *as Model Organism*

1.2.1 About the Worm

Caenorhabditis elegans is a small ($\sim$ 1 mm long) nematode which has long been an important system in the study of genetics, cell biology and neuroscience [3]. In understanding how neural networks generate behavior, it is important partly due to the facts that (1) it possesses only a small number of sensory neurons, often linked to specific stimuli [3], and (2) its range of behavioral responses are varied yet limited, confined to swimming, crawling, turning and performing chemotaxis, for instance. Although *C. elegans* is capable of a wide range of behaviors over various timescales [43], experimental studies suggest that these behaviors are fundamentally low-dimensional [110], and the behaviors of the worm can be understood as low-dimensional trajectories on attractors between which it will transition stochastically [109]. Thus it is reasonable that a complete model of its neuro-sensory integration capabilities could be posited.

Aiding in this effort is the near-complete connectivity data for the gap junctions and chemical synapses connecting the sensory neurons to the inter- and motor-neurons [122]. Moreover, current experiments measure the response of various neurons to input stimuli since a description of these responses cannot be drawn from the static connectivity data alone. These studies suggest that computational modeling can assist in describing neural dynamics and their relation to the connectome. Such efforts may allow for a theoretical characterization of the network biophysics and voltage activity that drives the neuro-sensory integration process in *C. elegans* and determine its ability to elicit behavioral responses [65, 102, 130].

1.2.2 The Connectome

The Connectome of the *C. elegans* hermaphrodite was first characterized by White, Southgate, Thomson and Brenner in 1986 [131]. In conjunction with previous studies detailing the nervous system's structure [2, 127, 128, 134, 132, 54], this work gave a fairly complete picture of the structure of the worm's nervous system. In it, the synaptic and gap junction

connectivities of the neurons were obtained via serial-section electron microscopy. From the resulting images, 5958 chemical synapses and 1106 gap junctions were manually identified and recorded between the 302 neurons [131, 122].

This dataset was pieced together from five separate reconstructions, each covering a different (but overlapping) region of the worm's body (see Figure A1 in [131]), from which a "canonical" Connectome was constructed. This contains some ambiguities: for example, much of the posterior connectivity data was gathered from an adult male, rather than a hermaphrodite (it is unclear how their neural connectivities differ). Technical difficulties also led to further incomplete or inconsistent connectivity data in the original report of the connectome.

Later work, both incorporating new data and re-analyzing the original data, improved the Connectome's completeness and self-consistency [55, 38, 122]. In [122], Varshney et al. incorporate this work and update over 3000 connections. This Connectome data is self-consistent and which they approximate as about 90% complete (with missing connections due to sparse sampling of certain regions for chemical synaptic connections and the relative difficulty of identifying gap junctions via conventional methods). We use this Connectome data from [122] throughout the manuscript.

As in [122], we restrict our efforts throughout this manuscript to the smaller somatic nervous system of 282 neurons (excluding the independent pharyngeal system of 20 neurons), and consider within that only the neurons which make synaptic connections (excluding the synaptically disconnected neurons CANL/R and VC06). This leaves a set of 279 neurons, which have 6393 chemical synapses and 890 gap junctions [122].

1.2.3 Reproducibility of the Connectome

The complete cell lineage of *C. elegans* from fertilized egg to adult has long been characterized, and has been shown to be invariant between animals [11, 22, 111]. Each wild-type hermaphrodite worm will have exactly 302 neurons at reproducible positions.

To study the consistency of the Connectome between animals, White et al.[131] acquired

overlapping connection data for two different animals around the pharynx, in which the majority of synaptic connections are located. R.M. Durbin [38] analyzed the consistency between the connectivity data for these two animals, and found the following: Let A, B denote an arbitrary pair of neurons in the first animal, and A', B' denote the corresponding neurons in the second animal. For an arbitrary pair A, B, there is only a 2.8% chance that A synapses onto B and a 1.3% chance that A and B are connected via a gap junction [122]. If there is a chemical synapse from A to B, however, there is then a 75% chance of a chemical synapse from A' to B'. Similarly, a gap junction between A and B implies a 79% chance of a gap junction between A' and B'. However, almost all inconsistency is within weak connections (with only one or two connections). If A has 3 or more synapses onto B (which is true for approximately two-thirds of connected pairs), then the probability that A' synapses onto B' is 92%, and the probability of a gap junction connection is 98%.

The fact that synaptic connectivity is not entirely invariant is an additional source of ambiguity, given that the current Connectome data is reconstructed from multiple worms. However, the relatively high reproducibility of connections suggests that it is reasonable to treat the current data as a "canonical" network.

1.2.4 A note on Neuron Names

The naming system typically used for *C. elegans* neurons comes from White et al.[131]. Each neuron group was assigned an arbitrary three-letter name (e.g. "AVM"). The third letter is sometimes replaced by a number (e.g. "VB01","VB02",...,"VB11") indicating a relationship between groups. Though arbitrary, some names may be chosen from physical properties of the neurons (e.g. "PLM" for posterior lateral microtubule cells [3]). Each of these groups consists of a set of neurons of the same class and related by geometrical symmetry, being very similar except for their position on the dorsal/ventral or left/right sides. Individual neurons within these groups are thus distinguished by additional letters indicating their position (e.g. neuron group PLM contains two neurons on the right/left sides, "PLML" and "PLMR"; "D" and "V" suffixes are similarly used to indicate dorsal/ventral symmetry).

1.3 Modeling Challenges

Simulations of *C. elegans* neural dynamics are challenging since (1) it it difficult to measure electrical parameters which characterize precisely the directionality and conductance of each connection, and (2) the single neuron dynamics do not appear to be characterized by standard spiking neuron models. Indeed, genomic sequencing and electro-physiological studies have consistently failed to observe classical Na+ action potentials in *C. elegans* neurons [50].

The failure to produce the stereotypical spike train dynamics normally associated with neuronal activity actually allow our model with graded electrical interaction to be more analogous with observed activity in physical systems such as power grids [113], thus broadening the scope of the work and its potential for impact in the physical sciences.

1.4 Low-Dimensional Behavioral Outputs

The fact that electrical parameters of the connections are not well characterized poses an obvious challenge for the creation of a simplified model for Connectome dynamics: without quantitative parameter accuracy, how can one expect to make a comparison to quantitative experimental results? In the time before these parameter values are available from experiment, there are a few possible approaches: (1) one could fit parameters to best match the quantitative output of the model. In large, nonlinear systems with many parameters, however, this is highly susceptible to overfitting, such that the knowledge gained from such an approach is often limited. (2) Alternately, one could impose rough, order-of-magnitude parameter estimates and see if meaningful input/output relationships or behaviors are reproduced *qualitatively*. This manuscript makes use of this latter approach.

Such qualitative comparisons are aided by the fact that the bodily dynamics of [**?**] is fundamentally low-dimensional. Stephens et al. [110] observed the behavior of *C. elegans* and found, by performing Principal Component Analysis on the worm's body shape dynamics, that its body shape is at all times well described by a linear combination of only six body-shape modes, called the "eigenworm" modes. As the worm crawls forward, it oscillates

between two eigenworm modes, forming a ring in phase space. Similar low-dimensional dynamics may be searched for in analyzing the output of computational neural models.

Chapter 2

NONLINEAR NETWORK DYNAMICS: A MODEL FOR THE CONNECTOME

In this Chapter, we develop a biophysical model of neuro-sensory integration in the model organism *Caenorhabditis elegans*. Building on experimental findings of the neuron conductances and their resolved connectome, we posit a full dynamic model of the neural voltage excitations that allows for a characterization of network structures which link input stimuli to neural proxies of behavioral responses. Full connectome simulations of neural responses to prescribed inputs show that robust, low-dimensional bifurcation structures drive neural voltage activity modes. Comparison of these modes with experimental studies allows us to link these network structures to behavioral responses. Thus the underlying bifurcation structures discovered, i.e. induced Hopf bifurcations, are critical in explaining behavioral responses such as swimming and crawling.

2.1 Modeling Connectome Dynamics

2.1.1 Single-Compartment Membrane Model

A model must be constructed for the graded response of neurons. Fortunately, it has been observed that many neurons in *C. elegans* are effectively isopotential, such that we can use the membrane voltage as a state variable for network simulations [50]. The time evolution of neuron i's membrane potential, V_i, is therefore given by the single-compartment membrane equation [133]:

$$C\dot{V}_i = -G^c(V_i - E_{cell}) - I_i^{Gap}(\vec{\mathbf{V}}) - I_i^{Syn}(\vec{\mathbf{V}}) + I_i^{Ext} \tag{2.1}$$

C is the whole-cell membrane capacitance, G^c is the membrane leakage conductance and E_{cell} is the leakage potential. The external input current is given by I_i^{Ext}, while neural interaction

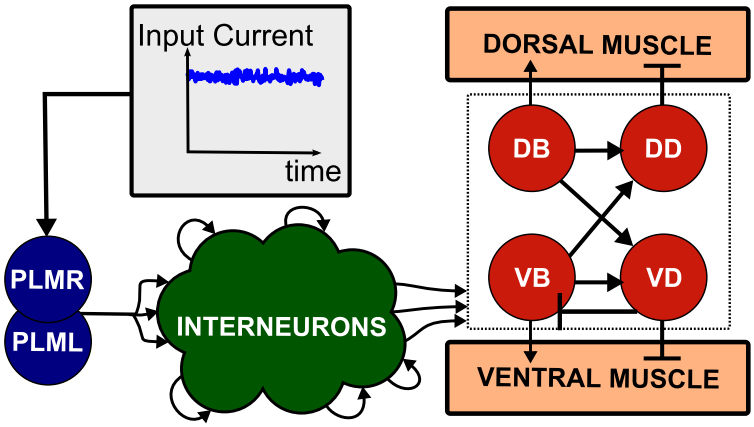

Figure 2.1: In our full network simulations, the neurons receive a physically realistic stimulus (input current) which varies in amplitude for different trials. For the example illustrated, PLM neurons are stimulated leading to signal propagation through a network of densely connected interneurons which activate the motorneuron subcircuits and low-dimensional neural response modes [110] that control forward locomotion [130].

via gap junctions and synapses is modeled by input currents $I_i^{Gap}(\vec{\mathbf{V}})$ (gap) and $I_i^{Syn}(\vec{\mathbf{V}})$ (synaptic). Their equations are:

$$I_i^{Gap} = \sum_j G_{ij}^g (V_i - V_j) \tag{2.2}$$

$$I_i^{Syn} = \sum_j G_{ij}^s s_j (V_i - E_j) \tag{2.3}$$

Gap junctions are taken as ohmic resistances connecting each neuron where G_{ij}^g is the total conductivity of the gap junctions between i and j. Synaptic current is proportional to the displacement from reversal potentials E_j. G_{ij}^s is the maximum total conductivity of synapses to i from j, modulated by the synaptic activity variable s_i, which is governed by:

$$\dot{s}_i = a_r \phi(V_i; \beta, V_{th})(1 - s_i) - a_d s_i \tag{2.4}$$

where a_r and a_d correspond to the synaptic activity's rise and decay time, and ϕ is the sigmoid function $\phi(x) = 1/(1 + \exp(-x))$, set here with width β and center V_i^{th}. Solving Equation 2.4 for its equilibrium value at $\dot{s}_i = 0$ yields:

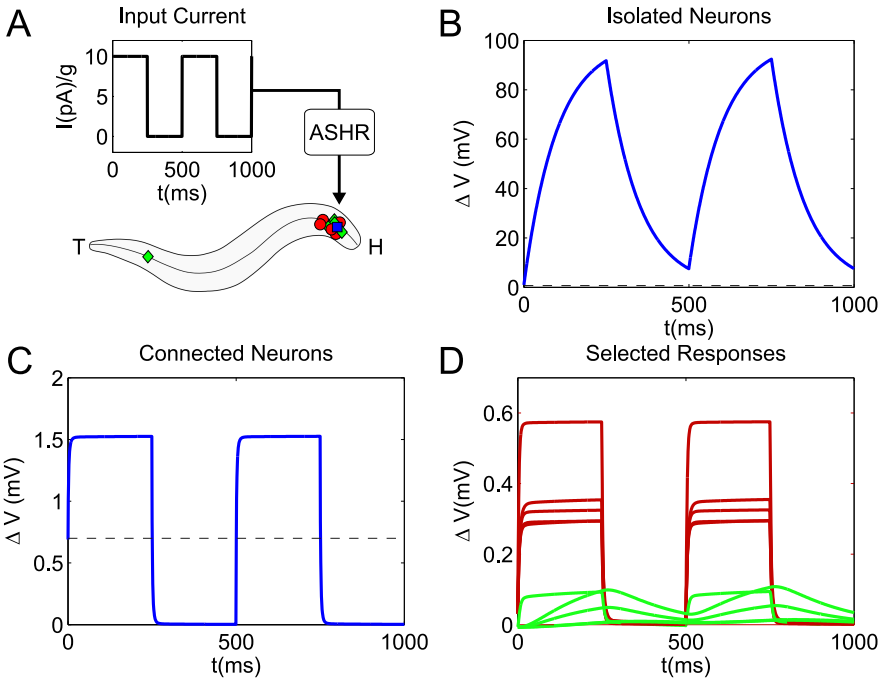

Figure 2.2: **A.** To demonstrate the behavior of the model, a square wave current input was injected into the neurons ASHR (units are pA divided by conductance constant $g = 100$ pS). Location of PLM neurons along the body of the worm are indicated by blue squares, and the location of 20 selected responding neurons are indicated by red circles and green diamonds. **B.** Resulting membrane voltage displacement from equilibrium in ASHR when isolated (i.e. no connections). **C.** Resulting membrane voltage displacement in ASHR when connected to the neuron network. Note the two order of magnitude reduction in peak voltage. **D.** Voltage responses of the 5 neurons most active in the presence of input current in red (dark gray) and the 5 neurons most active in the absence of input current in green (light gray). The ASHR neuron is chosen so as to illustrate a ubiquitous phenomena: downstream neurons do not necessarily get entrained to the time-response of the stimulated neuron.

$$\dot{s_i} = 0 \implies s_i = \frac{\phi(\beta(V_i - V_i^{th}) + \ln(1 + a_r/a_d))}{(1 + a_d/a_r)} \tag{2.5}$$

Thus the equilibrium value of s_i depends sigmoidally upon the membrane voltage V_i. Synaptic current I_i^{Syn} into neuron i depends upon the values of s_j for all presynaptic neurons j, and s_j depends directly upon V_j. Thus I_i^{Syn} depends directly upon the membrane voltages of presynaptic neurons.

2.1.2 Parameters

While the precise parameter values of each connection are unknown, we assume reasonable values as previously considered in the literature [122, 133]. We assume each individual gap junction and synapse has approximately the same conductance, roughly g =100 pS [122]. Each cell has a smaller membrane conductance (taken as 10 pS) and a membrane capacitance of about $C_i = 1$ pF [122]. Leakage potentials are all taken as $E_c = -35$ mV [133]. Reversal potentials E_j are 0 mV for excitatory synapses and -45 mV for inhibitory synapses [133]. For the synaptic variable, we choose $a_r = 1$, $a_d = 5$, and define the width of the sigmoid by $\beta = 0.125$ mV^{-1} [133]. V_{th} is found by imposing that the synaptic activation $\phi = 1/2$ at equilibrium [133]. Using this condition with Equations 2.1 and 2.4 allows us to solve for both the values of V_i^{th} and V_i^{eq}, which we call the "standard" equilibrium. Note that this equilibrium always exists, though is not necessarily the only fixed point within the system, nor is it always stable.

The directionality of the connections (i.e., inhibitory or excitatory) is estimated by the rough approximation that putative GABAergic neurons are inhibitory, while cholinergic and glutamatergic neurons are excitatory (as in [122]). This estimation of parameter values captures robust responses in the network dynamics and excludes from the simulation any responses which depend on more precise details of the network. The network that we simulate consists of 279 somatic neurons, where we exclude the 20 pharyngeal neurons and 3 additional neurons which make no synaptic connections, as in [122]. To validate the simulation and the

choice of parameters we tested for robustness by perturbing (±20%) individual connection strengths and each neuron's parameters, showing that dynamic functionality persists.

2.1.3 Model Discussion

The model does not include various extra-synaptic features known to drive or regulate responses. For example, there is evidence that self-sustained forward locomotion in *C. elegans* is regulated by proprioception within motor neurons [130] (compare how our model, lacking this, does not sustain oscillation in the absence of explicit external input). Computational modeling which includes stretch-receptive proprioception shows that such feedback loops can control behavioral features such as gait modulation between differing environments [10, 19]. The lack of such feedback mechanisms and other signaling mechanisms (such as various neuromodulators, monoamines and peptides [92, 123]), in combination with the simple neuron model and parameter assumptions, mean that specific responses to given inputs seen within the model can be encoded only within the network's connectivity. This reductive approach yields information as to how behavioral responses could be encoded within the structure of the connectome.

2.1.4 Simulated Steady-State Neuron Responses

In situ whole-cell voltage patch-clamp recordings have shown that the steady-state relationship between net membrane current and voltage is nonlinear [50]. Specifically, the steady-state I-V (current-voltage) curves of several *C. elegans* neurons were shown to have flat regions of high "phenomenological impedance." These experimentally obtained IV curves, from Goodman et al. [50], are reproduced in Panel D of Figure 2.3.

We investigated the ability of our model to produce similar nonlinearity in steady-state I-V relationships by performing simulated versions of these measurements. This was done by setting the membrane voltage of a neuron, then measuring that neuron's input current after allowing the system to equilibrate. This was repeated for several voltages to construct steady-state I-V curves.

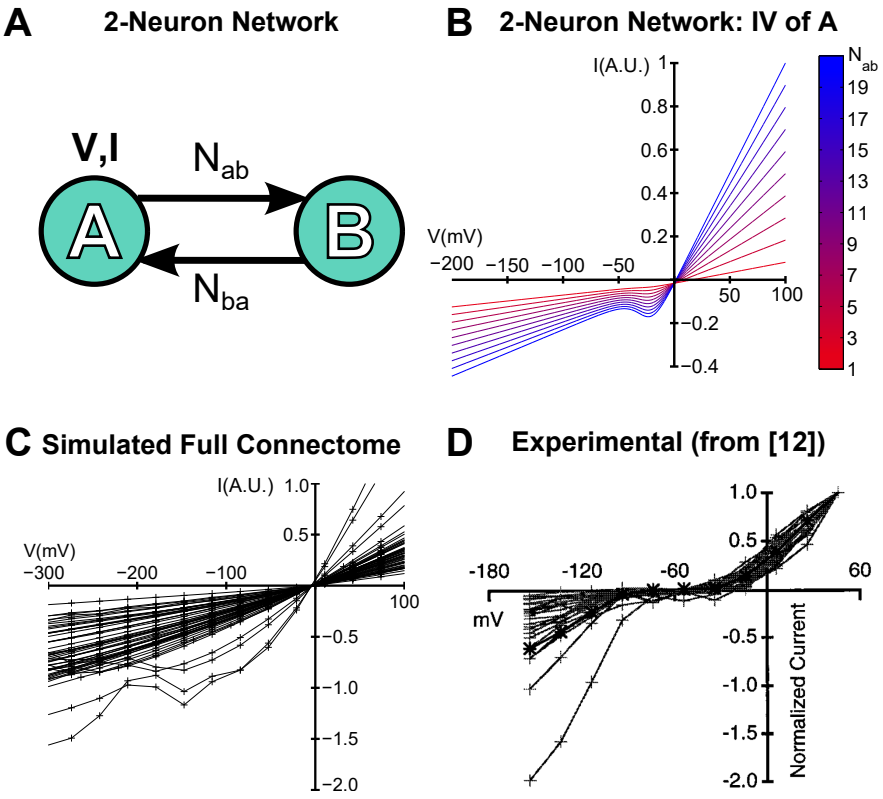

Figure 2.3: **A.** A 2-Neuron network, consisting of N_{ab} synaptic connections from A onto B and N_{ba} synapses from B onto A. **B.** The I-V curves for neuron A in this simple network. The colors represent different values of N_{ab}, with fixed $N_{ba} = 1$. **C.** A population of simulated I-V curves using the full connectome. Each curve corresponds to a different neuron. Note that the more elaborate connectivity and presence of gap junction connections can lead to more complicated nonlinearities. **D.** Experimental *C. elegans* I-V curves, reproduced from Goodman et al. [50]. Obtained by in situ whole-cell patch-clamp recordings for various (mostly unidentified) neurons.

We first considered a network of two neurons, named A and B, as illustrated in Panel A of Figure 2.3. The parameter values of previous subsection were used, and the connectivity consisted solely of N_{ab} synapses from A onto B and N_{ba} synapses from B onto A (with no gap junction connections). The resulting IV curves for various values of N_{ab} is shown in Panel B of Figure 2.3 (with fixed $N_{ba} = 1$). This shows that the connectivity of a neuron affects strongly the shape of its steady-state I-V curve, and that the model is capable of producing nonlinear I-V curves which include relatively flat central regions.

A more direct comparison with the experimentally-obtained curves comes from the same simulation performed with neurons in the full connectome. A population of I-V curves from neurons in the full network is shown in Panel C of Figure 2.3. We observe from this qualitative comparison (subject to different normalizations and uncertainty of which neurons were measured experimentally) that our neuron model, within a connected network, is indeed capable of producing non-trivial nonlinearities in neural steady-state I-V relationships.

2.2 Analysis of Simulated Dynamics

There are many ways to test the validity of the *C. elegans* model. Given the numerous stimuli response experiments [65, 102, 130], we can simply select a neuron of interest and interrogate the downstream neuronal response. For instance, the PLM neurons (PLML/R) are posterior touch mechanoreceptors. Activation of PLM by tail-touch causes a worm to move forward or, if already moving forward, to accelerate [25]. Thus stimulating these neurons should produce a downstream time-dependent neural-response resulting in a motorneuron response consistent with forward motion. Figure 2.1 illustrates a schematic for this neuro-sensory cascade from sensory activation by stimulation of the sensory neuron PLM that excites the motor-neurons associated with forward motion [102]. Characterizing such neural pathways are the key objective in this study.

A characteristic example of simulated full-network dynamics can be seen in Figure 2.2, in which the polymodal nociceptive neuron ASHR[3] is stimulated with a constant periodic input current. The system starts out in equilibrium before the excitation, and the voltages

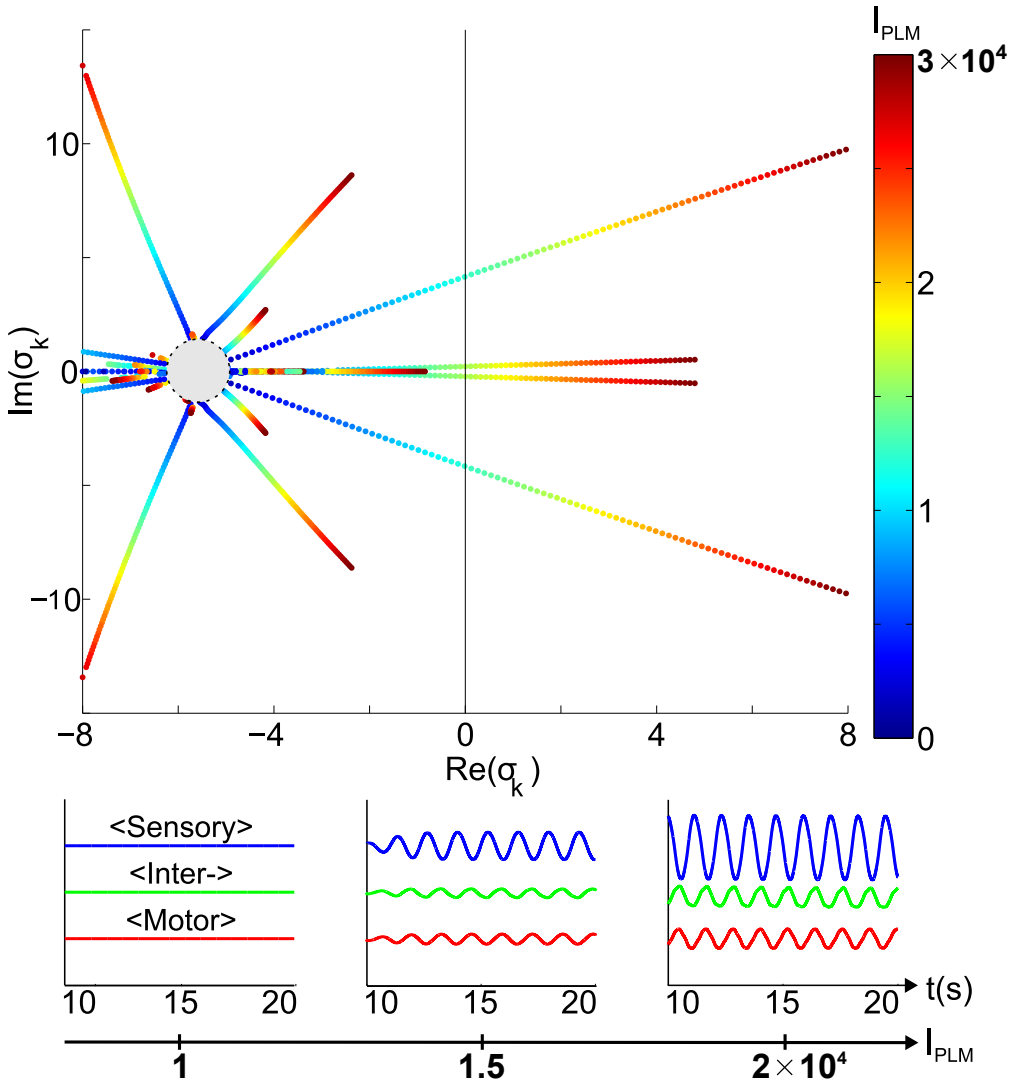

Figure 2.4: Jacobian eigenvalue spectrum as a function of PLM input amplitude. Input units are current normalized by conductance constant g. At inputs around 1×10^4, the system goes through a Hopf bifurcation and oscillatory motion results. Traces of average sensory, inter- and motor neuron voltage displacements from equilibrium are shown to illustrate this.

plotted are neuronal displacements from their equilibrium values. Panel B shows the response an unconnected neuron, whereas panel C shows the response of ASHR when connected to the network. Note that the characteristic timescale of the response changes due to the presence of connections. In panel D, the voltage responses are plotted for the 5 neurons which respond most strongly when input current is present, and for the 5 which respond most strongly when it is not. This illustrates that downstream neuron responses are not necessarily entrained to the stimulus, but may respond through different temporal modes.

2.2.1 Low-Dimensional Bifurcations

Behaviorally, crawling is known to be dominated by a two-mode stroke motion [110], i.e. the so-called eigenworm motion. Thus the motor-neuron response to PLM stimulation should produce a two-mode dominance in accordance with the eigenworm behavior given that the motor responses control muscle contraction [102]. We therefore intuitively anticipate that a constant input of sufficient strength, corresponding to sensory stimulus, should be able to drive two-mode oscillatory behavior in the forward motion motorneurons. To test if this is qualitatively captured by our model, we first seek oscillatory solutions by calculating the Jacobian matrix at equilibrium and looking for eigenvalues with positive real parts.

With zero external input, all Jacobian eigenvalues have a negative real part and the system is stable. However, eigenvalues with a positive real part are seen to exist for sufficiently high constant input amplitudes. Figure 2.4 shows the Jacobian spectrum as a function of PLM input amplitude. At certain threshold values, the system goes through Hopf bifurcations and oscillatory modes arise. The average voltage displacement within each neuron class is shown on the right of the figure, illustrating this.

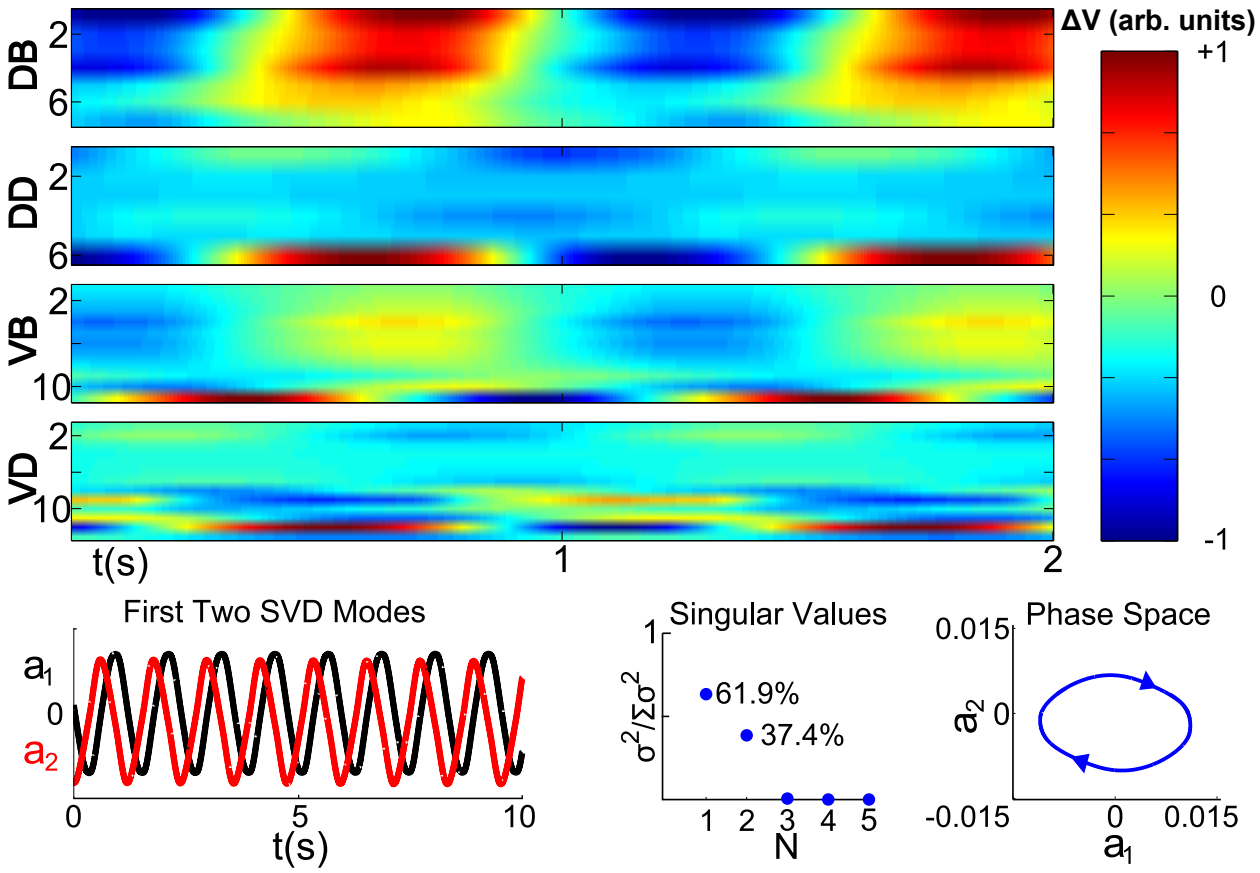

Figure 2.5: Resultant dynamics for a constant input 2×10^4 into PLM (with this input, the system is within an oscillatory regime). Raster plots of the voltage responses within the forward motorneurons are shown at top. The time evolution of the SVD modes along with their singular values are shown on below, establishing that the system response is indeed dominated by two modes. The trajectory of these modes in phase space correspond to a two-mode swimmer.

2.2.2 Singular Value Decomposition of PLM response

To obtain the modes that the motor neurons exhibit we collected time snapshots of motor-neuron voltages $\mathbf{V}_M(t)$ into a matrix and computed the singular value decomposition:

$$V = \begin{bmatrix} \mathbf{V}_M(t_0) & \mathbf{V}_M(t_1) & \dots \end{bmatrix} = \mathrm{P} \cdot \Sigma \cdot \mathrm{Q}^T \tag{2.6}$$

The columns of matrix P (the vectors $\mathbf{P_i}$) are the principal orthogonal components, which are weighted by the diagonal elements in Σ (the singular values σ_i). Decomposition of the voltage onto these provides the dynamical coefficients $a_k(t)$:

$$\mathbf{V}(t) = \sum_{k=1}^{N} a_k(t)\mathbf{P}_k$$

Figure 2.5 shows the time-dynamics of the motor neurons given constant PLM stimulation consistent with tail-touch at input amplitudes above the Hopf bifurcation level. As shown, there are two dominant response modes that produce periodic, laterally out-of-phase, voltage activity.

The analysis as shown in the bottom row of Figure 2.5 confirms that the motor activity is dominated by two time-dependent response modes (with the first and second modes possessing 61.86% and 37.36% of the energy respectively). Their dynamics are periodic and similar to physiological [110] and behavioral [65] studies that find low-number of modes that determine the motion (specifically, there are two dominant oscillatory modes which move through their phase space in a ring around the origin). Thus the model produces a proxy for this behavior through analyzing motor responses, although it does not produce *directly* the behavioral response.

2.3 Ablation

Experimental ablation studies [25] have observed that the ablation of the densely-connected AVB interneurons destroys the worm's ability to perform forward motion, whereas ablation of the similarly densely-connected AVA interneurons preserves it (affecting instead the ability

of the worm to perform backwards motion). If our model's PLM response modes do indeed serve as a proxy for this behavioral response, they should be similarly affected by such network modifications.

We explore the effect of ablation upon our response modes by removing the AVA/AVB interneurons from the network and repeating the analysis of Figure 2.5. Specifically, the neurons AVAL/AVAR (or AVBL/AVBR) were removed from the network, the dynamics in response to an identical constant PLM input were simulated, and the SVD was calculated. The resultant singular value distributions for these ablations is shown in Figure 2.6, which shows that the two-mode dominance is destroyed with the removal of AVB but remains intact with the removal of AVA. This serves as another confirmation that the response modes correspond to the experimental foward-motion modes.

2.3.1 Change in Response Modes

The top rows of Figure 2.6 show the dynamics of the first two SVD modes (i.e. the time evolution of $a_k(t) \cdot \mathbf{P_k}$ for modes $k = 1, 2$) before any ablations (for the "healthy" system) and after ablation of AVA, AVB, and AIZR (the latter being chosen because, experimentally, the ablation of AIZR does not inhibit forward motion[3]). The 37 neurons selected are the forward-motion motorneurons (those belonging to classes DB, DD, VB and VD, as in Figure 2.5). A time interval of 1 second was selected from each simulation such that the first modes of all cases were maximally in-phase. Note that the same observation as before can be made when qualitatively comparing the structures of the modes of the healthy, AVA ablated and AVB ablated case: when the AVA interneurons are ablated, the structure of the modes appears slightly altered, but similar, whereas ablation of AVB destroys the dominant mode.

2.3.2 Quantification of Response Similarity

To quantify the effect of ablations on the response modes and their dynamics we introduce two metrics. The first metric measures the similarity of the singular values by computing

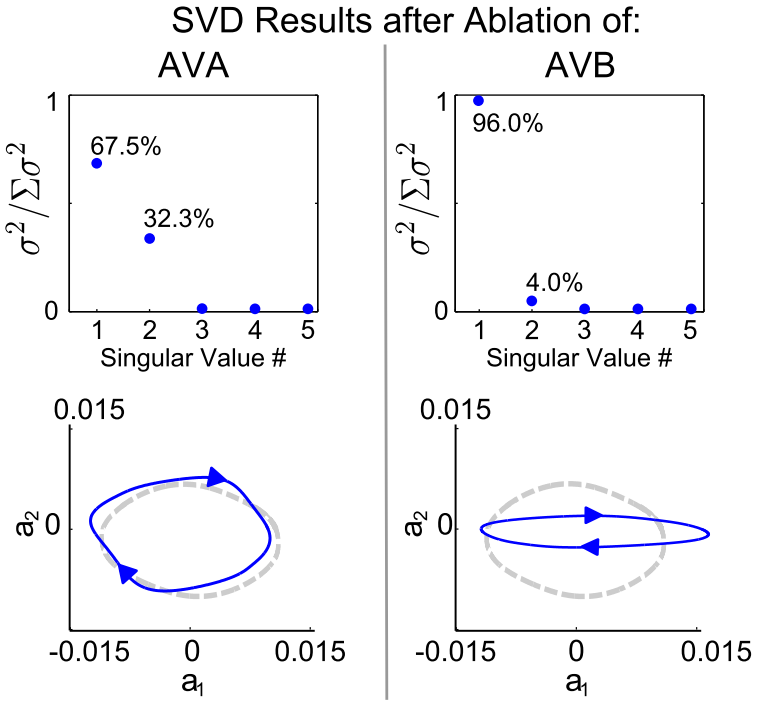

Figure 2.6: Singular value distributions when the analysis of Figure 2.5 is repeated with the removal of the AVA or AVB interneurons from the network. Experimental studies [25] show that the ability to perform forward motion is destroyed with the removal of AVB, but preserved under the ablation of AVA. The second row shows the new trajectories in phase space after these ablations (where the dashed gray line is the healthy trajectory, for comparison). Note that AVA ablation does distort the trajectory, but does so less than does the ablation of AVB. Within our model, the two-mode response in the forward-motion motorneurons is affected in the same way by these ablations (AVB ablation destroys two-mode dominance, AVA ablation preserves it). This further suggests that the simulated neural modes serve as proxies for forward motion.

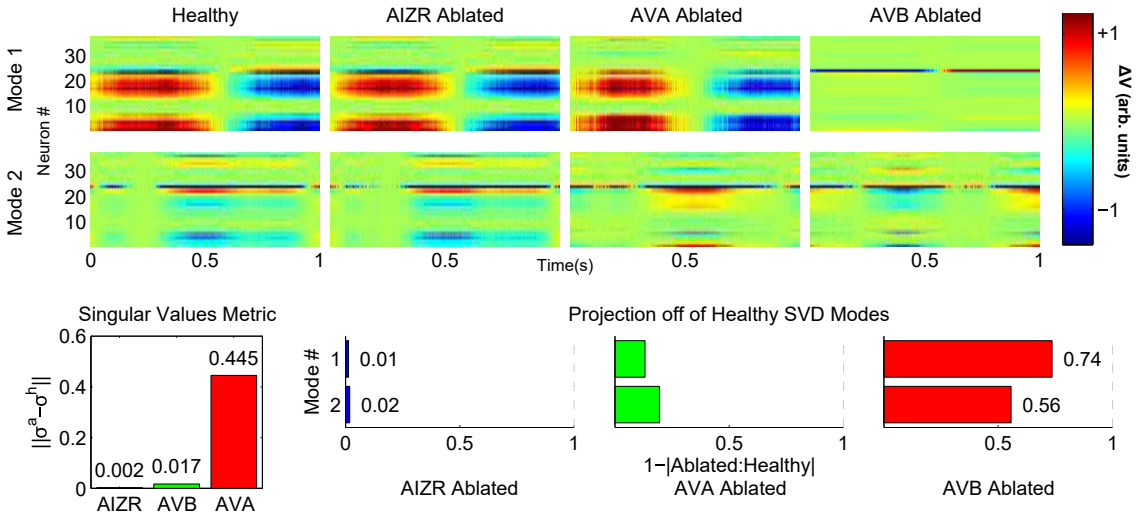

Figure 2.7: Modes after ablation of neurons compared to healthy modes. The raster plots at top show the first and second SVD modes of forward-motion motorneuron activity for the healthy case (full network, no ablations), along with the cases where AIZR, AVA and AVB are ablated. Phase-matched one-second intervals of each mode are shown. On the bottom row the metrics of Equations 2.7 and 2.8 are shown for the ablation modes compared to the healthy modes.

the l^2-norm between the ablated and healthy distributions. For ablated singular values σ^a and healthy singular values σ^h, we compute:

$$||\sigma^a - \sigma^h|| = \sqrt{\sum_{i=0}^{N} (\sigma_i^a - \sigma_i^h)} \tag{2.7}$$

The second metric computes the similarity between the mode dynamics. We take the one-second dynamics segments from Figure 2.6 (labeled here as matrix H for the healthy modes and A for the ablated), and we compute the absolute value of their Frobenius product:

$$|\mathrm{A} : \mathrm{H}| = \left| \sum_{i,j} \mathrm{A}_{ij} \cdot \mathrm{H}_{ij} \right| \tag{2.8}$$

where all matrices have been normalized to have a Frobenius norm of one.

Figure 2.6, on the bottom row, shows the values of these metrics for the ablations of AVA, AVB and AIZR. By these metrics, ablation of AIZR does not affect the network's

response to this stimulus, ablation of AVA affects only slightly the network's response, and ablation of AVB destroys the functionality of the network in response to a PLM stimulus. This suggests that such analysis can be used to computationally classify the roles played by specific neurons in the response of the network to given stimuli. Hence our model provides a computational framework in which to computationally classify, from no prior knowledge, the neural subnetworks responsible for behavioral responses to stimuli.

2.4 Discussion

In conclusion, we have developed a neuro-sensory integration model of the *C. elegans* nematode which describes the nonlinear, time-dependent, network voltage conductances. In our computational model, the entire 302 neuron network of sensory-, inter- and motor-neurons are dynamically coupled with the available biophysical connectome data. In the specific application of the tail mechanosensory neuron PLM stimulation, a complete neuro-sensory integration of this specific stimulus pathway is discovered whereby sensory information translates to downstream motor responses that are responsible for behavioral actions, in this case a two-mode swimmer dynamics. In theoretical terms, the input stimulus robustly induces a Hopf bifurcation in the network. Thus a low-dimensional bifurcation, which is ultimately responsible for behavior, is inscribed in the underlying network structure.

With the abundant current and on-going biophysical experiments on individual neuron stimulation in *C. elegans* (through opto-genetics, for instance), the current model presents a significant step forward in providing a theoretical platform to more accurately understand neuro-sensory encoding, processing, and integration. Specifically, we construct a biophysically inspired computational model and demonstrate that the underlying low-dimensional bifurcations of the network drive neural voltage modes which are responsible for low-dimensional movement and behavior. These neural modes can be linked to behavioral responses via comparison with the experimentally observed behavioral effects of neural network modification. Thus our model allows for the identification and characterization of behavioral responses which are encoded within low-dimensional bifurcation structures in the network. The iden-

tification of such bifurcation-encoded responses within the network allows for computational classification of neurons into the subnetworks responsible for those responses.

Our study thus allows one to study the structure and robustness of networks of voltage conductances for producing prescribed responses. More broadly, understanding the *C. elegans* model organism may help produce and promote bio-inspired network designs in other fields of scientific applications given the observed robust nature of such architectures. This study promotes a viewpoint of the broader potential for understanding what can be gained in modeling physical systems whose dynamics are driven by network connectivity and nonlinear dynamical systems. This can lead to bio-inspired design, quantification and engineering principles capable of producing robust functionality.

Chapter 3

EXISTENCE AND FUNCTION OF MULTISTABILITY AND TRANSIENT TRAJECTORIES

The neural dynamics of the nematode *C. elegans* are experimentally low-dimensional and correspond to discrete behavioral states, where previous modeling work has found neural proxies for such states. Experimental results further suggest that dynamics may be understood as long-timescale transitions between multiple low-dimensional attractors. To identify multistable regimes of our model, we develop a method for systematic generation of bifurcation diagrams and their analysis in an interpretable low-dimensional subspace, showing the existence and nature of multistable input responses at a glance. Stimulation of the PLM neuron pair, experimentally associated with forward movement and shown in simulation to drive a limit cycle, defines our low-dimensional projection space. We then obtain bifurcation diagrams for single-neuron excitation over a range of amplitudes and classify whether the dynamics in this projection space are associated with a limit cycle, fixed point, or multiple states. In the specific case of compound input into both the PLM pair and ASK pair we discover bistability of a limit cycle and a fixed point, and the transient timescales in approaching each of these states are much longer than other timescales of the system. This suggests consistency of our model with the characterization of dynamics in neural systems as long-timescale transitions between discrete, low-dimensional attractors corresponding to behavioral states. Our methodology thus prescribes a method for identifying these states and possible transitions in response to arbitrary input.

3.1 Motivation in Searching for Multistability

While the exact role of the connectome in neuronal computation remains unresolved and in general controversial, we have shown that simple models of *C. elegans* neural dynamics (combining specific connectivity data with simple unfit parameter estimates and dynamics) are capable of generating non-trivial, qualitatively correct responses to given stimuli [68] (see the results of Chapter 2). This suggests that such computational modeling can be informative in understanding how the system generates behavioral responses. It is therefore of interest to consider whether or not models capable of producing neural proxies for behavioral responses (as in Chapter 2) are further capable of characterizing experimentally observed multistable attractor dynamics.

One motive in the search for multistability is that we have found a neural proxy for behavior consisting simply of a single limit cycle within the system. However, the traditional dynamical systems view of neural phenomena in which behavior is described by dynamics on a stable attractor is limited, and experimental evidence suggests that many neural responses are better described as transient trajectories between multiple attractors [94]. Within this paradigm it is important not only that multistability should exist, but that transients with biophysically relevant behavioral timescales should exist.

In this chapter, we explore the input space of our full model for the neuronal network dynamics of *C. elegans* and find that various multistabilities arise in response to inputs. Performing direct neuronal simulations to reveal such multistable regimes is a formidable task, since the input space is large and neuronal simulations produce high-dimensional outputs which are difficult to interpret. We therefore develop a systematic methodology to explore responses to complex inputs and understand the dynamics within a framework of low-dimensional attractor dynamics. Specifically, for a chosen input vector we generate a bifurcation diagram (using the amplitude of the constant-in-time input as our bifurcation parameter) to identify multistability. We use an interpretable low-dimensional projection (as defined by forward motion) to characterize the dynamics for multiple attractors as identified

by the bifurcation diagrams.

With this framework we survey all input vectors corresponding to single-neuron current injections and find bifurcations corresponding to limit cycles and multiple attractors. For some of these input vectors which induce multistability, simulated transient dynamics are on much slower (on the order of seconds to tens of seconds) than any intrinsic neuronal timescales (which in our model do not exceed a few hundred milliseconds). The transient trajectories themselves are low-dimensional and could be associated with network-produced functionalities, such as neural proxies for movement. Thus our connectomic simulations are in agreement with behavioral observations of *C. elegans* and help support recent biophysical conjectures that the transients themselves are critical in understanding behavioral assays [94].

As a particular example, we choose input into the PLM neuron pair, which is known experimentally to excite forward motion [25] and within our model creates a two-dimensional limit cycle response [68]. We then use the low-dimensional PLM response plane to consider the dynamics of a compound input vector PLM+ASK, where ASK stimulation is known to facilitate transitions (i.e. turns) [53]. Our bifurcation analysis reveals that this induces bi-stability, in which the system goes either into a fixed point or a limit cycle. Transient timescales are shown to be considerably longer in this bistable case than the intrinsic timescales of the system. This allows for long timescales in the system in the presence of discrete, low-dimensional attractors corresponding to behavioral states, consistent with the experimentally-based framework [109]. This input scenario demonstrates how our bifurcation analysis methodology prescribes a generic approach for identifying multi-stable states and their transient timescales in response to arbitrary inputs. Since we model neurons as identical save for their connectivity, it further indicates that their connectivity alone can encode the creation and destruction of multiple behavioral attractors.

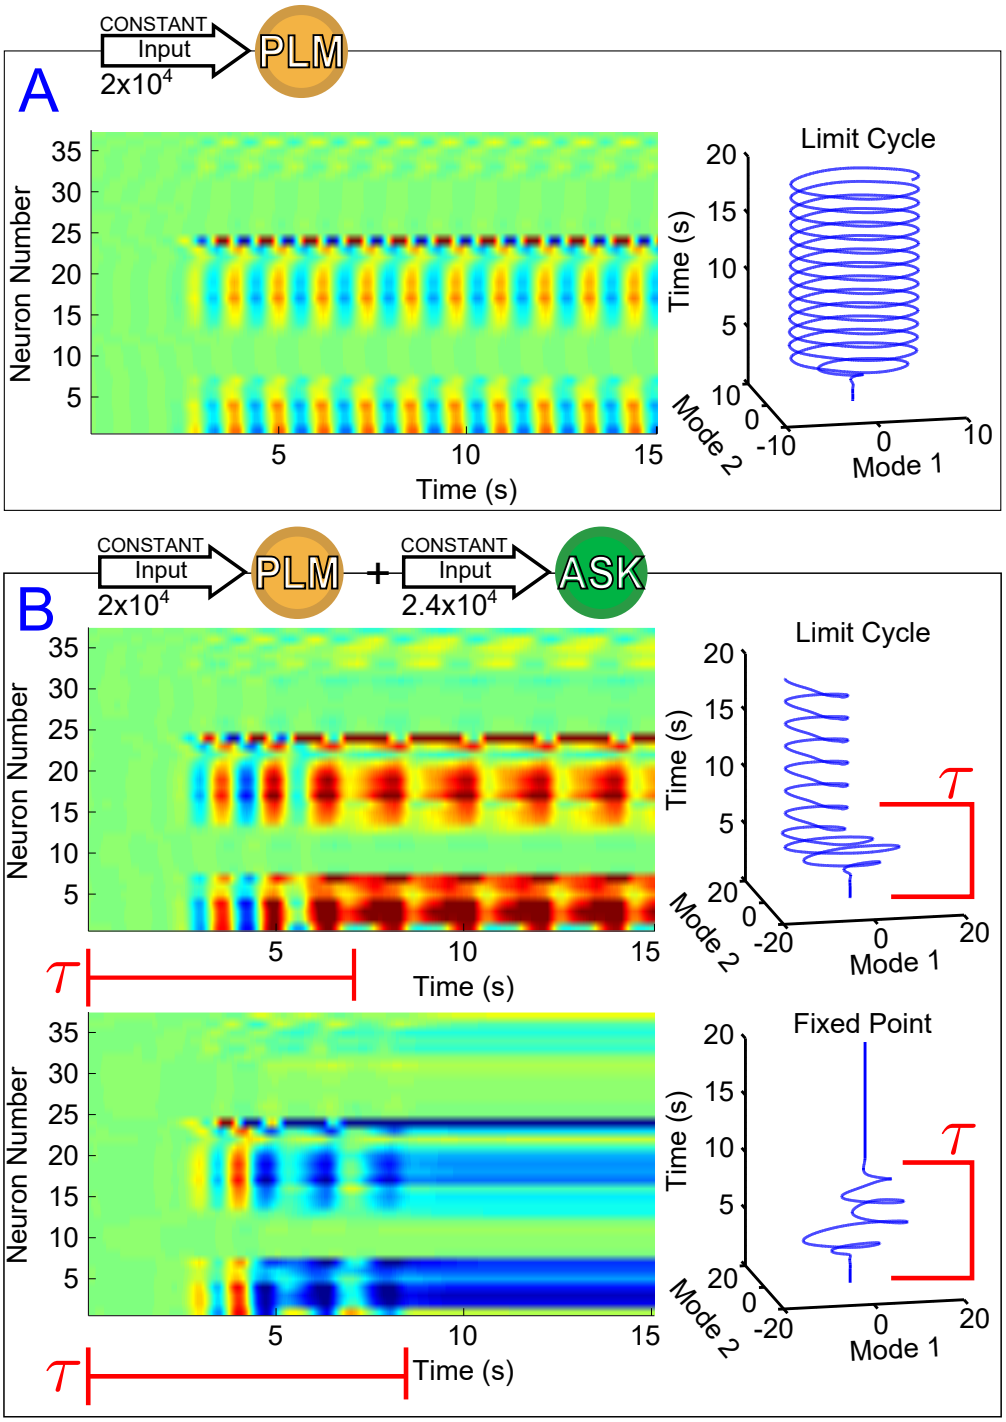

Figure 3.1: **A.** Voltage dynamics of forward-motion motorneurons (neurons of classes DB, DD, VB and VD) in response to the following sensory inputs: in Panel A, an input of 2×10^4 (Arb. Units) into the PLM sensory neuron pair (known experimentally to drive forward motion [25]). **B.** An input of 2×10^4 into the PLM pair with an additional input of 2.4×10^4 into the ASK sensory neuron pair (known experimentally to promote turning [53]). Simultaneous PLM+ASK stimulation causes bistability, with relatively long transient times τ. To the right of each raster plot is the trajectory within the Forward-Motion 2D Plane (defined by the trajectory in Panel A, and used for all subsequent projections).

Interaction	Timescale
Single-Neuron Membrane Leakage	100 ms
Gap Junctions	10 ms
Synaptic Connections	200 ms

Table 3.1: Orders of magnitude for various timescales within the system for the parameters chosen.

3.2 Simulation of C. elegans *Connectome Dynamics*

3.2.1 Model Timescales

Of particular relevance to this Chapter are the timescales within the system. From the first term in Equation 2.1, we see that the exponential free decay constant of an unconnected neuron (i.e. decay through the membrane leakage term alone, with $I_i^{Gap} = I_i^{Syn} = I_i^{ext} = 0$) would be $\tau_{free} = C/G^c = 100$ ms. Similarly, the time constant value given by gap junctions would be $\tau_{gap} = C/g = 10$ ms.

There are also timescales intrinsic to the synaptic dynamics. We approximate these by considering the dynamics when voltages are held constant, and thus $\phi(\beta(V_i - V_i^{th})) \equiv \phi_i$ is constant. Then Equation (2.4) becomes:

$$\dot{s_i} = a_r\phi_i - (a_r\phi_i + a_d)s_i \tag{3.1}$$

and thus the synapses will exponentially approach equilibrium with a time constant of $\tau_{syn} = 1/(a_r\phi_i + a_d)$. Since $a_r = 1$ s^{-1}, $a_d = 5$ s^{-1}, and $\phi_i \in (0, 1)$, synapses must have exponential time constants in the range $\tau_{syn} \in (166, 200)$ ms.

This collection of timescales is summarized in Table 3.1. It will be shown that, when the system is in a bistable regime, the timescales of transient dynamics within the system can be orders of magnitude above any of these intrinsic time constants within the system (on the order of 10 s, for example).

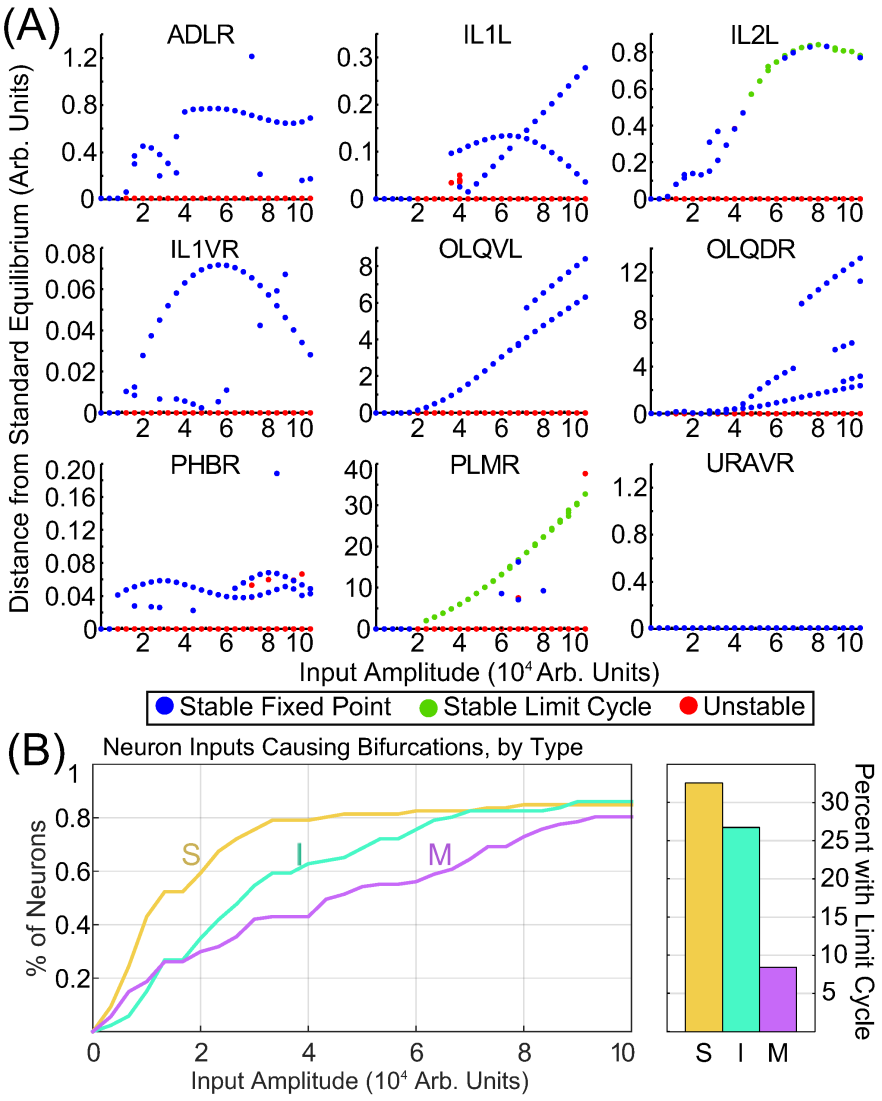

Figure 3.2: **A.** A selection of bifurcation diagrams for inputs into different neurons, showcasing the variety of multistable behaviors within the system. **B.** Input amplitude of the first bifurcation, by neuron type. The vertical axis shows the percentage sensory neurons, interneurons, and motorneurons for which the standard equilibrium is unstable at the corresponding input amplitude. On average, sensory neurons drive bifurcations at a lower input amplitudes than inter- or motorneurons. Motorneurons are much less likely to drive limit cycles within the system.

3.2.2 Response to PLM Stimulation Defines the Low-Dimensional Projection

As an example of the model's ability to generate proxies for behavioral responses encoded within the network, it was shown that stimulating the tail-touch mechanosensory neuron pair PLM, which experimentally leads to forward motion [25], gives rise via a bifurcation to a limit cycle within the forward-motion motorneurons. This limit cycle consists of only two modes, in agreement with the behavioral observation that the worm's body shape during forward motion is well-described by a similar two-mode oscillation [110]. The non-triviality of this agreement was established by showing that simulated ablation studies affected this response in agreement with experimental ablation studies (e.g., ablation of the AVB interneurons destroys the response both experimentally and in the model [25]).

Figure 3.1 shows the response of forward motion motorneurons to various inputs as a function of time. Panel A of the figure shows a raster plot of motorneuron voltages in response to PLM input (through the I^{Ext} term in Equation (2.1)), for which the two-mode oscillatory response can be observed [68]. The trajectory of these two leading modes are plotted as a function of time on the right. We use this same low-dimensional space (defined as the two forward-motion motorneuron modes which oscillate during PLM activation) throughout the paper. In other words, we use the same projection for the low-dimensional trajectories in Panel B and in all further figures.

Specifically, we calculate this plane (as in [68]) by taking time snapshots of forward-motion motorneuron membrane voltages $\vec{V}_M(t)$, collecting them into a matrix V (the matrix plotted in the raster plot of Figure 3.1(A)), and taking that matrix's singular value decomposition. That is:

$$V = [\vec{V}_M(t_0), \vec{V}_M(t_1)...] = P \cdot \Sigma \cdot Q^T, \tag{3.2}$$

where P and Q are unitary and Σ is diagonal. The columns of P are the principal orthogonal modes. As in [68], the first two of these modes (the first two columns of P) almost entirely capture the dynamics of the system within this subspace under constant PLM stimulation. Projection of the full-system dynamics onto this plane consists of projecting the system's

motorneuron dynamics onto these modes.

Importantly, all of the neuron membrane voltages which have a nonzero projection onto this plane are forward-motion motorneurons (that is, motorneurons of class DB, DD, VB and VD). All neurons outside of these classes (and thus all sensory and interneurons) belong to the null space of this projection; that is, only motorneuron dynamics project onto this plane. All low-dimensional trajectories within this study can therefore be understood as corresponding to motorneuron dynamics, aiding in their potential biological interpretability (in the sense that motorneuron dynamics should map onto muscular dynamics).

This specific two-dimensional plane represents the dominant motorneuron neural response modes. These modes correspond to muscle deformation through additional biophysical processes. Each stable fixed point within this plane corresponds to a static neural response, whereas a limit cycle corresponds to an oscillatory trajectory. Therefore, oscillatory trajectories which we identify in this space indicate possible neural proxies for low-dimensional body movements in *C. elegans*.

3.2.3 Response to Compound Inputs: Multistability

The response to PLM stimulation alone consists of a single possible state (i.e. a limit cycle trajectory), but if the model is capable of describing the dynamics in terms of long-timescale transitions between states under the same input, then we wish to find inputs which allow multiple states and transitional dynamics. We find that such inputs indeed exist. As an example, we consider the response to simultaneous stimulation of the PLM neuron pair along with the ASK neuron pair. We choose this stimulation since excitation of ASK neurons have been shown experimentally to promote turning [53] and their ablation greatly increases the duration of periods of forward motion [125].

As we show in Figure 3.1B, for this combined input there coexist two different attractors, i.e. the system is bistable. The two trajectories plotted are in response to the same constant input amplitudes into PLM and ASK, and differ only by their initial conditions. Note that the transients before convergence into the eventual fixed point or limit cycle have long

timescales (relative to the intrinsic timescales of the system as discussed in Section 3.2.1). The model therefore does exhibit multistability for this given input, but given the large dimensionality of the input space the discovery, identification and interpretation of these multistable regimes is not trivial. Since we wish to understand the neural dynamics as consisting of long-timescale transients in the presence of multiple discrete attractors, we develop a method for (1) identifying the existence and nature of attractors in response to arbitrary inputs, (2) characterizing transient timescales, and (3) providing interpretable biophysical meaning to calculated trajectories via projection onto a meaningful low-dimensional space.

3.3 Bifurcation Diagrams for State Identifications

3.3.1 Calculating Bifurcation Diagrams

Motivated by observational studies which describe *C. elegans* behavioral dynamics in terms of low-dimensional attractor dynamics [109], we wish to understand our simulated neural dynamics in the presence of multiple discrete attractors. We therefore propose to construct bifurcation diagrams that depict attractors existing within the system under arbitrary inputs. By fixing the direction of the input vector I^{Ext} in Eq. 2.1 and using its amplitude as our bifurcation parameter, such diagrams will show us at a glance the set of states created in response to a given input, and provide us with a method of identifying induced multistability.

Figures 3.4 and 3.5 show examples of such bifurcation diagrams, in which we plot the furthest L^2 distance from standard equilibrium (within the 2D Forward-Motion Plane) of all attractors present as a function of input amplitude. In principle, such diagrams could be calculated by simply performing a large ensemble of simulations and projecting the results into the 2D plane, but such simulations are relatively time-consuming. We can take advantage of the fact that, for this model, it is easy to compute its Jacobian matrix at any given point for any constant input. We therefore use Newton's method when possible, supplementing with simulations to explore the input space and find additional attractors or when Newton's method does not converge to the desired result. Full detail on the algorithm used to generate

such diagrams can be found in Appendix 3.4.1.

We generated these diagrams for all 279 of the single-neuron inputs into the system. These figures can quickly give an indication of the likely number of states for each range of inputs. Thus these diagrams give a means of identifying what attractors will exist within the system for a broad range of arbitrary inputs, and of easily identifying regions of multistability in the input space.

Generating these diagrams for all possible single inputs allows for the qualitative comparison of features within each neuron's bifurcation diagram. Similar features in the bifurcation diagrams of neurons may suggest similar functionalities. As a simple example, in Figure 3.2, we compare the input amplitude at which the standard equilibrium first becomes unstable for sensory neurons, interneurons and motorneurons. The majority of sensory neurons are seen to drive bifurcations in the system at lower input levels than for most interneurons, which in turn require lower inputs than most motorneurons. Intuitively, this suggests that the system is typically more sensitive to input into sensory neurons than it is to interneuron or motorneuron inputs. Furthermore, for each group of neurons we compute the percentage of single neuron inputs which promote limit cycle attractors. We find that within our input range, 32.6% of sensory neurons and 26.7% of interneurons give rise to oscillatory dynamics, whereas only 8.4% of motorneurons result in oscillation when stimulated. This points to the sensitivity and particular ability of sensory neurons to drive complex dynamics within the network. Such results serve as a demonstration of the ability of these bifurcation diagrams to provide meaningful and intuitive information about the functionality of neurons and the behavior of the system.

3.4 Algorithm for Calculating Bifurcation Diagrams

3.4.1 Bifurcation Diagram Generation Algorithm

Panel A of Figure 3.3 gives a flowchart which describes the specific implementation of the algorithm used to automatically generate the bifurcation diagrams as pictured in Figure 3.2.

This algorithm was performed for each input level (that is, each horizontal position) of each diagram. This procedure, as summarized in Figure 3.3(A), is as follows:

The loop is initialized with a set of "guess" points $X_0 = [\vec{x_{01}}, \vec{x_{02}}, ...]$, with the goal of, ideally, finding stable solution(s) near each. Fixed points found at this input level are added to X_i. If this is the first set of calculations (i.e. the first input level), X_0 contains only the standard equilibrium point. If previous input levels have been calculated, the set of points found for the previous input level are used. The standard equilibrium for the new input level is added to this set if it is missing.

The members of X_0 generate the initial guesses for Newtons method; specifically, Newtons method is initialized from a point displaced along the least stable direction from the point in X_0 (i.e. the direction corresponding to the Jacobian eigenvalue with the maximum real component). In practice, this is seen to lead to a much faster convergence of Newtons method, and compared to a random perturbation is more successful in leading Newtons method to new fixed points (rather than just collapsing onto the already-found fixed point).

If Newton's method is not successful (either failing to converge within the maximum number of cycles, or converging to a point already in X_i) then a simulation is performed starting at a random point in the full-dimensional space. If Newtons method is successful but converges to a previously-found unstable point, a simulation is performed from the neighborhood of this unstable point. If it converges to a new, different unstable point, the algorithm records this new point and tries Newtons method again. If Newtons method is successful at finding a stable fixed point, it is recorded in X_i and the loop ends.

Simulations have ongoing convergence and periodicity checks (as detailed in Section 3.4.1), along with a maximum number of timesteps. If the maximum number of timesteps is exceeded without finding a stable solution, then Newtons method is tried again from last simulation point (which may be, ideally, within the basin of attraction of a new stable solution). If a stable fixed point or limit cycle is detected in a simulation, it is recorded in X_i, and the loop ends.

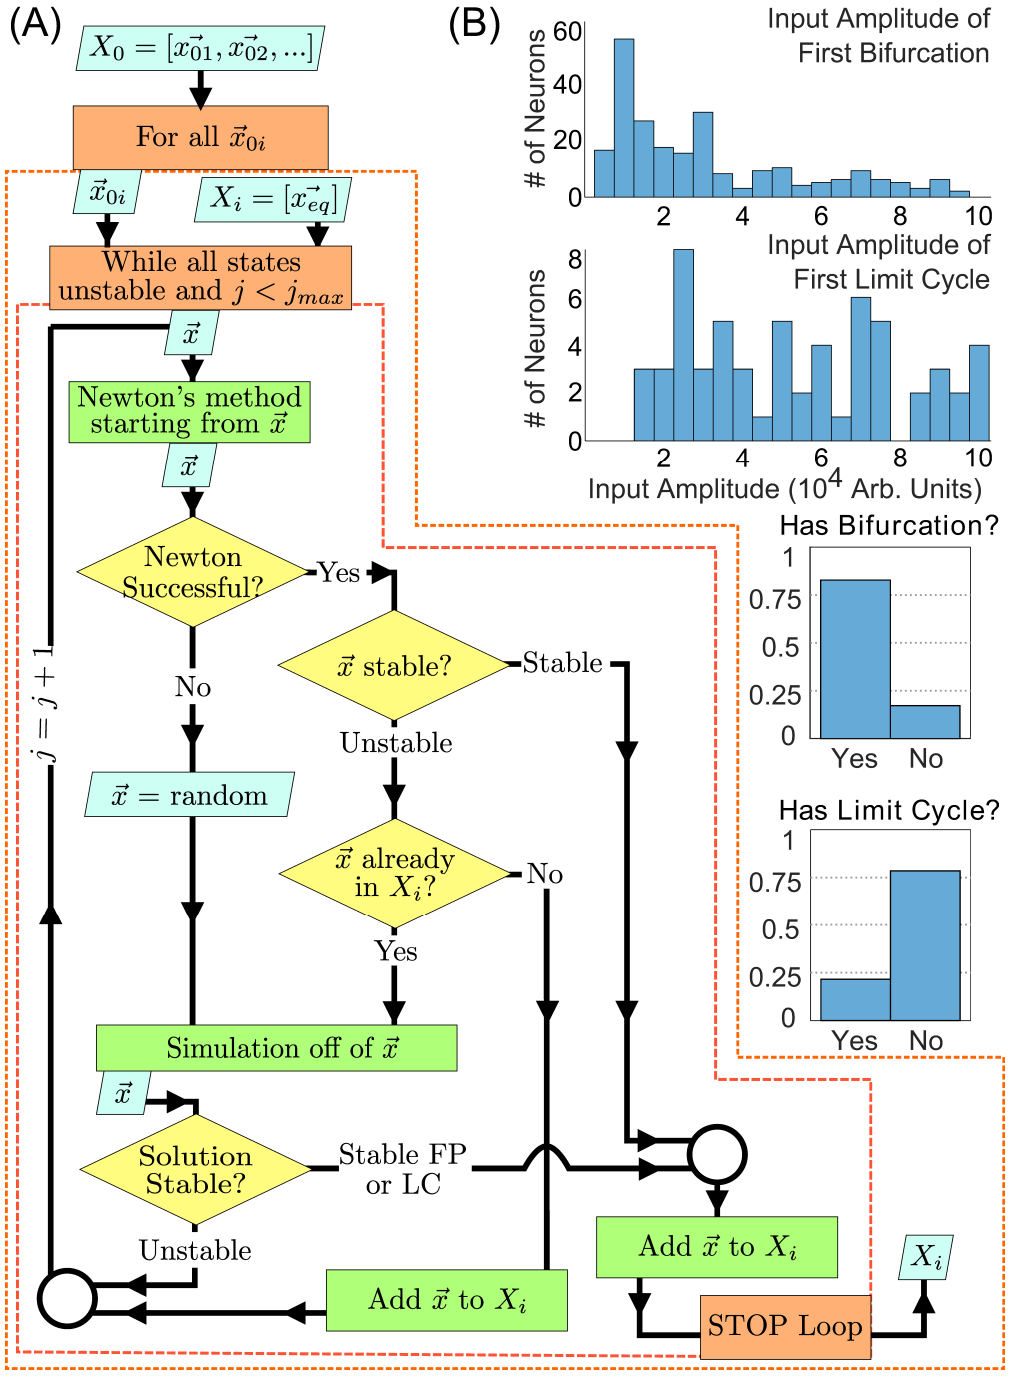

Figure 3.3: **A.** Outline of the algorithm used at each input level of each diagram to automatically generate bifurcation diagrams such as the ones in Figure 3.2. We discuss this algorithm in Section 3.4.1. **B.** We summarize some features from the diagrams generated for inputs into each of the 279 individual neurons.

3.4.2 Comments on Algorithm and Interpreting Results

Note that the simulations periodicity check is performed as follows: trajectory positions are recorded as reference points at random time intervals. After the trajectory leaves a small neighborhood around the reference point, its distance from the point is continuously monitored, and trajectories which return to the point's neighborhood are labeled as periodic (starting after the trajectory has left the neighborhood of the reference point, to exclude fixed point solutions). Clearly, the distance threshold for detection must be sufficiently small so as to exclude transient spirals towards fixed points. For any threshold, there could exist a sufficiently slow transient decay such that the trajectory was detected to be periodic, but the perspective of this Chapter would consider very slow transients (as might re-enter a sufficiently small neighborhood of a reference point) to be functionally periodic, and so this method of checking periodicity suffices.

Note that the algorithm, by design, continuously searches for new stable solutions which are not necessarily close to old branches of the diagram. These exploratory processes are unnecessary if the number and type of fixed points within an input region are already known. The higher resolution bifurcation diagrams, such as those within Figures 3.4 and 3.5, were therefore generated by initially generating a low-resolution diagram, then performing simulations to fill in each branch. For example, the fact that there appears to be only one stable solution for all inputs of Figure 3.4 means that the diagram can be generated simply by running a single simulation at each input point. Similarly, for Figure 3.5, one can use the basin of attraction as in Figure 3.6 to choose two initial points, in the first and third quadrant, which are likely to converge to the fixed point and limit cycle solution respectively. Thus increasing the resolution of these branches was seen to typically require only two simulations per input level.

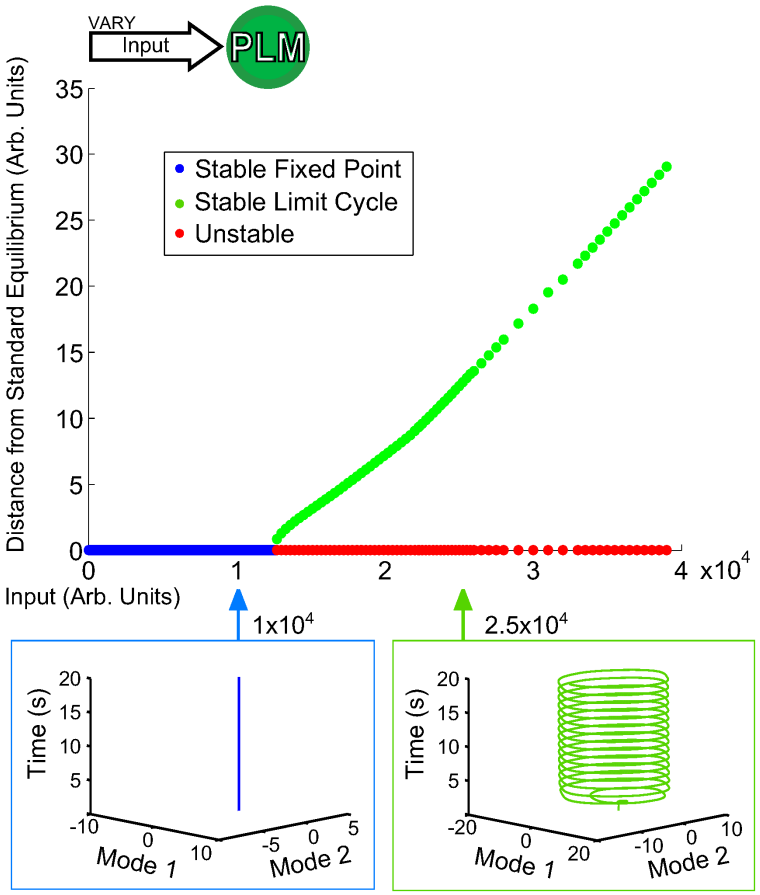

Figure 3.4: Bifurcation diagram for constant PLM stimulation of varying amplitude. Below an input of 1.2×10^4 the system goes to a stable fixed point very close to the standard equilibrium, but beyond that input level the system goes to a stable limit cycle (where the plotted point gives the furthest distance from standard equilibrium on the limit cycle). The diagram shows the two qualitatively distinct regions of interest for PLM inputs: the low input level in which the system remains at a fixed point, and the higher input level beyond which the system enters into a limit cycle (which in this case can be considered to serve as a proxy for forward motion [68]).

3.4.3 A Defining Example: Response to PLM Input

Figure 3.4 shows a low-dimensional bifurcation diagram for constant PLM input. The figure shows the creation of a stable limit cycle in response to input into the neurons PLML/R. By evaluating this bifurcation diagram we can identify the regions of interest which have qualitatively distinct responses (in this case, the region with a lone attractor which is a stable fixed point and the second region with a lone stable attractor which is a limit cycle after the fixed point attractor becomes unstable). For each region we can perform simulations which are then projected onto the low-dimensional plane (the PLM limit cycle being what defines this plane). Given the correspondence of this limit cycle to forward motion, as in [68], these low-dimensional trajectories are readily interpretable: the fixed point corresponds to a static worm, and the limit cycle corresponds to oscillatory motion of the body of the worm.

3.4.4 Characterizing Bistable Dynamics

Of greater interest are responses to compound activations; that is, more complicated inputs leading to more complicated responses. We consider as an example the dual input into the PLM and ASK neuron pairs as discussed in Section 3.2.3. We keep a constant input of 2×10^4 into the PLM pair and use as our bifurcation parameter the input into the ASK pair. Figure 3.5 shows the resulting bifurcation diagram. At inputs below 1.5×10^4, the limit cycle remains relatively undisturbed. At greater inputs, however, a series of bifurcations occur such that there is a sudden jump in the distance of the limit cycle, and at about 1.7×10^4 the system becomes bistable with the addition of a new fixed point. Thus we are able to immediately identify from this figure multistability within the system, which we may then go on to investigate further. Specifically, we are interested in the further investigation of transient timescales of the system.

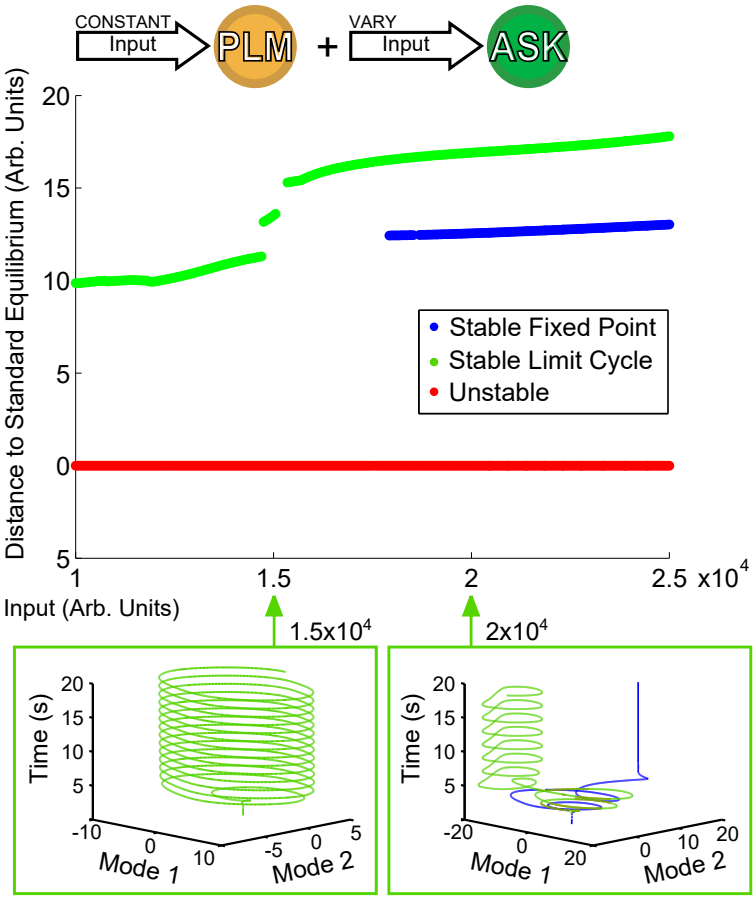

Figure 3.5: Bifurcation Diagram for varying amplitude of input into the ASK pair. Input into PLM is fixed at 2×10^4. Note that as input into ASK increases, the forward-motion limit cycle remains relatively undisturbed until it reaches about 1.5×10^4, after which the distance jumps and a fixed point becomes stable, giving rise to a bistability within the system.

3.5 Long Transient Timescales

In Figure 3.6 we investigate spatial and temporal aspects of the convergence into one of the two bistable states. An ensemble of 200 simulations (with random initial conditions in the neighborhood of the standard equilibrium) were performed for each ASK input level. From those, the trials converging to the fixed point solution were taken and the convergence time τ was calculated by calculating, for each fixed point trial, the time after which all points of the trajectory are within a distance ϵ of the final value (using here $\epsilon = 0.004$). The average and standard deviation of these convergence times are shown in the top right of Figure 3.6. Convergence times for the limit cycle solutions are qualitatively similar when comparing trajectories such as those in the upper-left of the figure. Note that these convergence times are considerably longer than other timescales within the system (comparing, for example, the model's free neuron decay constant of 10 ms [122], or other trajectory timescales such as the limit cycle period, which remains approximately two seconds regardless of ASK input).

Shown also are the basins of attraction for trajectories starting on the low-dimensional plane, on a grid of initial conditions centered at the standard equilibrium point (which we choose as our origin). The size of the grid is chosen to be within a small neighborhood of zero (within the range $(-4, 4) \times 10^{-6}$) since we find that trajectories initiated farther away are first attracted towards the zero point before being rerouted to the fixed point or limit cycle attractors.

Note that this structure appears only when initializing on the plane itself (that is, when displaced from the standard equilibrium point only along the two modes, rather than in a random direction). Random initializations about the standard equilibrium, when projected onto the plane, do not show this convergence structure. Likewise, displacing along the plane from a point other than the standard equilibrium (chosen since this is where the system will fall under no input) will show a different structure. Therefore these basin of attraction plots indicate distinct regions in which initial conditions starting on this plane will, under compound PLM+ASK input, go towards either solution (i.e. initializing in the most of the

upper-right portion of the plane leads to the fixed point solution, whereas initializing the the lower-left portion of the plane leads to the limit cycle solution). Since this is the plane on which the PLM response limit cycle proceeds, these figures show which portions of forward movement (i.e. which segment of the PLM-driven limit cycle trajectory, which goes through both of these regions) are more prone to ASK-driven transitions into a fixed point.

The existence of bistability between a limit cycle and fixed point within our model given simultaneous PLM stimulation (which promotes forward motion) and ASK stimulation (which promotes turning) is thus suggestive of an interpretation in which the worm's motorneuron activity will either, depending on the neural state upon the onset of this stimulus, oscillate or approach a fixed state.

3.6 Discussion

We explored the input space of a *C. elegans* neural dynamic model which incorporates its fully-resolved connectome and demonstrated that various multistabilities arise in response to inputs. Using a low-dimensional projection space based upon forward motion, we are able to systematically explore responses to complex inputs and understand them in a framework of low-dimensional attractor dynamics. In our study, the bifurcation diagram is constructed by using the constant-in-time input as our bifurcation parameter. We show that such diagrams are capable of revealing and mapping multiple attractors within the system by using a low-dimensional projection space which guides the search for attractors, identifying their stability and their effect upon forward movement. Furthermore, the low-dimensional projection helps in the interpretation of the dynamics upon the discovered attractors, especially the dynamics associated with multistability. We characterize such multistable dynamics, noting specifically that when the system enters into a multistable regime, transient timescales within the system can be very long relative to intrinsic neural timescales (comparing, for example, the three orders of magnitude between the $\sim$100 ms neural timescales in Section 3.2.1 to the $\sim$10 s transient lengths in Figure 3.5).

The fact that multistability within the connectomic dynamical system is capable of gen-

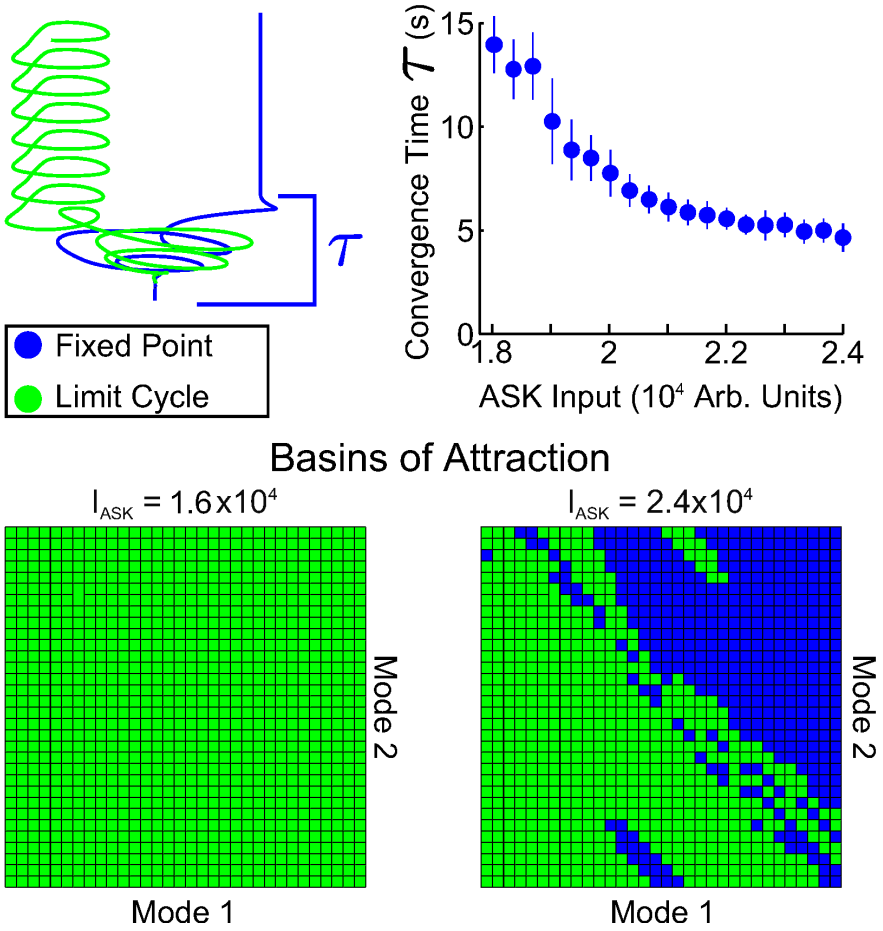

Figure 3.6: Spatial and temporal properties of convergence for PLM+ASK input (i.e. the bistable region of Figure 3.5). The upper-right plot shows fixed point convergence times as a function of input amplitude (from 200 trials at each point). Note the relatively long transient timescale. The second row shows the spatial basins of attraction for different inputs. Each grid covers a small region around the standard equilibrium, plotting on $(-4, 4) \times 10^{-6}$ for both modes. At an ASK input of 1.6×10^4 all initial conditions converge to a limit cycle, but initial conditions on the plane are split between the limit cycle and fixed points at higher inputs such as 2.4×10^4.

erating such long transient timescales has critical implications. These longer timescales, on the order of many seconds, are on a similar order to many behavioral timescales such as forward crawling survival time [109]. This suggests that various behavioral responses could be associated not with the attractor itself, but rather with the transient leading to that attractor. This is consistent with theoretical constructions and experimental observations of transient orbits between attractors [94, 95, 93]. Importantly, this viewpoint is supported *independently* and in a completely different theoretical framework by direct connectomic simulations from biophysically appropriate neuron dynamics within the worm, i.e. the multistability of attractors and long-time transients are not engineered in the model to fit the data and observations, rather they naturally arise from the dynamics associated with the connectome. Future work may analyze the structure of perturbations which drive transitions between different states, to predict inputs which result in said behavioral transients.

This study suggests that neural computations can consist of both dynamics on attractors (as in our PLM-driven limit cycle) and of long-timescale transients towards multiple attractors which may arise in the system (as we show in the long-timescale transients approaching the multistable states from PLM+ASK input). We have demonstrated that both dynamical features can arise by applying simple, identical neuron models onto the *C. elegans* connectome data, suggesting that these responses are encoded within the connectome itself.

More broadly, many networked dynamical systems across the engineering, physical and biological sciences may also be dominated by patterns of activity and long-time transients induced by the structure of the network architecture. Understanding the basic principles of such behaviors is critical for optimizing performance and controlling deleterious effects. One can easily imagine scenarios in which suppressing transients would be important, such as the observed power-grid swing instability [113]. The analysis above may be able to help understand how the network architecture encodes such deleterious patterns of activity when combined with relevant dynamics. In contrast, one might desire to generate a network architecture to induce a transient that is beneficial for some purpose relative to an application (for instance, a crawling motion in the case of the *C. elegans*). Understanding how the

network connectivity graph drives such activity would be critical for inducing such beneficial patterns of activity, perhaps even suggesting network control protocols for achieving desired results. The theoretical framework presented here highlights the rich and complex dynamics that emerge with networked architectures.

Chapter 4

ROBUSTNESS OF INJURED NETWORKS: LINKING BEHAVIORAL DEFICITS TO NEURAL CIRCUIT DAMAGE

With the emergence of a dynamic connectomic model for the nematode *C. elegans*, we address how biological network architectures and their functionality are degraded in the presence of focal axonal swellings (FAS) arising from neurodegenerative disease and/or traumatic brain injury. Using biophysically measured FAS distributions and swelling sizes, we are able to simulate the effects of injuries on the neural dynamics of *C. elegans*, showing how damaged networks degrade the functionality as encoded in low-dimensional representations of its behavior, i.e. eigenworm modes. We show that there is a diversity of functional deficits that arise from the same level of injury on a connectomic network. Functional deficits are quantified using a statistical shape analysis, a procrustes analysis, for deformations of the limit cycles that characterize key behaviors such as forward crawling. Such studies are directly relevant to understanding, and potentially treating, cognitive deficits arising from TBI and neurodegenerative diseases. More critically, these studies can help close the enormous gap that currently exists between theory and experiment in understanding how biological signals are processed in high-dimensional networks.

4.1 Author Summary

Neurodegenerative diseases such as Alzheimer's disease, Creutzfeldt-Jakob's disease, HIV dementia, Multiple Sclerosis and Parkinson's disease are some of the leading causes of cognitive impairments and death worldwide. Similarly, traumatic brain injury, the signature injury of the Iraq and Afghanistan wars, affects an estimated 57 million people worldwide. All of these conditions are characterized by the presence of focal axonal swellings (FAS) throughout the

brain. On a network level, however, the effects of FAS have remained unexplored. With the emergence of models which simulate an organisms full neuronal network, we are poised to address how neuronal network performance is degraded by FAS-related damage.

With our model for the full-brain dynamics of the nematode *Caenorhabditis elegans*, we are able to explore the loss of network functionality as a function of increased neuronal swelling. The relatively small neuronal network generates a limited and tractable set of functional behaviors, and we develop metrics which characterize and quantify how these behaviors are impaired by injuries to the network. These metrics quantify the severity of TBI and/or neurodegenerative disease, and could be potentially used to construct diagnostic tools capable of identifying a variety of cognitive deficits.

4.2 Introduction

All networks, and particularly biological networks, are susceptible to pathological and/or traumatic events that might compromise their performance. In neuronal settings, this may be induced by neurodegenerative diseases [31, 67, 119], concussions, traumatic brain injuries (TBI) [85, 103, 135] or aging. This chapter develops a computational model to investigate behavioral impairments in the nematode *C. elegans* when the underlying neuronal network is damaged. Specifically, we consider how the low-dimensional population codes are compromised under the impact of an injury. Characterizing the resulting cognitive and behavioral deficits is a critical step in understanding the role of network architecture in producing robust functionality.

A hallmark feature of damaged neuronal networks is the extensive presence of Focal Axonal Swellings (FAS). FAS has been implicated in cognitive deficits arising from TBI and a variety of leading neurological disorders and neurodegenerative diseases. For instance, FAS is extensively observed in Alzheimer's disease [67, 119], Creutzfeldt-Jakob's disease [73], HIV dementia [1], Multiple Sclerosis [44, 118] and Parkinson's disease [46]. Most concussions and traumatic brain injuries also lead to FAS or other morphological changes in axons [56, 57, 75, 81, 107, 126]. Such dramatic changes in axon geometry may disrupt axonal transport [116,

117], and can potentially hinder the information encoded in neural spike train activity [76, 77, 78]. Injured axons thus provide an important diagnostic marker for the overwhelming variety of cognitive and behavioral deficits [31, 76, 83], in animals and humans [9, 28, 51, 81].

The massive size of human neuronal networks and their complex activity patterns make it difficult to directly relate neuronal network damage to specific behavioral deficits. The nematode *C. elegans*, in contrast, has only 302 neurons, and its stereotyped connectivity (i.e. the worm's "Connectome") is known [122]. This relatively small neuronal network generates a limited and tractable set of functional behaviors (see Table 1 of [43]), with much of its behavior approximately confined to five observable motor states related to forward and backward crawling, omega turns, head sweeps and brief pause states. Furthermore, these behaviors are well described as a superposition of only a few principal component body-shape modes [110]. The combination of a fully-resolved neuronal network and a tractable low-dimensional output space makes the *C. elegans* an ideal model organism for studying the impact of network damage on behavioral deficits. Indeed, it is the only such neuronal network model currently available allowing for such a direct translational study of network damage (injury) to behavioral responses.

More precisely, our computational model of *C. elegans* connectome dynamics successfully generates motorneuron outputs that can be related to behavior, allowing for interpretable outputs even without accounting for muscular, mechanical or environmental factors. Figure 4.1(1) demonstrates a simulation of the forward crawling behavior in the *C. elegans* along with its projection into principal component body-shape modes [110]. In this perspective, forward crawling is understood as a limit cycle in principal component space. Extending this framework to *damaged* networks as in Fig. 4.1(3) allow us to explore how axonal pathologies lead to impaired functionality and behavioral deficits. Even in our idealized injury simulations, the network's impaired activity displayed significant variability. This highlights one of the most challenging aspects of the field: the need for effective *metrics* to distinguish different types of behavioral deficits. We propose such a criterium by using techniques borrowed from statistical shape analysis to quantify distortions in the main features of dynamical

activity. This leads to experimentally-testable predictions about the effects of neuronal network-damage to the *C. elegans* crawling motion and potentially new avenues for clinical diagnostics. Indeed, our studies show that network damage leads to a diversity of behavioral and cognitive deficits.

4.3 Results

4.3.1 Low-dimensional Signatures for Crawling Behavior

We investigate how network distributed FAS as illustrated in Fig. 4.1(3) may affect its ability to generate desired responses to an input. Network features associated with behavioral outcomes are best understood in model organisms such as the *C. elegans* since it has a limited repertoire of functional responses that include forward and backward crawling, omega turns, head sweeps and brief pause states. Our focus in these studies will be on the behavior of forward crawling since a variety of experimental ablation studies have identified key neurons associated this functionality. For instance, experiments show that stimulation of PLM neurons excites densely-connected interneurons, which in turn, activate motorneurons responsible for forward body motion.

Details of the underlying neurocircuitry were found by a series of ablation studies, where the functional role of a neuron is evaluated by disconnecting it from the network and observing behavioral deficits [25, 68]. The coordinated body motion of a crawling worm is well documented in videos and was revealed by principal component analysis to be a simple alternation of its two most dominant modes [110]. It is remarkable that an apparently complex functional behavior can be well described by low-dimensional structures. More specifically, by circular trajectories (limit cycles) on the phase-space of its principal components. An analogous mathematical form is present in the collective motorneuron dynamics following PLM stimulation [68].

This commonality suggests that observed behaviors do retain fundamental signatures of the underlying network dynamics. We show such a trajectory for (simulated) motorneuron

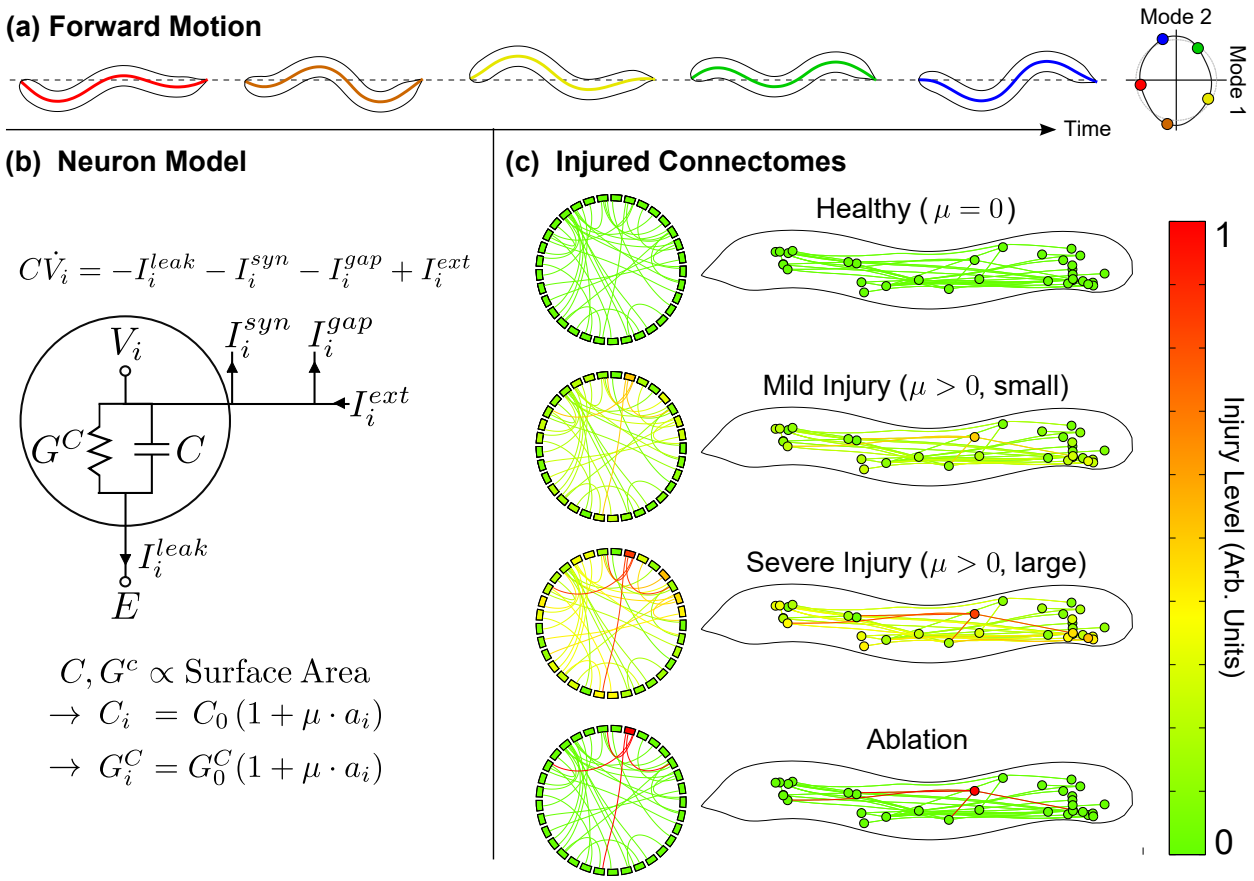

Figure 4.1: **(a)** In our model, stimulating PLM neurons drives two-mode motorneuron oscillations. We project all dynamics onto these modes. We map these projected dynamics onto the behavioral modes from [110] to reconstruct theoretical body motion. **(b)** We model neuron injury by scaling membrane capacitance and conductance with surface area. The relative swelling of neuron i is set by pulling its individual swelling level a_i randomly from a distribution, and scaling all swellings by an overall "injury amplitude" μ. **(c)** We refer to a particular set of a_i values as the same "injury". Here we illustrate the same injury in three different regimes of μ. Compare the common experimental case of ablation, in which only one neuron is injured very severely (as opposed to our distributed injuries).

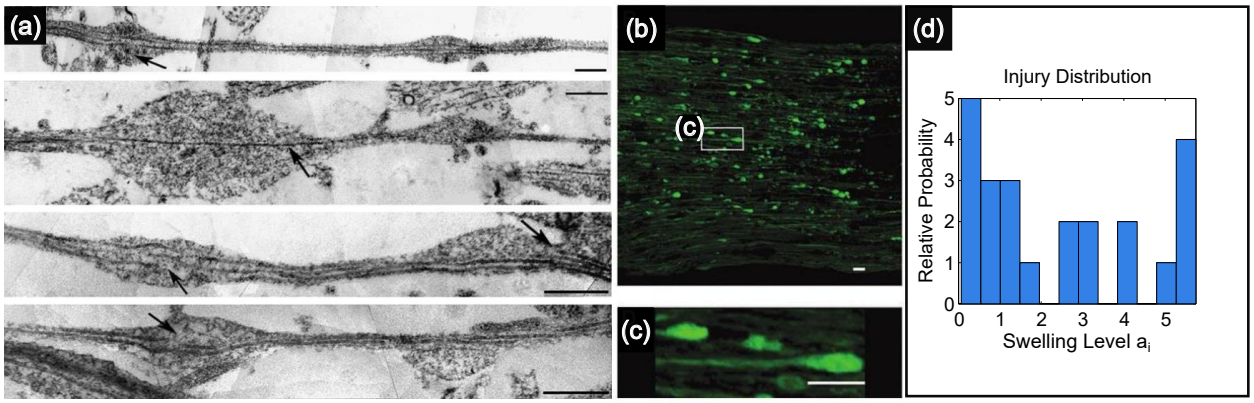

Figure 4.2: We draw swelling values from experimental distributions. **(a)** Transmission electron microscopy image of axons 3 hr. after stretch injury, modified from [116]. **(b** and **c)** Confocal microscopy images of injury-induced swellings in the optic nerve of Thy1-YFP-16 mice, modified from [126] (scale bar 20μm). **(c)** It is equally simple to use any swelling distribution under our approach. We choose axonal swelling data from [126] (as in (b) and (c)), taken 12 hr. post-injury, from which we calculate the above probability distribution for neuron swelling levels a_i.

responses to PLM excitation in Figure 4.1(1). This low-dimensional representation captures 99.3% of the total energy of the system, and can be artificially mapped to crawling body-shape modes. Although this mapping is still far from a mechanistic description of the worm's coordinated body movement, we believe it captures important aspects of the crawling behavior. See the Materials and Methods for details. Importantly, functional deficits of the *C. elegans* dynamics are understood as excursions/perturbations from the ideal limit cycle trajectory. Damaged networks will be shown to fail to produce the low-dimensional output codes necessary for generating the optimal forward crawling limit cycle.

4.3.2 Modeling Injured Connectomic Dynamics

The robustness of the dynamical signatures (population codes) associated with behavior are investigated in injured neuronal networks. Our injury statistics and FAS models are drawn from state-of-the-art biophysical experiments and observations of the distribution and size of FAS. Figure 4.2 shows prototypical FAS injuries from stretching [116] and TBI in the optic

nerve of mice [126]. Figure 4.2(4) shows a histogram of the probability of injury and size of the FAS. These are used in our computational model [68]. In a simulated injury, we assign to each affected neuron an axonal swelling from the distribution in Figure 4.1(2). Values are scaled by an (overall) injury intensity parameter μ, such that

$$\mu \propto \mathbf{E}\big[\text{axonal swelling area}\big]. \tag{4.1}$$

Figure 4.1(3) exemplifies different injury settings: $\mu = 0$ reproduces the original (uninjured) network, and lower/higher values of μ correspond to mild/severe injuries. The presence of axonal swellings ultimately distorts the forward-motion limit cycle dynamics. Figure 4.3 shows dynamical anomalies for different connectome injuries. Notice how they induce qualitatively different changes to the closed orbit regarding location, size and shape. Figure 4.3(3) reproduces the specific simulated ablations from [68], leading again to different dynamical effects.

A much larger ensemble of simulations (95 randomly-chosen injuries) were performed and their corresponding effects to fundamental low-dimensional structures were analyzed. The code for generating additional injury simulations is open-source and freely available from the authors. Increasing values of μ typically shrink and shift the limit cycles within the plane. In all simulations, there was always a sufficiently high injury level in which

$$\mu^* = \{\text{injured limit cycle collapses into a stable fixed point}\}. \tag{4.2}$$

This occurs for instance, in Figure 4.3(2) when $\mu = 6.6$. Recent blast injury studies on *C. elegans* show that many of the nematodes display temporary paralysis before recovering to crawling behaviors [6]. We would suggest that during the peak of the FAS, the injury levels on many of the nematodes are above μ^*, thus leading to a collapse of a limit cycle to a fixed point where no behavior is possible, i.e. it is in a paralyzed state.

4.3.3 Distinguishing Signatures of Different Behavioral Deficits

Despite their common statistical distribution, randomly drawn injuries induce qualitatively different changes in the *shape* of the limit cycle. Additional distorted sets are shown in the

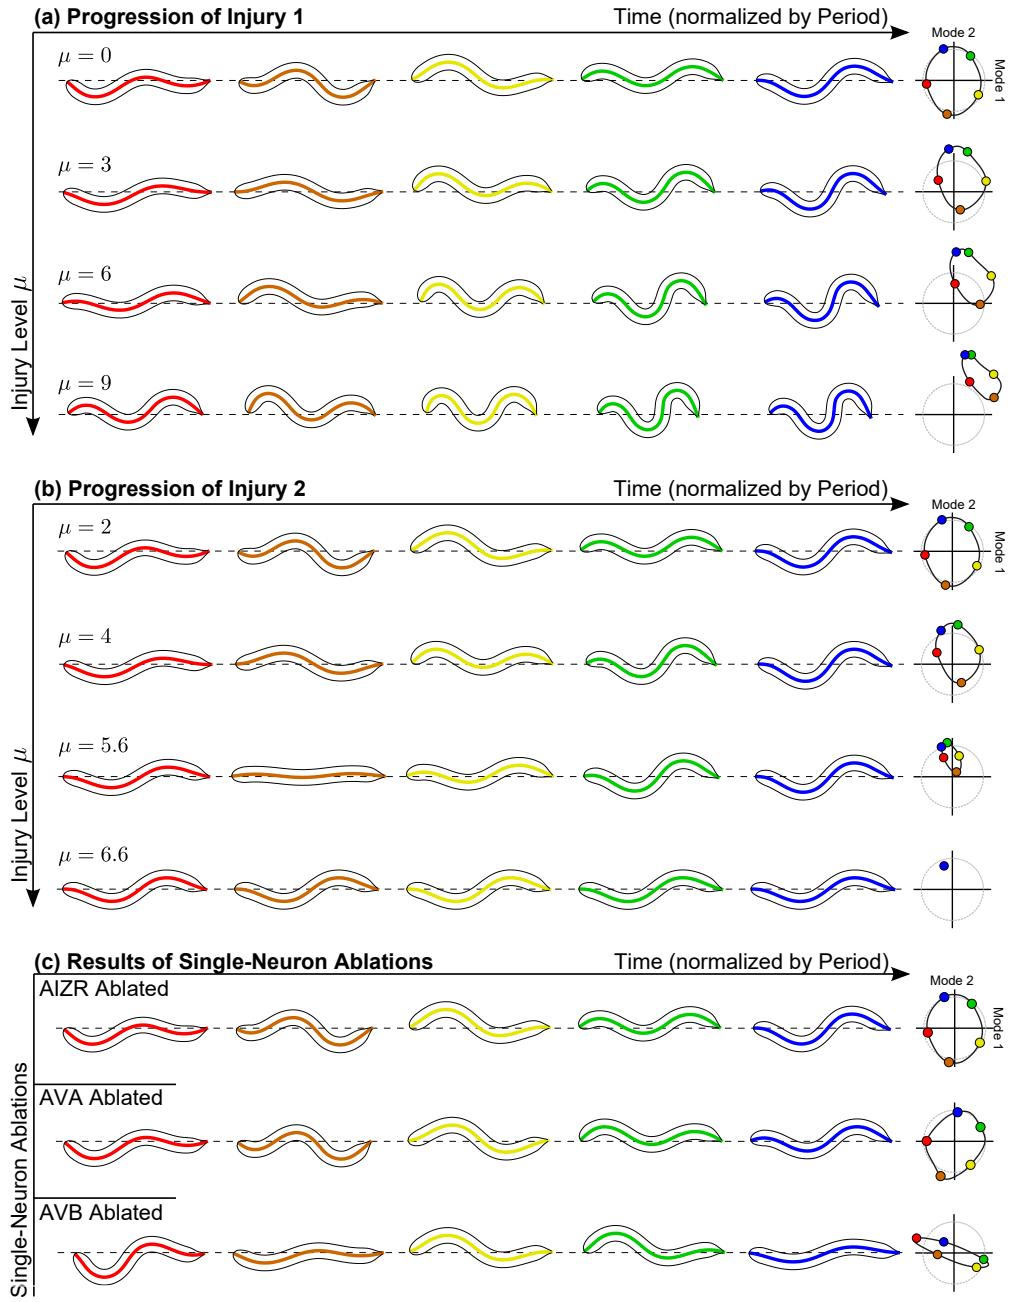

Figure 4.3: Response to PLM stimulation for two different distributed injuries in **(a)** and **(b)**. Resultant trajectories were mapped onto the two neural modes, which were then mapped onto the two behavioral modes to reconstruct theoretical body dynamics. Within each panel, injury amplitude μ is increased in subsequent rows. Different injuries yield qualitatively different injury progressions as μ is increased, and sufficiently high μ values lead to the cycle collapsing into a fixed point. **(c)** For comparison, consider the all-or-nothing injury effects of isolated single-neuron ablations.

rows of Figure 4.4 (along with 95 random-injury simulation sets). Thus, random injuries of equitable strength can lead to significantly different behavioral deficits. Importantly, the deformation of the two-dimensional limit cycle can be used to characterize such functional differences. To distinguish dynamical signatures of potentially different functional deficits, we evaluate the *Procrustes Distance* (PD) between healthy and injured limit cycles. The PD is an important tool from statistical shape analysis to measure the similarity between two shapes after discounting effects due to translation, uniform scaling, or rotation. Figure 4.4 depicts PD values for pairs of healthy/injured limit cycles as a function of injury level μ. All curves are plotted until the injured limit cycle collapses into a fixed point ($\mu = \mu^*$), and the colored dots in the rightmost plots correspond to the same-colored limit cycles on the left plots. Increasing μ-levels distort the limit cycle with potentially different trends, which are well captured by three classes of PD curves:

(1) *Rising Curves.* Here the PD rise smoothly and mostly monotonically as a function of μ, ending with a small decrease when the limit cycle collapses.

(2) *Peaked Curves.* These PD curves have a large and smooth monotonic increase followed by an extended decreasing regime.

(3) *Discontinuous Slopes.* These PD curves are nearly monotonic but display a kink after which the PD rises much more sharply.

We show another example of a PD curve with a discontinuous slope in Figure 4.4(4): for some μ values, the limit cycle crosses itself and form an inner loop. The PD curve increases dramatically when the limit cycle is not simple.

4.3.4 Procrustes Distance Classes and Behavioral Dynamics

Recent experimental work which induced mild TBI in *C. elegans* found that increasing the number of shock waves to which the worm was exposed reduced the worm's average speed

and, in many cases, led to temporary paralysis [6]. The results of our simulations can be compared to these results:

- *Speed Reduction:* In our model, injury tends to slow the oscillation of the limit cycle. Specifically, the temporal frequency of the limit cycle was reduced by an average of 17% in the highly-injured interval $(0.9\mu^*, \mu^*)$. However, a slower limit cycle frequency does not necessarily imply slower movement. As the frequency changes, so does the amplitude and shape of the limit cycle, and these will also affect the forward movement speed. Without a coupled mechanical model for the body movement of the worm and for the environment in which it moves, we are unable to calculate how these trajectory distortions affect forward movement speed.

- *Temporary Paralysis:* Simulated neural patterns are static at the fixed point when $\mu > \mu^*$. Neglecting extra-connectomic effects (e.g. electrical coupling within muscles themselves) this may imply that the worm is not moving. The point at which the trajectory stops should correspond to the fixed shape of the worm (note that this depends on the full-dimensional location of the point, not simply our plane projection). Thus, if the fixed points are clustered according to their PD class, the PD class should carry information about the final body shape of the worm for a given injury (i.e. different shapes that the worm assumes when paralyzed should correspond to the PD class of the injury).

In Figure 4.5 we plot the location of the fixed points into which limit cycles collapse (at injury level $\mu = \mu^*$). Each fixed point is colored according to its PD curve class. Some separation between classes is visually apparent and confirmed by various data clustering techniques: Linear/Quadratic Discriminant Analysis (LDA/QDA), k-Means and Gaussian mixture models (see Figure 4.5 for details). All methods perform significantly better than the expected accuracy of a random classification of normal distributed data points.

We believe these three classes of PD curves carry dynamical signatures for different functional deficits. We currently do not have experimental evidence to match these results with mechanistic predictions of impaired body movement. Still, the same artificial map (used in Figure 4.2) may provide insights on how distorted limit cycles relate to anomalous crawling behavior. Videos illustrating these distorted limit cycles are available online, to be included with the publication of the paper corresponding to this Chapter. These results could lead to novel functional diagnostics for neural networks under pathological conditions.

4.4 Methods

4.4.1 Governing Equations

As in previous chapters, we assume that the membrane voltage dynamics of neuron i is governed by:

$$C_i\dot{V}_i = -G_i^c(V_i - E_{cell}) - I_i^{Gap}(\vec{\mathbf{V}}) - I_i^{Syn}(\vec{\mathbf{V}}) + I_i^{Ext}. \tag{4.3}$$

As before, the parameter C_i represents the whole-cell membrane capacitance, G_i^c the membrane leakage conductance and E_{cell} the leakage potential of neuron i. The external input current is given by I_i^{Ext}. In all simulations within this Chapter, we set I_i^{Ext} to be constant for the PLM neuron pair and zero for all other neurons. This assures that densely connected interneurons will stimulate the motorneuron subcircuits responsible for forward crawling behavior. Neural interaction via gap junctions and synapses are modeled as before, with unchanged parameters.

We model the effect of injuries by altering the values of cell membrane conductance and capacitance, which have uninjured values set to $G_i^c = G_H^c = 10$ pS and $C_i = C_H = 1$ pF. Note that in uninjured simulations, all neurons are modeled as identical except for their connectivity and the assignment of them as excitatory or inhibitory (where E_j will have one of two values corresponding to these classes).

The model is valuable for this study because (as detailed in Chapter 1 it generates a low-dimensional neural proxy for behavioral responses. Specifically, constant stimulation of

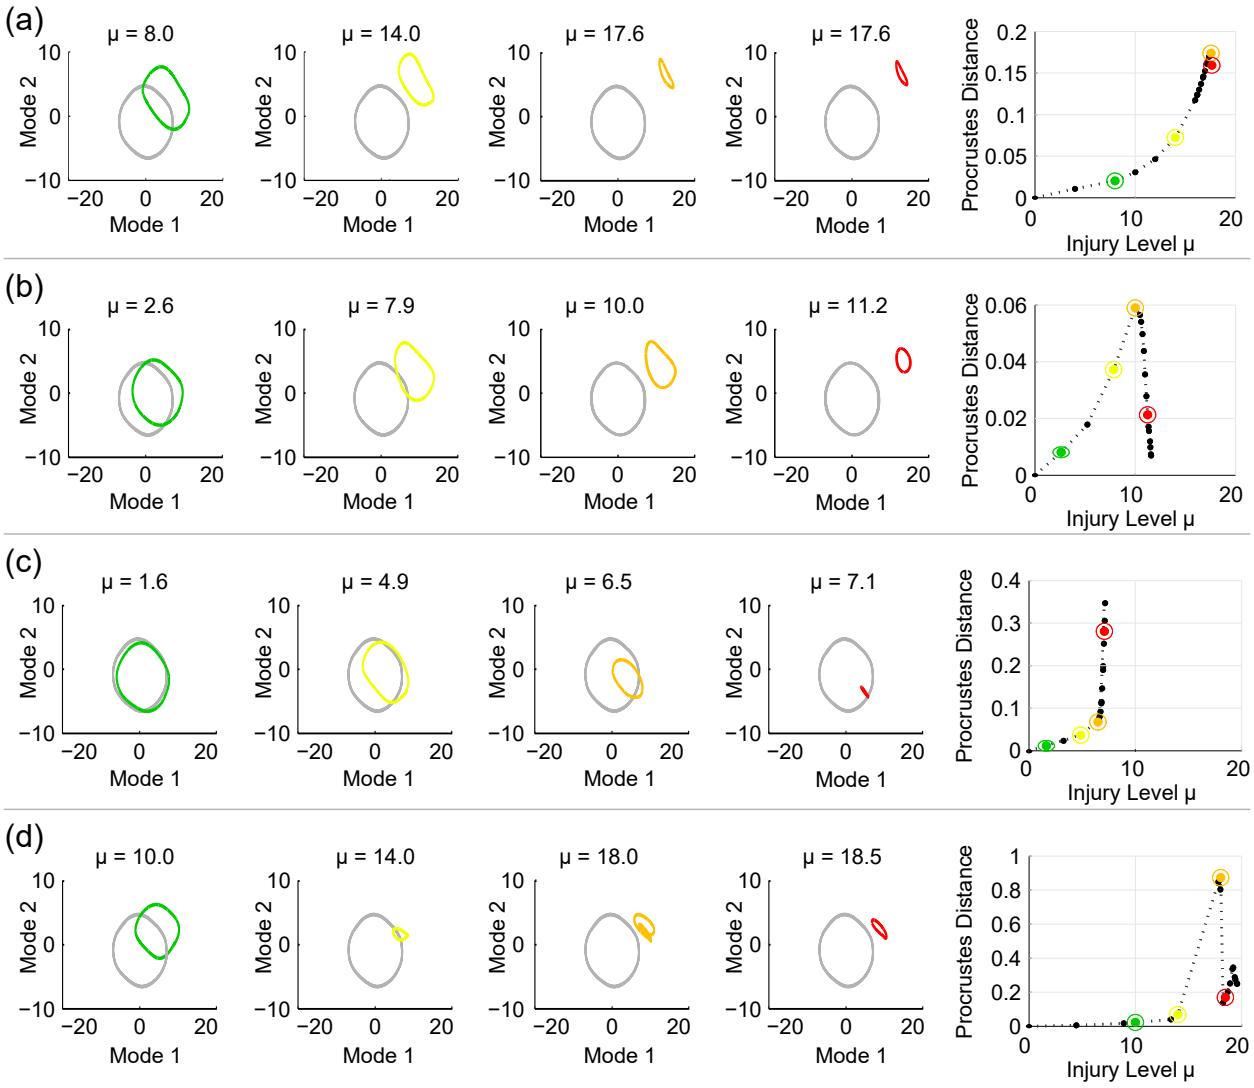

Figure 4.4: Each row depicts projected neural responses for different types of injury with increasing intensity μ. Distortion of each cycle is quantified via the Procrustes Distance (PD), a metric that compares shapes ignoring translation, rotation or uniform scaling. The PD curves terminate when $\mu = \mu^*$, level for which the cycle collapses into a fixed point (i.e. paralysis). Colored dots on the rightmost plots correspond to the same-colored limit cycles on the left plots. By visual classification of our ensemble of 100 Procrustes curves (for different injuries), there are three typical cases: **(a)**, in which the PD curve is nearly monotonic; **(b)**, for which there is an extended drop-off region in the PD curve preceding the collapse of the cycle; and **(c)** and **(d)**, PD curves with a discontinuous slope or an apparent discontinuity.

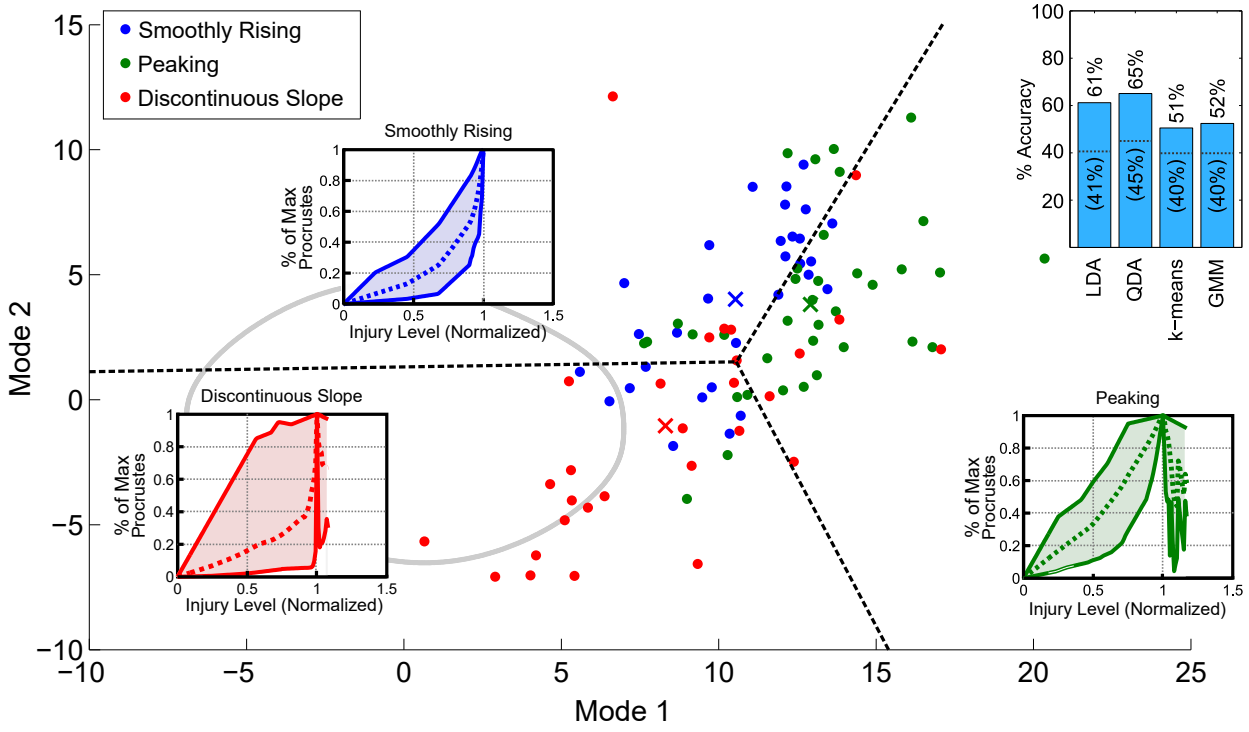

Figure 4.5: The Procrustes curves carry information about the functional outcome of an injury. Plotted on our neural plane are the locations of the fixed points into which injured cycles eventually collapse, for sufficiently high μ (color-coded by Procrustes curve class). Classification lines from the linear discriminant analysis are plotted, and the centers of each class are labeled by an "x" marker. For both supervised and unsupervised classification techniques, prediction of Procrustes curve class from the fixed point location is significantly better than random (see the Methods section "Classification of Deficient Behavioral Responses" for additional details). This indicates that analyzing the shape of the Procrustes curve as μ increases may help predict the form of paralysis resulting from that specific injury.

the tail-touch mechanosensory pair PLM creates a two-mode oscillatory limit cycle in the forward motion motorneurons [68]. This same dynamical signature was revealed in video analysis of the body-shape of the crawling worm [110]. Thus the model is consistent with the observed biophysics.

4.4.2 Modeling FAS in Neuronal Network Simulations

Focal Axonal Swellings (FAS) increase the volume of an axon, which in turn, should alter the cell's capacitance and leakage conductance. If we approximate a neuron by a single cable of length l and constant cross-section a, the circuit parameters will scale with the axonal volume, i.e.,

$$C \propto a \cdot l \tag{4.4a}$$

$$G^c \propto a \cdot l \tag{4.4b}$$

When an axon swells, its healthy cross-sectional area a_H will increase to some swollen value $a_i > a_H$. The healthy values for capacitance C_H and conductance G^c_H will also change according to

$$C_i = C_H \cdot (1 + a_i/a_H) = C_H \cdot (1 + \mu \cdot m_i) \ , \tag{4.5a}$$

$$G^c_i = G^c_H \cdot (1 + a_i/a_H) = G^c_H \cdot (1 + \mu \cdot m_i) \ . \tag{4.5b}$$

The individual damage of neuron i is proportional to the relative swelling area, i.e., $m_i \propto a_i/a_H$. Values of m_i are computed from the experimentally derived distributions in Figure 4.2. Specifically, we construct FAS from the axonal swelling data of Wang et al. [126], which used confocal microscopy to measure injury-induced swellings in the optic nerve of Thy1-YFP-16 mice. The damage to an individual axon, if damaged, is then scaled by a *global injury intensity* defined as follows:

$$\mu = \frac{\langle a_i/a_H \rangle}{\langle m_i \rangle} \tag{4.6}$$

Mild traumatic brain injuries yield small values of μ indicating that the average area of swollen axons is small. Severe brain injuries yield high values of μ, indicating that large swellings are more common. We leave the PLM pair of neurons receiving input uninjured. All other neurons have their m_i values assigned from the experimental statistical distributions. The governing equation for an injured neuron is now

$$C_H \dot{V}_i = -G_H^c(V_i - E_{cell}) - (I_i^{Gap}(\vec{\mathbf{V}}) + I_i^{Syn}(\vec{\mathbf{V}}))/(1 + \mu \cdot m_i). \tag{4.7}$$

We can readily interpret the limiting cases: when $\mu \cdot m_i = 0$, the original governing equation is recovered, and thus $\mu = 0$ corresponds to the healthy case. When $\mu \cdot m_i$ is large, gap junction and synaptic currents have no effect. The neuron's voltage decays exponentially to its leakage potential, effectively isolating it from the network.

Note that our random assignment of swelling values neglects any spatial structure of the injury. This could be easily modified by using a distribution which depends on the spatial location of the neuron. Furthermore, this is a very simple model for neuronal swelling, in keeping with our simple model for neurons. It necessarily neglects the actual geometry of swelling. The use of a multi-compartment model would enable this in future studies. Ultimately, there is currently limited biophysical evidence for making more sophisticated models. As such, we have tried to capitalize on as many biophysical observations as possible so as to make a model that is consistent with many of the key experimental observations.

4.4.3 Numerical Simulations and Convergence Criteria

We use MATLAB to solve the system of neuronal dynamical equations. We consider an ensemble of 95 different types of injury (set of targeted neurons), for which the global intensity μ may vary from 0 (uninjured) to a critical value μ^*. When the intensity exceeds μ^* (found by a bisection algorithm), the limit cycle collapses to a fixed point. To obtain intermediate values, we perform five simulations linearly spaced throughout $(0, 0.9\mu^*)$ and ten additional simulations throughout $(0.9\mu^*, \mu^*)$.

We classify the resulting injured trajectories as a Fixed Point or a Periodic Orbit according to the following criteria:

(1) *Fixed Point:* when the trajectory is always confined to a circular region of radius of 0.01 (about three orders of magnitude below the uninjured radius).

(2) *Periodic Orbit:* when the trajectory escapes the circular region but re-enters it periodically.

Note that these criteria classify very small periodic orbits as fixed points, since their behaviors are very similar. The method may also classify sufficiently slow, long-timescale oscillatory transients as periodic.

4.4.4 Artificial Mapping of Dynamical Signatures to Behavioral Modes

Stephen et al. [110] found that the forward crawling motion of a *C. elegans* is well described by two principal component body-shape modes called *eigenworm* modes. When moving forward, the modes alternate within its phase space forming a limit cycle. Kunert et al. [68] (as outlined in Chapter 1) also found a two-dimensional limit cycle, but for the collective motorneuron activity after PLM stimulation. They interpret this similar dynamical signature as a neuronal analog to the observed behavioral pattern.

To interpret the distorted neural activity caused by our simulated injuries, we construct a map from the neuronal activity plane onto the eigenworm plane. The body-shape modes were extracted from Figure 2(3) of [110]. We first calculate the optimal linear mapping of the healthy trajectory onto a circle (see Fig. 4.3a). We then use this calibration for all other trajectories. This artificially translates anomalous neuronal responses to anomalous body motions. Our procedure has a number of limitations, for which we list a few:

(1) The behavioral limit cycle in [110] is approximately circular, but the relative rotation between the two planes is unknown. This parameter could be inferred by observing the motion of injured or ablated worms.

(2) It is unclear that the mapping would hold for injured worms, especially without accounting for body-shape modes (eigenworms) from impaired crawling behavior.

(3) We consider only the first two (healthy) behavioral modes. Thus, lack of motion within this plane does not necessarily imply that the worm is not moving. The injured body-shape dynamics could evolve along different modes leaving no traces on the original two.

The lack of direct neuronal analogs for injured network modes limits our ability to interpret arbitrary impaired behavioral responses. Further computational work could also find neuronal proxies for additional behavioral modes so as to enable a more complete mapping. Recent work on blast injuries of worms [6] could potentially help extend the analysis by providing injured *eigenworm* mode projections.

4.4.5 Procrustes Shape Analysis

Procrustes Distance (PD) measures the dissimilarity between shapes, and in our context, we wish to compare the *shape* of the trajectories of the healthy neural responses (circular orbits in the phase plane) with the distorted ones produced after simulated injuries. For that, we use the function `procrustes.m` from MATLAB's Statistics and Machine Learning Toolbox. We collect N data points from each trajectory and annotate their (x, y) coordinates in a $(2 \times N)$ shape matrix S. The PD between two distinct shapes S_A and S_B is given by

$$PD = \min_{b,R,c} \|S_B - b \cdot S_A \cdot R + \vec{c}\|_2. \tag{4.8}$$

In other words, it finds the optimal (2D) rotation matrix R, scaling factor $b > 0$, and translation vector $\vec{c}$ to minimize the sum of the squares of the distances between all points. Intuitively, it compares shapes discounting translation, rotation, or scaling. To calculate the PD curves as in Figure 4.4, we use the uninjured ($\mu = 0$) limit cycle as our first shape S_A. The second shape S_B is the limit cycle calculated for each injury at the indicated value of μ.

We pre-process the trajectories to extract data points only within a single period. Since injuries usually distort the trajectory length, we use MATLAB's `spline.m` function to interpolate them and collect the same number of data points. Both limit cycles must also be phase-aligned, which we achieve by finding the phase that minimizes the Procrustes Distance.

4.4.6 Classification of Deficient Behavioral Responses

We hypothesize that PD curves contain meaningful signatures of behavioral outcomes of a given injury. For example, there is always a critical injury level $\mu = \mu^*$ in which the injured response collapses into a fixed point. Our artificial map suggests that this corresponds to a paralyzed worm. We searched for (visually apparent) qualitative features of these curves to avoid overfitting the relatively small number of data points. We increase the intensity μ of each simulated random injury and visually inspect the resulting PD curves for similar trends. We describe three different classes in the Results Section: (1) Rising Curves, (2) Peaking Curves and (3) Discontinuous-Slope Curves. We use the following standard cluster-detection algorithms to verify if these classes are indeed informative: Linear/Quadratic Discriminant Analysis (which find optimal lines/quadratics to separate the labeled data points), k-Means and Gaussian Mixture Models. We use all these methods as implemented in MATLAB's Statistics and Machine Learning Toolbox. Their accuracy values were calculated as follows:

(1) We trained LDA and QDA classifiers on a randomly selected subset of 75 points, and then calculated how well they classify the remaining 20 points (totalizing 95 injury-types). We shuffled the training subsets 1000 times. The average accuracy value is 61% for LDA and 65% for QDA.

(2) k-Means and GMM do not take the PD classes into account in their classification. Accuracy is calculated then by simply comparing their 3-cluster classification to the actual classes, giving 51% for k-means and 52% for GMM.

These accuracy levels imply that the fixed points are, to some extent, clustered within these PD classes. To estimate the expected classification accuracies if these points were not clustered, we repeated this analysis on 100 sets of 95 randomly-labeled, normally-distributed data points. The average classification accuracies for these random sets were as follows: 41% for LDA, 45% for QDA, and 40% for k-means. Our data points exceed these values for each method. Therefore, the PD class (as defined by the shape of the Procrustes curve) contains information about which cluster the trajectory collapses into.

4.5 Discussion

This study introduces a tractable framework for analyzing how biophysically-inspired injuries distributed across a physical neuronal network induce behavioral deficits. The specific injuries we consider arise from FAS which has been implicated in most leading neurodegenerative diseases, aging and TBI. By identifying low-dimensional population codes within our model which correspond to a known behavior, a proxy metric for cognitive deficit can be constructed. Specifically, limit cycles in our dominant features serve as a neural proxy for actions such as forward motion in the *C. elegans*. Such trajectories can be artificially mapped onto experimental body-shape modes, and suggests a behavioral interpretation of the distorted limit cycles resulting from an injury. Our analysis also suggests that there is a diversity of functional deficits that arise from the same level of injury on a connectomic network.

The ability to provide a theoretical understanding of functional, cognitive and behavioral deficits due to connectomic injuries is a the forefront of TBI and neurodegenerative disease studies. Both have an enormous societal impact and implications. Specifically, TBI is annually responsible for millions of hospitalizations [42, 62], with reports estimating that 57 million people worldwide experienced some form of TBI [135]. It was also manifest in around 15% of all veterans of the Iraq and Afghanistan wars, with blast injuries being the signature wound of these conflicts [62, 135]. Numerous studies show that even mild concussions, if induced repeatedly, can lead to permanent brain damage; the issue is constantly

debated in the sports media, but especially in football [40]. Neurodegeneration affects orders of magnitude more people than TBI through diseases such as Alzheimer's disease [67, 119], Creutzfeldt-Jakob's disease [73], HIV dementia [1], Multiple Sclerosis [44, 118] and Parkinson's disease [46]. Thus, any study that can help understand how FAS compromises cognitive function is of great value.

4.5.1 Diagnostic Tools

Simulated injuries distort dynamical signatures in the network's activity, such as limit cycles. Our Procrustes Distance metric quantifies how much the shape of the limit cycle is distorted, compared to the healthy cycle. Our results indicate that as different injuries evolve, this metric follows qualitatively different trends (as in Fig. 4.4). In all trials, a sufficiently high injury level $\mu = \mu^*$ collapses the limit cycle into a stable fixed point. The shape of the PD curve carries information about where this fixed point will be (as in Fig. 4.5). This suggests that the shape of the PD curve, as the injury evolves, may predict the eventual behavioral outcome (e.g., the body shape the worm will assume during temporary paralysis).

More broadly, the metrics described in this work can potentially be used to construct diagnostic tools capable of identifying a variety of cognitive deficits. Moreover, the severity of a TBI injury and/or neurodegenerative disease can be quantified by measuring its metric distance from the normal/healthy performance. Our work gives clear mathematical tools capable of formulating such diagnostic tools for assessing injuries and functional deficits.

4.5.2 Suggested Experiments

We analyze simulated data for impaired network and behavioral activities. However, our methods could be directly applied to experiments:

(1) Record the motion of a healthy worm and calculate its leading two body-shape modes (as in [110]).

(2) Injure the worm (e.g. using a shock wave to induce mild TBI such as in [6])

(3) Record post-injury motion and project onto the two body-shape modes found in (1).

(4) Monitor the PD metric of the experimental trajectory, in real-time, in the time after the injury.

Doing this for many worms will create an experimental ensemble of low-dimensional trajectories and PD curves which can be directly compared to those in this study. As discussed in the Results section, different PD curves should be classified and compared to the body shapes assumed by the worm when it becomes temporarily paralyzed.

4.5.3 Limitations

The present study has many limitations, many due to the lack of biophysical evidence required to build better models. We predict the results of injuries only on the two "forward-motion" motorneuron modes, ignoring other modes potentially associated with impaired behaviors. Furthermore, the exact mapping of our motorneuron voltage modes onto these body-shape modes is ambiguous. The model lacks muscles and body features of the worm which limits our ability to make more general predictions. We also neglect external feedback mechanisms required for sustained and spontaneous forward motion, and assume that tail-touch neurons are constantly stimulated. It is uncertain how such feedback mechanisms would alter the trajectory. The order-of-magnitude parameter estimates of our model parameters also make direct quantitative comparisons difficult.

4.5.4 Future Work Related to Network Damage

We believe the merit of this study lies not so much on the specific results presented, but on the new directions and methodologies it opens for future work. In fact, computational and experimental studies on the effects of network injury are still at their infancy for the *C. elegans* and other models. Many limitations of this work could be overcome with a more detailed model for the *C. elegans* neuronal network both before and after injury. Coupling

this with an external, mechanical model would allow for more general predictions. This could be accomplished with simplified mechanical models for locomotion (such as in [10]) or with more complete, future "in-silica" models such as OpenWorm [114]. Our analysis techniques could be applied directly to such models. Experimental studies would not only test our model, but also in, in conjunction with our work, provide a new testbed for models of injured connectomic dynamics. Our Procrustes Distance metric, shown here to carry information about the eventual outcome of an injury, may also be useful in the real-time analysis of injury progression. Thus our study provides a way forward in monitoring behavioral outcomes of injured networks.

Ultimately at present, limitations in biophysical measurements and neural recordings make it extremely difficult to identify more sophisticated underlying mechanisms responsible for dysfunctions in neural networks, especially when circuits display intrinsically complex behavior and functional activity. We believe the rapid advancement of recording technologies in neuroscience will significantly help refine the model presented here.

Given that the modeling of neuronal networks is one of the most vibrant fields of computational neuroscience [33, 39, 47], our contribution provides a comprehensive study of how the effects attributed to FAS jeopardize the network functionality, opening new possibilities and objectives for the study of network architectures.

Chapter 5

PROPRIOCEPTION, SUSTAINED RESPONSES, AND NETWORK STRUCTURE

Using a computational model of the *Caenorhabditis elegans* connectome dynamics, we show that proprioceptive feedback is necessary for sustained dynamic responses to external input. This is consistent with the lack of biophysical evidence for a central pattern generator, and recent experimental evidence that proprioception drives locomotion. The low-dimensional functional response of the *Caenorhabditis elegans* network of neurons to proprioception-like feedback is optimized by input of specific spatial wavelengths which correspond to the spatial scale of real body shape dynamics. To explore how the connectomic dynamics produces the observed two-mode, oscillatory limit cycle behavior from a static fixed point, we probe the fixed point's low-dimensional structure using Dynamic Mode Decomposition. This reveals that the nonlinear network dynamics encode six clusters of dynamic modes, with timescales spanning three orders of magnitude. Two of these six dynamic mode clusters correspond to previously-discovered behavioral modes related to locomotion. These dynamic modes and their timescales are encoded by the network's degree distribution and specific connectivity. This suggests that behavioral dynamics are partially encoded within the Connectome itself, the connectivity of which facilitates proprioceptive control.

5.1 Broader Context

The nematode *C. elegans* lives a complex and rich life despite having only 302 neurons. The full connectivity between these neurons (its Connectome) has been measured, making it an ideal model system for understanding how neural processing generates behavior. However, unlike most animals, it doesnt appear to have neural circuits dedicated to generating

rhythmic motion. Even the simple rhythmic behavior of forward locomotion is known to be ultimately driven by stretch-receptive proprioception as the worm changes body shape.

Computational models of Connectome dynamics can capture behavioral responses, making them ideal for studying neural dynamics. We consider a model of full-Connectome dynamics which, without external input, is in a static neural state. We show that proprioception-like input can bring the system into a locomotion-like rhythmic state. We then show that arbitrary external inputs across all neurons can only excite some combination of six neural patterns, and that the shape and dynamics of these patterns are determined by the Connectome. Furthermore, two of these patterns correspond to forward locomotion, such that the Connectome facilitates locomotion-like responses. We thus find that the Connectome itself is not random, but is ideally tuned for generating behavior.

5.2 Introduction

The exact process through which the nematode *Caenorhabditis elegans* (*C. elegans*) generates the rhythmic activity necessary for locomotion remains unclear [48]. In many other species, a Central Pattern Generator (CPG) is typically the source of rhythmic activity [13, 36, 80, 86, 79]. There is insufficient experimental evidence to support the existence of a CPG in the *C. elegans* neuronal network [87, 30]. Experimental and computational evidence shows that proprioception within motorneurons plays an important role in driving and modulating forward locomotion [130, 10], and it has been hypothesized that this proprioceptive feedback is what ultimately generates rhythmic locomotion [10], rather than any dedicated circuitry in the neuronal network. Using a computational model for the connectome dynamics of *C. elegans* [68], we provide strong theoretical and computational support, through the emerging method of dynamic mode decomposition, for the hypothesis that proprioception within motorneurons does indeed encode and drive rhythmic activity.

Critical to assessing how sustained, low-dimensional dynamic activity is generated, is understanding the role the network's connectivity graph (its "Connectome") plays in generating rhythmic motion. The structure of a neuronal network's connectivity often determines how

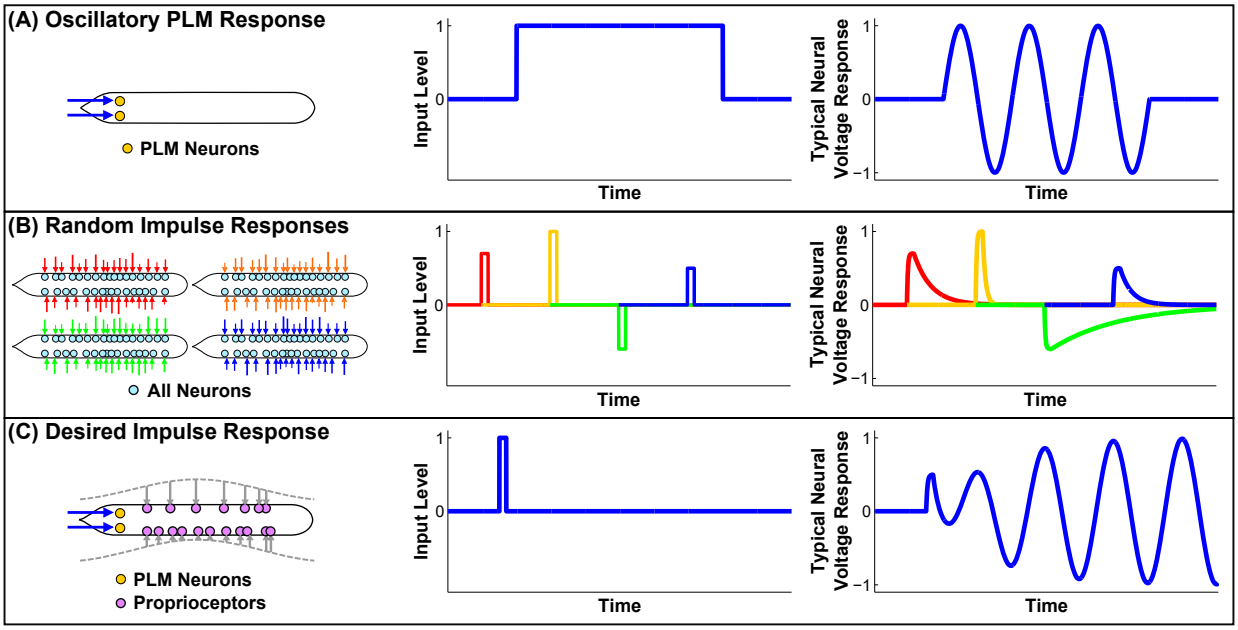

Figure 5.1: **A.** Illustration of the oscillatory response as demonstrated in [68]. Unrealistically, the system requires constant stimulation or it will collapse into a fixed point. This is consistant with evidence that proprioceptive feedback is necessary for sustained dynamic responses to external input. **B.** Illustration of the response to momentary random stimuli. After any stimulus, the system will decay back to the fixed point, albeit at different timescales. We will use this to probe the dynamical structure of the fixed point. **C.** Illustration of a potentially more realistic response. We seek a mechanism for proprioceptive feedback which produces sustained responses to momentary stimuli. We will investigate the consistency of our model with such a framework.

the network operates as a whole [129, 89], encoding key behavioral responses characterized by low-dimensional patterns of activity [12, 23, 29, 74, 71, 32]. However, the exact importance of the *specific* connectivity of a network is unclear, and neuronal network dynamics are often computationally modeled using uniform random networks[121, 4, 16, 124, 37, 84, 58, 96]. In *C. elegans*, however, the structure of the Connectome is clearly not random, and it may further play a critical role in helping to generate or facilitate rhythmic responses. This is suggested by the fact that computational models of the Connectome can generate motorneuron oscillations related to forward locomotion in response to constant stimuli even without proprioception (and even when modeling neural dynamics alone, with no coupled muscular, bodily or environmental modeling) [68]. This suggests that oscillatory, stereotyped responses are, at some level, encoded within the Connectome.

There is, however, an important caveat to this result: oscillatory output occurs only due to an unrealistic stimulus, consisting of a constant input into the tail-touch mechanoreceptor pair PLM [25]. In the absence of constant stimulus, the neural state will collapse onto a static, stable fixed point, i.e. a state of no movement. This is illustrated in Panel (A) of Figure 5.1. This is clearly not realistic; the actual worm is not constantly receiving tail-touch stimulus during every moment at which it crawls forward. As illustrated in Panel (B), the system will quickly decay back to static equilibrium after *any* random stimulus. A more realistic response to an impulse may perhaps look more like Panel (C): if the worm is in a pause state, a momentary stimulus should be capable of driving it into sustained motion. This lack of sustained oscillation can be explained by the model's lack of feedback, specifically by the lack of stretch-receptive proprioception within B-class motorneurons, which is known to drive and regulate locomotion [130].

This study thus considers the following questions about the model in [68]: Is it consistent with a framework of proprioception-driven locomotion? If so, do the low-dimensional output patterns encoded by the Connectome facilitate proprioceptive control? In other words, does the system's equilibrium have a low-dimensional dynamical structure which facilitates responses related to locomotion?

In this Chapter, we demonstrate that proprioceptive feedback is indeed necessary and sufficient for sustained dynamic responses to external input. This is consistent with the lack of biophysical evidence for a central pattern generator driving locomotion, and the evidence that proprioception drives locomotion. Explicitly, we use the spatial location of specific motorneurons to drive them with a sinusoidal traveling wave, approximating strech-receptive proprioception during locomotion. The functional response of the network to this proprioception-like input is optimized by specific spatial wavelengths, specifically optimal locomotion responses are driven by input with spatial scales consistent with *C. elegans* body shape dynamics, i.e. eigenworm-like structures [110]. By applying Dynamic Mode Decomposition to the network data, we discover that the dynamics encode six clusters of dynamic modes with timescales spanning three orders of magnitude. Two of these six dynamic mode clusters correspond to previously-discovered behavioral modes related to locomotion. The dynamic modes and timescales are encoded by the network's degree distribution and specific connectivity. This suggests that behavioral dynamics are partially encoded within the Connectome itself, the connectivity of which facilitates proprioceptive control. Thus our results suggest a framework in which the neural network is not the *source* of spontaneous oscillation, but rather is structured to facilitate specific proprioception-driven oscillation responses. More broadly, our application of Dynamic Mode Decomposition to network dynamics demonstrates its utility at discovering, from activity data alone, the responses which a network may be encoded to promote or inhibit.

5.3 Results

5.3.1 Perturbation Response and Dynamic Modes

Given the lack of evidence for a CPG within the network, it is interesting that the system is able to generate oscillation in response to a non-oscillatory input, and that this oscillation appears related to locomotion. However, it is clearly unrealistic that such oscillation would require a constant, explicit external input, and would otherwise collapse to a fixed point

(i.e. a static neural pattern). The dynamical structure of this fixed point, from which we wish to drive the system into sustained oscillatory motion, can be investigated through impulse-response experiments.

In each of several trials, we model the dynamics of the full somatic nervous system of 279 neurons (where there are 302 neurons total, 282 within the somatic nervous system, and 279 of those which make synaptic connections [122]). We perturbed the system from equilibrium with a short stimulus distributed randomly across all 279 neurons. The system was then allowed to freely decay back to the fixed point, and the decaying neuron voltages were recorded (providing data as shown in Figure 5.2). We observed that, in all trials, the system decayed back to the same fixed point regardless of input stimulus.

We find that these dynamics are well-described by a few modes (i.e. specific neural voltage patterns), each of which decay exponentially bringing the system back to the fixed point. Applying Dynamic Mode Decomposition to the data gives us both these modes and their decay time constants. Interestingly, we find the following: (1) in all trials the dynamics are well-described by only six modes, (2) DMD gives approximately the same six modes regardless of the random stimulus direction, and (3) the time constants of the modes are well-separated and span three orders of magnitude. The distributions of decay constants for each of these modes is plotted in Figure 5.4.

These modes can be interpreted as the components of a low-dimensional manifold to which the dynamics are constrained around the fixed point. In other words, an arbitrary stimulus into all 279 neurons can effectively only excite some combination of these six neural patterns. This is what we mean by the fixed point having "low-dimensional structure".

5.3.2 Relation of Dynamic Modes to Forward Motion

How does this low-dimensional structure relate to the previously-observed, locomotion-like oscillatory response? To answer this, we note that the PLM response in [68] is characterized by three modes (the "PLM modes"): (1,2) the two modes defining the plane in which the limit cycle proceeds (the "PLM plane"), which we call the "plane modes", and (3) the

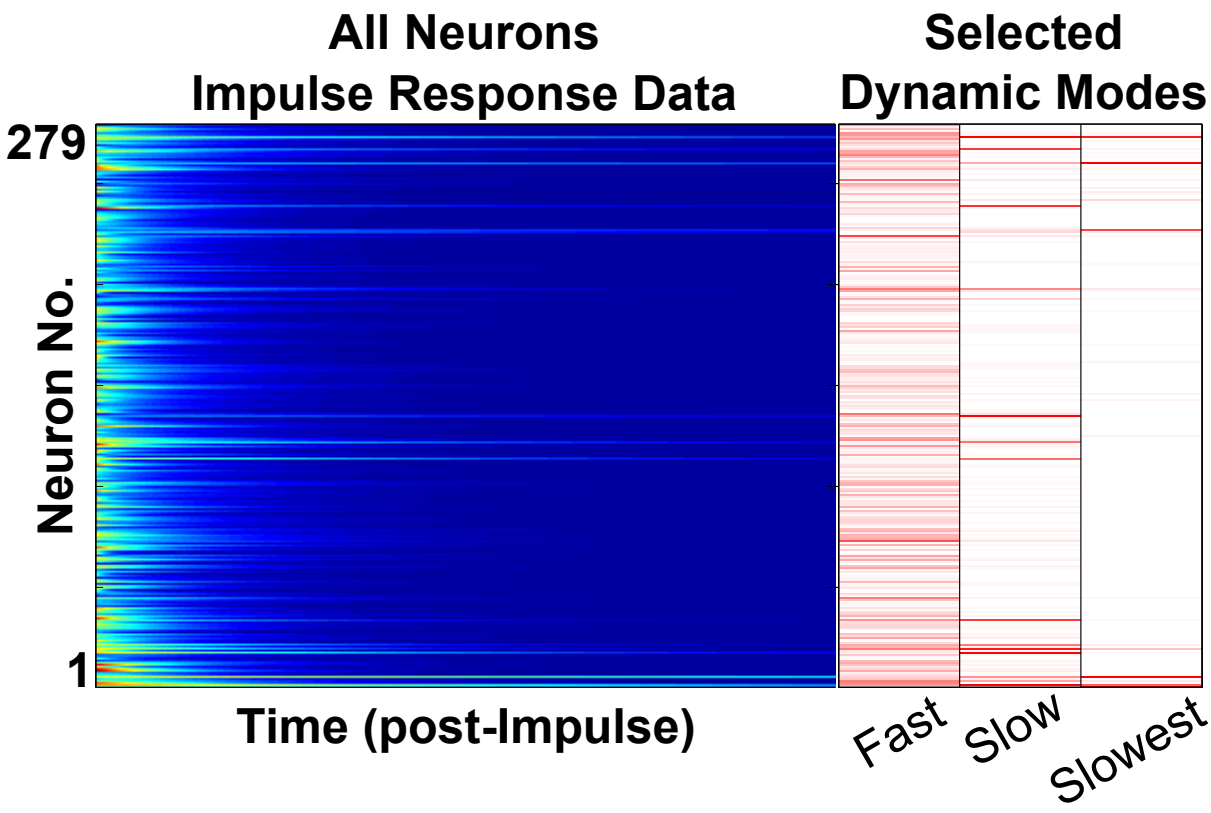

Figure 5.2: The raster plot at left plots shows a single trial of neuron voltage responses to a random impulse. The nonlinear network dynamics encode six clusters of dynamic modes with timescales spanning three orders of magnitude. Shown at right are slowest of these dynamic modes (later referred to as Modes 4, 5 and 6). One can see the modal dynamics within the raster plot (e.g. the raster plot resembles the "Slowest" mode near the end of the trial).

displacement between the equilibrium fixed point and the center of the limit cycle, which we call the "displacement mode". These modes and their relation to the fixed point and limit cycle are depicted in the "Phase-Plane Dynamics" illustration of Figure 5.3.

We investigate the biological meaning of our dynamic modes by calculating their projections onto the PLM modes. This reveals that two of the six dynamic modes correspond to previously-discovered PLM modes. Figure 5.4 shows the magnitude of each dynamic mode's projection onto the displacement mode and PLM plane. We compare these against the projections of 1,000 random modes. Both the displacement mode and the PLM Plane have a single dynamic mode projecting strongly onto them. Therefore the low-dimensional structure of the fixed point facilitates responses in both the displacement mode and PLM Plane directions, with a highly distinct timescale for each response.

Mode 4, which projects strongly onto the displacement mode, is particularly interesting. It has the most consistent timescale between trials (see Table 5.1). Additionally, all other dynamic modes have a particularly *low* projection onto the displacement mode (i.e. a significantly lower median projection than random modes). This suggests that the low-dimensional structure of the fixed point facilitates responses in that direction with a particularly consistent timescale. Thus periodic perturbations of the correct timescale could perturb the system off of the fixed point in the direction of the limit cycle; we discuss these implications further in the Discussion.

5.3.3 Dynamic Modes Result from Connectivity

The role of the Connectome, the experimentally validated network connectivity, was investigated by repeating the perturbation experiments but with randomly changed network connectivities. We considered the following variations: (A) a network with the same degree distribution, with node degrees and connections randomly assigned; (B) random connectivity with the same total number of edges.

Results from these cases are summarized in Figure 5.5. In (A), where the degree distribution is maintained, there are still six modes. However, timescales vary somewhat from the

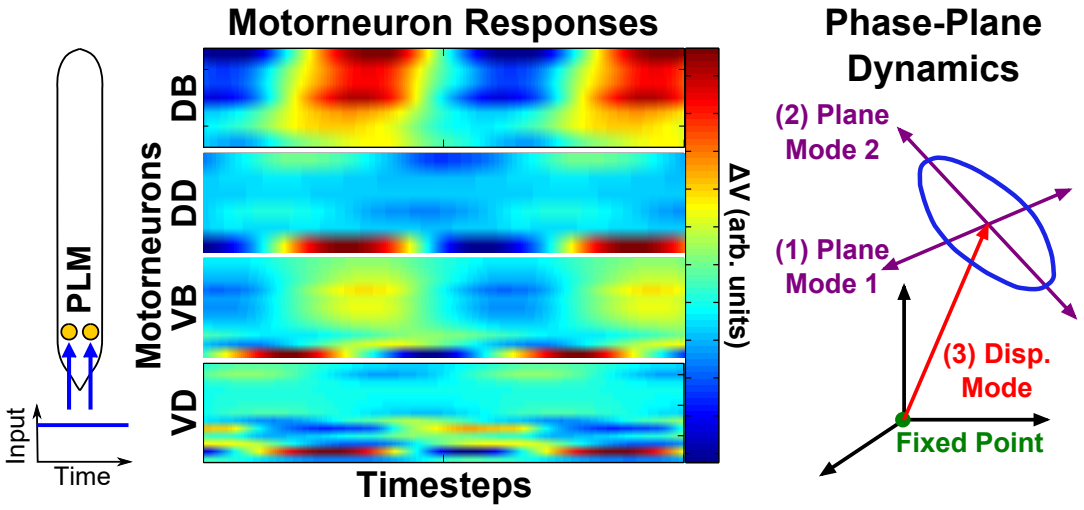

Figure 5.3: Illustration of the PLM Response as in [68]. Constant stimulation of PLM neurons (corresponding to tail-touch) causes oscillation in body-wall motorneurons. This oscillation can be described as a 2D limit cycle, consistant with the observed 2D body shape dynamics of forward motion [110]. The center of this limit cycle is displaced in the full-dimensional space from the zero-input fixed point. We refer to the two oscillatory modes as the "plane modes", and the displacement from the fixed point to the center of oscillation as the "displacement mode".

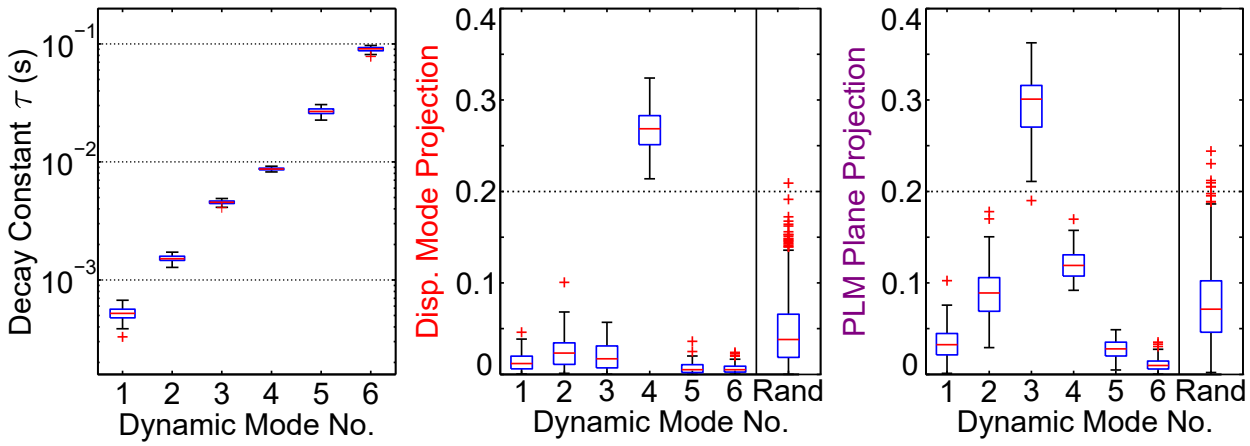

Figure 5.4: Two of the six dynamic mode clusters correspond to previously-discovered behavioral modes related to locomotion. DMD results from 100 random impulse trials are plotted as box-and-whisker plots showing each mode's decay constant, as well as projections onto the displacement mode and PLM plane. Projections from 1000 random modes are compared. Timescales of each mode are highly distinct and consistant, and vary over three overs of magnitude. Mode 4 and Mode 3 project strongly onto the Displacement Mode and PLM Plane, respectively. See Tables 5.1, 5.2 and 5.3 for corresponding numerical values.

original, and the projections are completely changed. In (B), which has a different degree distribution, the number of modes and all of their properties are qualitatively different. This establishes that the dynamic modes and timescales are encoded by both the network's degree distribution and specific connectivity.

5.3.4 Driven Oscillatory Response

We have established that random impulses can drive the system in the direction of the displacement mode. However, given the apparent global stability of the fixed point, an additional mechanism is required for sustained dynamic responses to external input. Proprioception may allow initial perturbations to grow into the desired limit cycle associated, for instance, with forward motion.

We thus investigated the following question: could stretch-receptive proprioceptive feedback within B-class motorneurons give rise to motorneuron oscillations which are qualitatively similar to PLM-driven oscillations? Since locomotion consists approximately of sinusoidal bends propagating along the body [52, 64], it suffices to drive B-class motorneurons with sinusoidal inputs, as determined by their location along the axis of the body and their position on the dorsal/ventral side (noting that the *C. elegans* lays on its side while it crawls [64]).

Figure 5.6 shows the motorneuron dynamics resulting from different sinusoidal inputs into B-class motorneurons. For certain spatial wavelengths, a limit cycle does occur which is qualitatively similar on the PLM plane. Note that there is a smooth transition between the rows of Figure 5.6, and that the middle row (for $k = 0.886$) is the most similar to the PLM-driven cycle. Temporal frequency did not have an effect on the shape of the limit cycle, consistent with the experimental observations that the spatial wavelength of *C. elegans* locomotion does not depend on temporal frequency [64]. As discussed further in the Discussion, this suggests that our model, if integrated with a mechanical body model in future work, could be made consistent with a system of feedback-driven oscillations.

5.4 Discussion

In this Chapter, we have introduced the Dynamic Mode Decomposition as a diagnostic tool to characterize impulse-response experiments on a nonlinear networked system. This revealed that the network is structured to generate a low-dimensional response at distinct timescales ranging over several orders of magnitude, and that two of these dynamic modes are related to the previously-characterized "forward motion" response to PLM-stimulation. It is possible that proprioceptive feedback could *sustain* a limit cycle but not be sufficient to bring the system to said limit cycle from the equilibrium fixed point. In other words, the limit cycle would need to be "jumpstarted", with a separate mechanism transporting the system from the fixed point near to the cycle. Were this the case, it would suggest a physiological purpose for the low-dimensional fixed-point structure which we detect: stimuli of the correct timescales could selectively perturb the system towards the limit cycle, to a point from which the proprioceptive feedback could be effective. In this view, it is interesting and suggestive that Mode 4, associated with the Displacement Mode, has the most tightly-constrained timescales of all the modes.

Repeating this analysis for different connectivities suggested that these dynamic modes and timescales are encoded by both the network's degree distribution and specific connectivity. A random graph, with the same number of nodes and connections but a different degree distribution, leads to a completely different number of modes. This suggests that the number of dynamical timescales is encoded by the degree distribution, as six timescales are recovered for any network with the same degree distribution. However, the specific timescale values and the neuronal makeup of these modes is not preserved. The degree to which each mode projects onto our biophysiologically-relevant directions, and with what specific dynamical timescale, depends on the specific wiring of the Connectome. Thus behavioral dynamics are partially encoded within the Connectome itself, the connectivity of which facilitates proprioceptive control. Said another way, the stereotyped worm Connectome seems to be optimized for its behavioral repertoire.

The usefulness of these insights as they apply to the actual system, however, depend on the model's compatibility with a framework of proprioception-generated oscillation. Thus we further show that sinusoidal input into the putatively proprioceptive B-class motorneurons does, indeed, drive a limit cycle at certain spatial wavelengths. Given that the worm crawls with a sinusoidal body shape [52, 64], this suggests that motorneuron proprioception could indeed drive the limit cycle, which in turn could drive sinusoidal movement. A proprioceptive mechanism such as this is necessary for sustained dynamic responses to external input. Despite this apparent consistency, the development of such a feedback rule remains nontrivial. Without a coupled biomechanical model that includes muscle activation, any feedback rule which we might implement on the present model would be no less artificial than our direct sinusoidal stimulus, which is biophysiologically reasonable.

However, modeling the worm's body and environment is ultimately crucial to fully understanding its behavior [87, 8, 41, 115, 104, 10]. This study prescribes multiple studies for future computational Connectome models which *are* fully integrated with bodily and environmental models (as exemplified by projects such as OpenWorm [114]). Specifically, it introduces the following questions: (1) When motorneuron proprioception and other external feedback is turned off within a model, does the system decay into a fixed point? If so, an identical study can be performed to probe that fixed point's low-dimensional structure. (2) Do the dynamic modes relate to the oscillatory dynamics which occur during locomotion? (3) If proprioception/feedback is turned back on while the system is in its fixed point, does the system proceed into a spontaneous limit cycle, and if so, how? Is periodic noise or other stimulation of a specific timescale necessary for such a transition?

More broadly, this work demonstrates the utility of Dynamic Mode Decomposition in relating the specific connectivity of a network to the multi-scale, low-dimensional structure of its dynamical responses. The methods of this Chapter are able to directly relate connectivity to dynamics even for large, nonlinear networked systems. Future work will further investigate this relationship, with implications for the design of nonlinear networks.

5.5 Methods

5.5.1 Modeling the C. elegans Connectome

Simulations were performed in MATLAB via Euler's method, using timesteps of $h = 10^{-6}$s. The data was downsampled by recording $v(t)$ every $\Delta t = 3 \times 10^{-5}$s, yielding a data matrix:

$$\mathbf{V} = \begin{bmatrix} | & | & & | \\ \mathbf{v}(t_1) & \mathbf{v}(t_2) & \cdots & \mathbf{v}(t_{m-1}) \\ | & | & & | \end{bmatrix}, \tag{5.1}$$

where $t_{k+1} - t_k = \Delta t$. The value of Δt was chosen to be sufficiently low so as to not affect the outcome of the analysis.

5.5.2 Random Perturbations

For each random perturbation simulation, a random external input I^{Ext} was applied to all neurons for a duration of 10^{-5}s, after which the system was allowed to decay. Output was recorded from all neurons after the cessation of input. Each I_i^{Ext} was drawn from a Gaussian distribution, after which the total I^{Ext} was then normalized to have a fixed input amplitude.

5.5.3 Dynamic Mode Decomposition

This section describes the method of Dynamic Mode Decomposition[101, 99, 120, 27, 70, 82, 18, 90], which we apply to our simulated neural voltage data $\mathbf{V}$. Specifically, we use it to relate the voltages at timestep t_k to the following timestep t_{k+1} as follows:

$$\mathbf{v}(t_{k+1}) \approx \mathbf{A}\mathbf{v}(t_k), \tag{5.2}$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the linear operator which is the best-fit solution for all pairs. Note that this does not imply that the underlying dynamics are linear; DMD is connected to nonlinear dynamical systems through the Koopman operator[82]. We can express this relationship in

matrix form by constructing two data matrices $\mathbf{X} \in \mathbb{R}^{n \times (m-1)}$ and $\mathbf{X'} \in \mathbb{R}^{n \times (m-1)}$ as follows:

$$\mathbf{X} = \begin{bmatrix} | & | & & | \\ \mathbf{v}(t_1) & \mathbf{v}(t_2) & \cdots & \mathbf{v}(t_{m-1}) \\ | & | & & | \end{bmatrix}, \tag{5.3}$$

$$\mathbf{X'} = \begin{bmatrix} | & | & & | \\ \mathbf{v}(t_2) & \mathbf{v}(t_3) & \cdots & \mathbf{v}(t_m) \\ | & | & & | \end{bmatrix}. \tag{5.4}$$

This allows us to write Equation (5.2) as:

$$\mathbf{X'} \approx \mathbf{AX}. \tag{5.5}$$

The dynamic mode decomposition of the data matrices $(\mathbf{X}, \mathbf{X'})$ is given by the leading eigendecomposition of the matrix $\mathbf{A}$, which is defined as follows:

$$\mathbf{A} = \mathbf{X'X}^{\dagger}, \tag{5.6}$$

where $\dagger$ denotes the Moore-Penrose pseudoinverse[120]. The pseudoinverse of $\mathbf{X}$ can be found by calculating its singular value decomposition, truncated at r singular values:

$$\mathbf{X} \approx \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^{*}. \tag{5.7}$$

Here * denotes the complex conjugate transpose, $\tilde{\mathbf{U}} \in \mathbb{R}^{n \times r}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{m-1 \times r}$ are matrices with orthonormal columns, and $\tilde{\mathbf{\Sigma}} \in \mathbb{R}^{r \times r}$ is diagonal. The diagonal entries of $\mathbf{\Sigma}$ are the singular values, and are proportional to the percentage of energy within each mode. We choose the smallest set of r modes which capture 99% of the energy.

We can thus approximate the linear operator $\mathbf{A}$ as follows:

$$\mathbf{A} \approx \bar{\mathbf{A}} = \mathbf{X'}\tilde{\mathbf{V}}\tilde{\mathbf{\Sigma}}^{-1}\tilde{\mathbf{U}}^{*}. \tag{5.8}$$

We are interested in the dynamics projected upon the lower-dimensional subspace as defined by the first r columns of $\tilde{\mathbf{U}}$. Rather than calculating the $n \times n$ matrix $\bar{\mathbf{A}}$, we project onto the low-dimensional subspace to calculate the $r \times r$ reduced order operator $\tilde{\mathbf{A}}$:

$$\tilde{\mathbf{A}} = \tilde{\mathbf{U}}^{*}\mathbf{X'}\tilde{\mathbf{V}}\tilde{\mathbf{\Sigma}}^{-1}. \tag{5.9}$$

The eigendecomposition $\tilde{\mathbf{A}}\mathbf{W} = \mathbf{W}\mathbf{\Lambda}$ gives the eigenvectors $\mathbf{w}_j$ and eigenvalues λ_j of the reduced-order system. The eigenvalues are equal to those of the full-dimensional $\bar{\mathbf{A}}$, and the corresponding eigenvectors can be used to exactly calculate the full-dimensional dynamic modes of the system[120]. For $\lambda_j \neq 0$, the dynamic mode corresponding to $\mathbf{w}_j$ is:

$$\phi = \mathbf{X}'\tilde{\mathbf{V}}\tilde{\mathbf{\Sigma}}^{-1}\mathbf{w}. \tag{5.10}$$

Since these dynamic modes correspond to the eigenvectors of the low-dimensional system, the modes give the dynamically-decoupled low-dimensional patterns which will exponentially growth/decay and/or oscillate with timescales given by their respective eigenvalues λ_j. We are primarily interested in analyzing the exponential decay of our system onto a fixed point. Thus we calculate each mode's exponential decay constant τ_j:

$$\tau_j = \frac{-\Delta t}{\ln\left(\mathrm{Re}(\lambda_j)\right)}, \tag{5.11}$$

where $\Delta t = t_{k+1} - t_k$ is the timestep between datapoints.

5.5.4 Dynamic Mode Properties

The properties of the resultant modes are summarized in the boxplots of Figures 5.4 and 5.5. These were generated from MATLAB function `boxplot.m`. Default settings are used in Figure 5.4, and in Figure 5.5 the settings are changed such that no points are treated are outliers. Projections are calculated by taking the dot product of each trial's DMD modes with the PLM modes (calculated by taking the singular value decomposition of the PLM-driven limit cycle, as in [68]). The random projections in Figure 5.4 are calculated by taking the dot product with a random mode, in which each element is chosen from a Gaussian distribution and the mode is then normalized.

5.5.5 Sinusoidal Driving of B-Class Motorneurons

In approximating proprioceptive input, we sinusoidally drove all B-class motorneurons, using an external input of the following form:

τ	Upper	75th Pct.	Median	25th Pct.	Lower	Spread
Mode 1	0.00067	0.00057	0.00052	0.00048	0.00039	27.6 %
Mode 2	0.00172	0.00159	0.00151	0.00146	0.00128	14.6 %
Mode 3	0.00489	0.00466	0.00455	0.00445	0.00412	8.5 %
Mode 4	0.00920	0.00886	0.00872	0.00856	0.00822	5.6 %
Mode 5	0.03065	0.02816	0.02681	0.02568	0.02264	14.9 %
Mode 6	0.09705	0.09324	0.09094	0.08736	0.08126	8.7 %

Table 5.1: Numerical values of the timescales in Figure 5.4. The consistency of timescales resulting from different trials is given by the "spread", calculated as $(Upper - Lower)/(2 \times Median)$.

$$I_i^{Ext} = \pm A \sin(\omega t - kx) \tag{5.12}$$

Input sign was given based on the dorsal/ventral location of the motorneuron. Input amplitude A affected only the amplitude of the cycle and was set at $A = 30$ Arb. Units to yield a qualitatively similar cycle amplitude. Temporal frequency ω appeared to affect the response only by changing the cycle period. Spatial wavelength k varied between trials (as shown in Figure 5.6). x was assigned to each neuron based on its soma position.

Soma position data originates from [122], and was retrieved from the "Neuronal Wiring" section of WormAtlas [3]. The use of the soma position is a simplification: proprioception in B-class motorneurons is believed to be due to stretch reception within the long axons posterior to the soma [130].

Disp.	Upper	75th Pct.	Median	25th Pct.	Lower
Mode 1	0.03858	0.01996	0.01178	0.00597	0.00043
Mode 2	0.06829	0.03442	0.02325	0.01089	0.00120
Mode 3	0.05698	0.03109	0.01701	0.00717	0.00021
Mode 4	0.32392	0.28287	0.26850	0.25110	0.21387
Mode 5	0.02002	0.01063	0.00514	0.00204	0.00013
Mode 6	0.01659	0.00891	0.00507	0.00248	0.00017
Random	0.13578	0.06584	0.03813	0.01850	0.00012

Table 5.2: Numerical values of the Displacement Mode Projections in Figure 5.4.

Plane	Upper	75th Pct.	Median	25th Pct.	Lower
Mode 1	0.07575	0.04464	0.03250	0.02130	0.00117
Mode 2	0.15038	0.10588	0.08901	0.06892	0.02930
Mode 3	0.36243	0.31583	0.30094	0.27032	0.21088
Mode 4	0.15740	0.13069	0.11906	0.10745	0.09184
Mode 5	0.04889	0.03529	0.02791	0.02014	0.00476
Mode 6	0.02739	0.01454	0.00976	0.00588	0.00021
Random	0.18629	0.10235	0.07133	0.04601	0.00214

Table 5.3: Numerical values of the PLM Plane Projection in Figure 5.4.

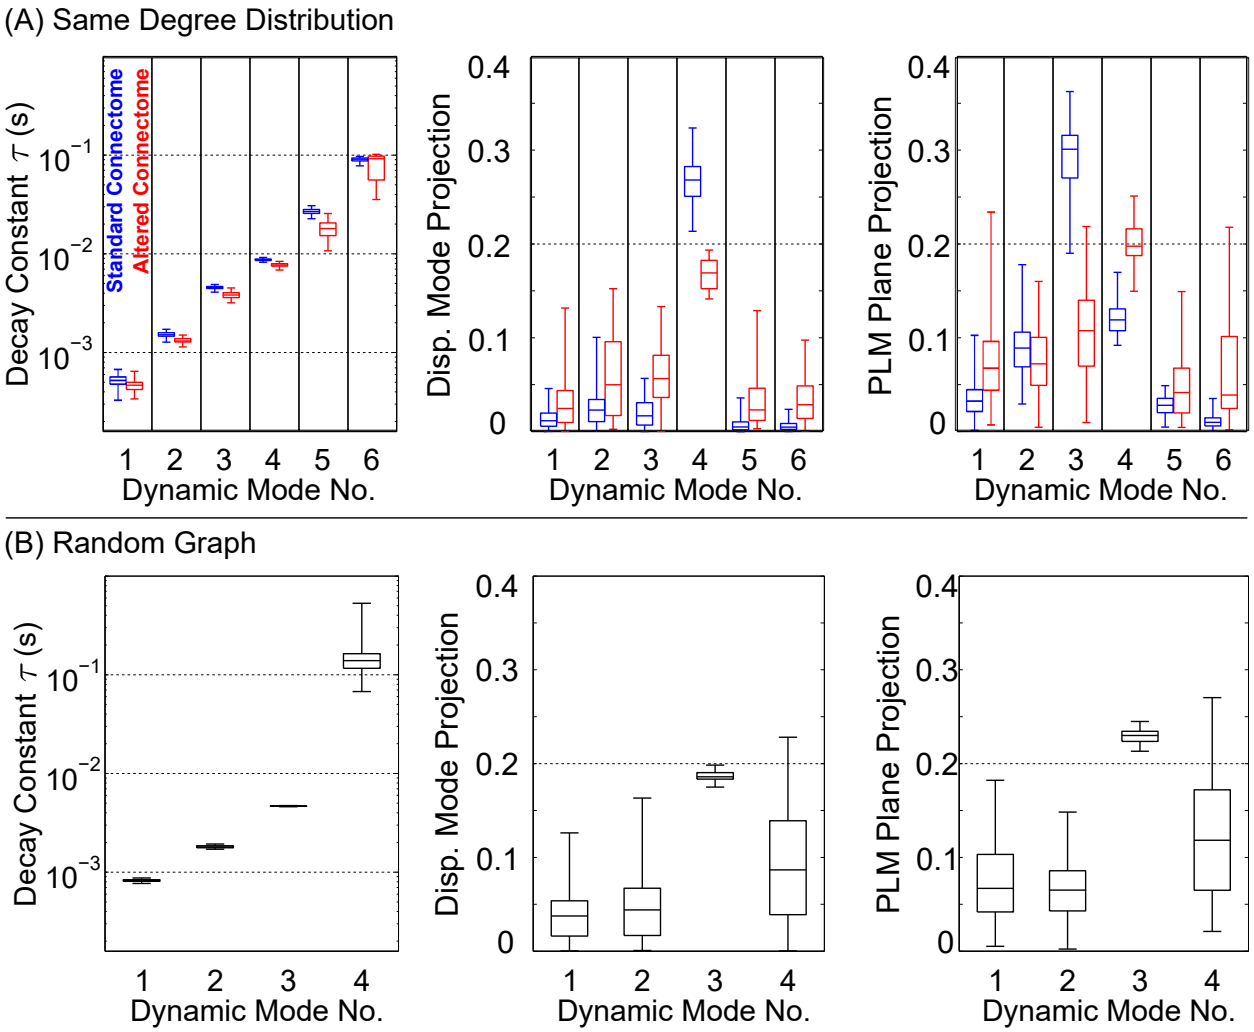

Figure 5.5: Analysis was repeated for different network connectivities. Each new boxplot takes data from 10 trials of 5 different randomly-generated networks: **(A)** Results from networks generated to have an identical degree distribution. Timescales are similar but projection values are qualitatively dissimilar. **(B)** Results from fully random networks. Changing the degree distribution changes the number of modes and their timescales. This shows that the dynamic modes and timescales are encoded by the network's degree distribution and specific connectivity. This suggests that behavioral dynamics are partially encoded within the Connectome itself, the connectivity of which facilitates proprioceptive control.

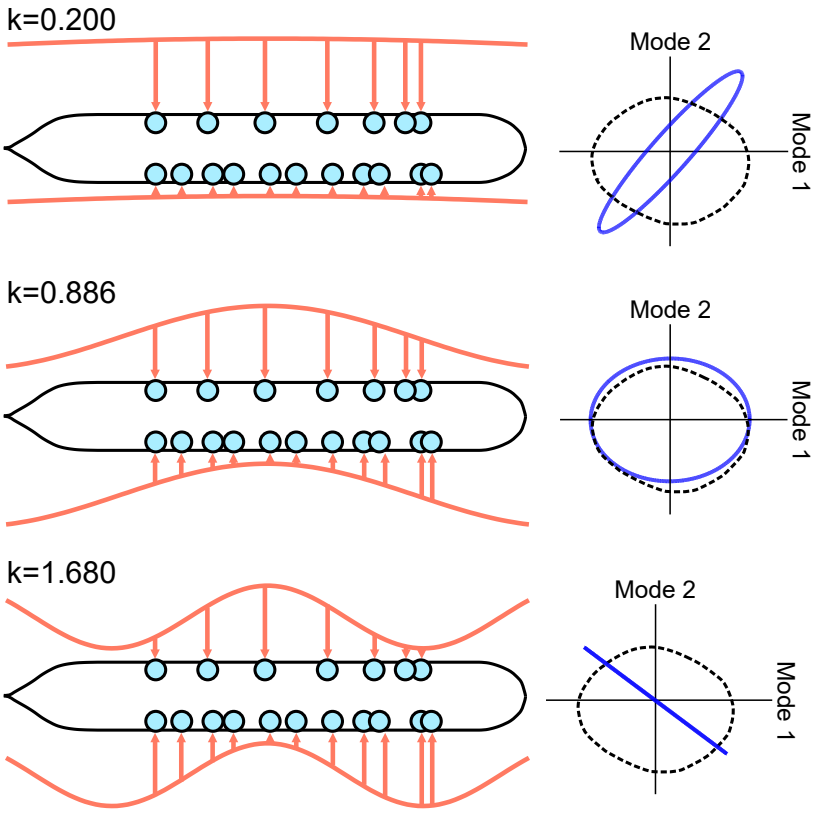

Figure 5.6: Resulting dynamics from sinusoidally driving B-class motorneurons, approximating proprioception. Dynamics are shown projected onto the PLM plane and plotted as solid blue lines (with the black dashed lines showing the PLM response cycle). Certain spatial wavelengths give rise to qualitatively similar limit cycles.

BIBLIOGRAPHY

[1] H Adle-Biassette, F Chretien, L Wingertsmann, C Hery, T Ereau, F Scaravilli, M Tardieu, and M Gray. Neuronal apoptosis does not correlate with dementia in hiv infection but is related to microglial activation and axonal damage. *Neuropathology and Applied Neurobiology*, 25:123–133, 1999.

[2] DG Albertson and JN Thomson. The pharynx of caenorhabditis elegans. *Phil. Trans. R. Soc. Lond.*, 275:299–325, 1976.

[3] Z.F. Altun, L.A. Herndon, C. Crocker, and D. H. Hall. Wormatlas. http://www.wormatlas.org, 2002-2016.

[4] D Amit and N Brunel. Dynamics of a recurrent network of spiking neurons before and following learning. *Network*, 8:373404, 1997.

[5] DJ Amit and N. Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7:237–252, 1997.

[6] NB Angstman, MC Kiessling, H-G Frank, and C. Schmitz. High interindividual variability in dose-dependent reduction in speed of movement after exposing *C. elegans* to shock waves. *Front. Behav. Neurosci.*, 9:12, 2015.

[7] DJ Bakkum, U Frey, M Radivojevic, TL Russel, M J, M Fiscella, H Takahashi, and A Hierlemann. Tracking axonal action potential propagation on a high density micro electrode array across hundreds of sites. *Nature communications*, 4:2181, 2013.

[8] S Berri, JH Boyle, M Tassieri, IA Hope, and N Cohen. Forward locomotion of the nematode c. elegans is achieved through modulation of a single gait. *HFSP J.*, 3:186193, 2009.

[9] PC Blumbergs, G Scott, J Manavis, H Wainwright, DA Simpson, and AJ. McLean. Topography of axonal injury as defined by amyloid precursor protein and the sector scoring method in mild and severe closed head injury. *Journal of Neurotrauma*, 12:565–572, 1995.

[10] JH Boyle, S Berri, and N. Cohen. Gait modulation in *C. elegans*: An integrated neuromechanical model. *Frontiers in Computational Neuroscience*, 6:10, 2012.

[11] S. Brenner. The genetics of behaviour. *British Medical Bulletin*, 29:269–271, 1971.

[12] B Broome, V Jayaraman, and G. Laurent. Encoding and decoding of overlapping odor sequences. *Neuron*, 51(4):467–482, 2006.

[13] TG Brown. The intrinsic factors in the act of progression in the mammal. *Proc. R. Soc. Lond. B Biol. Sci.*

[14] KD Browne, XH Chen, DF Meaney, and DH. Smith. Mild traumatic brain injury and diffuse axonal injury in swine. *Journal of Neurotrauma*, 28(9):1747–1755, 2011.

[15] N. Brunel. Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. *Journal of Physiology*, 94:445–463, 2000.

[16] N Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*, 8:183208, 2000.

[17] BW Brunton, LA Johnson, JG Ojemann, and JN Kutz. Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *J. Neurosci. Methods*, 258:1–15, 2015.

[18] BW Brunton, LA Johnson, JG Ojemann, and JN Kutz. Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *Journal of Neuroscience Methods*, 258:1–15, 2016.

[19] J. Bryden and N. Cohen. Neural control of caenorhabditis elegans forward locomotion: the role of sensory feedback. *Biol. Cybern.*, 98(4):339–51, 2008.

[20] D Bucher and JM. Goaillard. Beyond faithful conduction: short term dynamics, neuromodulation, and log-term regulation of spike propagation in the axon. *Progress in Neurobiology.*, 94:307–346, 2011.

[21] SN Burke and CA. Barnes. Neural plasticity in the aging brain. *Nature Reviews Neuroscience.*, 7:30–40, 2006.

[22] L Byerly, RC Cassada, and RL Russell. The life cycle of the nematode *Caenorhabditis elegans*: I. wild-type growth and reproduction. *Developmental Biology*, 51:23–33, 1976.

[23] M Byron, J Cunningham, G Santhanam, S Ryu, K Shenoy, and M. Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Adv. Neural Inform Process Syst*, pages 1881–1888, 2009.

[24] S. H. Chalasani, N. Chronis, M. Tsunozaki, J. M. Gray, D. Ramot, M. B. Goodman, and C. I. Bargmann. Dissecting a circuit for olfactory behaviour in caenorhabditis elegans. *Nature*, 450:63–70, 2007.

[25] M. Chalfie, J. E. Sulston, J. G. White, E. Southgate, J. N. Thomson, and S. Brenner. The neural circuit for touch sensitivity in caenorhabditis elegans. *J. Neurosci.*, 5(4):956–964, 1985.

[26] B. L. Chen, D. H. Hall, and D. B. Chlovskii. Wiring optimization can relate neuronal structure and function. *Proc. Natl. Acad. Sci.*, 103:4723–4728, 2006.

[27] KK Chen, JH Tu, and CW Rowley. Variants of dynamic mode decomposition: Boundary condition, koopman, and fourier analyses. *J. Nonlin. Sci.*, 22:887–915, 2012.

[28] CW Christman, MS Grady, SA Walker, KL Hol-Loway, and JT. Povlishock. Ultrastructural studies of diffuse axonal injury in humans. *Journal of Neurotrauma*, 11:173–186, 1994.

[29] M Churchland, J Cunningham, M Kaufman, J Foster, P Nuyujukian, S Ryu, and K. Shenoy. Neural population dynamics during reaching. *Nature*, 487:51–56, 2012.

[30] N. Cohen and T. Sanders. Nematode locomotion: dissecting the neuronal-environmental loop. *Curr. Opin. Neurobiol.*, 25:99–106, 2014.

[31] M. Coleman. Axon degeneration mechanisms: commonality amid diversity. *Nature Reviews Neuroscience.*, 6(11):889–898, 2005.

[32] JP Cunningham and BM Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17:1500–1509, 2014.

[33] P Dayan and LF Abbot. *Theoretical neuroscience*. MIT Press.

[34] D. Debanne. Information processing in the axon. *Nature Reviews Neuroscience*, 5(4):304–316, 2004.

[35] D Debanne, E Campanac, A Bialowas, E Carlier, and G. Alcaraz. Axon physiology. *Physiological Reviews*, 91:555–602, 2011.

[36] F Delcomyn. Neural basis of rhythmic behavior in animals. *Science*, 210:492–498, 1980.

[37] A Destexhe and D Contreras. Neuronal computations with stochastic network states. *Science*, 314:85, 2006.

[38] RM Durbin. *Studies on the development and organisation of the nervous system of Caenorhabditis elegans*. PhD thesis, University of Cambridge, 1987.

[39] GB Ermentrout and DH Terman. *Mathematical Foundations of Neuroscience*. Springer.

[40] M Fainaru-Wada and S Fainaru. *League of Denial: The NFL, Concussions, and the Battle for Truth*. Crown Archetype.

[41] C Fang-Yen, M Wyart, J Xie, R Kawai, T Kodger, S Chen, Q Wen, and ADT Samuel. Biomechanical analysis of gait adaptation in the nematode caenorhabditis elegans. *Proc. Natl. Acad. Sci. U.S.A.*, 107:2032320328, 2010.

[42] M Faul, L Xu, MM Wald, and VG Coronado. *Traumatic brain injury in the United States: emergency department visits, hospitalizations, and deaths*. Centers for Disease Control and Prevention, National Center for Injury Prevention and Control, 2010.

[43] S Faumont, TH Lindsay, and SR Lockery. Neuronal microcircuits for decision making in c. elegans. *Curr. Opin. Neurobiol.*, 22(4):580–91, 2012.

[44] B Ferguson, MK Matyszak, MM Esiri, and VH. Perry. Axonal damage in acute multiple sclerosis lesions. *Brain.*, 120:393–399, 1997.

[45] T. C. Ferrée and S. R. Lockery. Computational rules for chemotaxis in the nematode *C. elegans*. *J. Comp. Neurosci.*, 6:263–277, 1999.

[46] JE Galvin, K Uryu, VM Lee, and JQ. Trojanowski. Axon pathology in parkinsons disease and lewy body dementia hippocampus contains α-, β-, and γ -synuclein. *Proceedings of National Academy of Science*, 96:13450–13455, 1999.

[47] W Gerstner, RN Werner, M Kistler, and L Paninski. *Neuronal Dynamics*. Cambridge University Press, 2014.

[48] J Gjorgjieva, D Biron, and G. Haspel. Neurobiology of caenorhabditis elegans locomotion: Where do we stand? *BioScience*, 64(6):476–486, 2014.

[49] J.I. Gold and M.N. Shadlen. Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404:390–394, 1999s.

[50] M. B. Goodman, D. H. Hall, L. Avery, and S. R. Lockery. Active currents regulate sensitivity and dynamic range in *C. elegans*. *Neuron*, 20:763–772, 1998.

[51] MS Grady, MR Mclaughlin, CW Christman, AB Valadaka, CL Flinger, and JT. Povlishock. The use of antibodies against neurofilament subunits for the detection of diffuse axonal injury in humans. *Journal of Neuropathology and Experimental Neurology*, 52:143–152, 1993.

[52] J Gray and HW Lissmann. The locomotion of nematodes. *J. Exp. Biol.*, 41:135–154, 1964.

[53] J.M. Gray, J.J. Hill, and C.I. Bargmann. A circuit for navigation in caenorhabditis elegans. *PNAS*, 102:3184–3191, 2005.

[54] DH Hall. *The posterior nervous system of Caenorhabditis elegans*. PhD thesis, California Institute of Technology, 1977.

[55] DH Hall and RL Russell. The posterior nervous system of the nematode caenorhabditis elegans: Serial reconstruction of identified neurons and complete pattern of synaptic interactions. *J. Neurosci.*, 11:122, 1991.

[56] MA Hemphill, BE Dabiri, S Gabriele, L Kerscher, C Franck, JA Goss, PW Alford, and KK Parker. A possible role for integrin signaling in diffuse axonal injury. *PLos ONE*, 6(7):e22899, 2011.

[57] MA Hemphill, S Dauth, CJ Yu, BE Dabiri, and KK. Parker. Traumatic brain injury and the neuronal microenvironment: A potential role for neuropathological mechanotransduction. *Neuron*, 86(6):1177–1192, 2015.

[58] J Hertz. Cross-correlations in high-conductance states of a model cortical network. *Neural Comput*, 22:427447, 2010.

[59] P. J. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge.

[60] VE Johnson, W Stewart, and DH. Smith. Axonal pathology in traumatic brain injury. *Experimental Neurology*, 246:35–43, 2013.

[61] L.M. Jones, A. Fontanini, B.F. Sadacca, P. Miller, and D.B. Katz. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc. Natl. Acad. Sci. USA*, 104:18772–18777, 2007.

[62] RE Jorge, L Acion, T White, D Tordesillas-Gutierrez, R Pierson, B Crespo-Facorro, and VA Magnotta. White matter abnormalities in veterans with mild traumatic brain injury. *American Journal of Psychiatry*, 169(12):1284–1291, 2012.

[63] C.C. Karaaslanli. *Bifurcation Analysis and Its Applications*. 2012.

[64] J Karbowski, CJ Cronin, A Seah, JE Mendel, D Cleary, and PW Sternberg. Conservation rules, their breakdown, and optimality in caenorhabditis sinusoidal locomotion. *J. Theoretical Biol.*, 242(3):652–669, 2006.

[65] A. Kocabas, C. H. Shen, Z. V. Guo, and S. Ramanathan. Controlling interneuron activity in caenorhabditis elegant to evoke chemotactic behavior. *Nature*, 490:273–277, 2012.

[66] KV Kolaric, G Thomson, JM Edgar, and AM. Brown. Focal axonal swellings and associated ultrastructural changes attenuate conduction velocity in central nervous system axons: a computer modeling study. *Physiological reports*, 1(3), 2013.

[67] D Krstic and I. Knuesel. Deciphering the mechanism underlying late-onset alzheimer disease. *Nature Reviews Neuroscience*, 9(1):25–34, 2012.

[68] J Kunert, E Schlizerman, and JN. Kutz. Low-dimensional functionality of complex network dynamics: Neuro-sensory integration in the caenorhabditis elegans connectome. *Phys Rev E.*, 89:052805, 2014.

[69] J. N. Kutz. *Data-Driven Modeling & Scientific Computation: Methods for Complex Systems & Big Data*. Oxford University Press.

[70] J. N. Kutz, X. Fu, and S. L. Brunton. Multi-resolution dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 15(2):713–735, 2016. Preprint. Available: arXiv:1506.00564.

[71] G. Laurent. Olfactory network dynamics and the coding of multidimensional signals. *Nat. Rev. Neurosci.*, 3:884–895, 2002.

[72] G. Laurent, M. Stopfer, R.W. Friedrich, M.I. Rabinovich, A. Volkovskii, and H.D.I. Abarbanel. Odor encoding as an active, dynamical process: Experiments, computation and theory. *Annu. Rev. Neurosci.*, 24:263–297, 2001.

[73] PP Liberski and H. Budka. Neuroaxonal pathology in creutzfeldt-jakob disease. *Acta Neuropathology*, 97:329–334, 1999.

[74] C Machens, R Romo, and C. Brody. Functional, but not anatomical, separation of what and when in prefrontal cortex. *J Neurosci*, 30(1):350–360, 2010.

[75] MH Magdesian, FS Sanchez, M Lopez, P Thostrup, N Durisic, W Belkaid, D Liazoghli, P Grutter, and DR Colman. Atomic force microscopy reveals important differences in axonal resistance to injury. *Biophysical Journal*, 103(3):405–414, 2012.

[76] PD Maia, MA Hemphill, B Zehnder, C Zhang, KK Parker, and JN. Kutz. Diagnostic tools for evaluating the impact of focal axonal swellings arising in neurodegenerative diseases and/or traumatic brain injury. *Journal of Neuroscience Methods*, 253:233–243, 2015.

[77] PD Maia and JN. Kutz. Compromised axonal functionality after neurodegeneration, concussion and/or traumatic brain injury. *Journal of Computational Neuroscience*, 27:317–332, 2014.

[78] PD Maia and JN. Kutz. Identifying critical regions for spike propagation in axon segments. *Journal of Computational Neuroscience*, 36(2):141–155, 2014.

[79] E Marder. Motor pattern generation. *Current Opinion in Neurobiology*, 10:691698, 2000.

[80] E Marder and RL Calabrese. Principles of rhythmic motor pattern generation. *Physiological Reviews*, 76:687–717, 1996.

[81] WL Maxwell, JT Povlishock, and DL. Graham. A mechanistic analysis of nondisruptive axonal injury: A review. *Journal of Neurotrauma*, 17(7):419–440, 1997.

[82] Igor Mezic. Analysis of fluid flows via spectral properties of the Koopman operator. *Ann. Rev. Fluid Mech.*, 45:357–378, 2013.

[83] S Millecamps and JP. Julien. Axonal transport deficits and neurodegenerative diseases. *Nature Reviews Neuroscience*, 14(161):161–176, 2013.

[84] A Morrison, A Aertsen, and M Diesmann. Spike-timing-dependent plasticity in balanced random networks. *Neural Comput*, 19:14371467, 2007.

[85] B Morrison, BS Elkin, JP Dolle, and ML. Yarmush. In vitro models of traumatic brain injury. *Annual Reviews in Biomedical Engineering*, 13(1):91–126, 2011.

[86] OJ Mullins, JT Hackett, JT Buchanan, and WO Friesen. Neuronal control of swimming behavior: Comparison of vertebrate and invertebrate model systems. *Progress in Neurobiology*, 93:244–269, 2011.

[87] E Niebur and P. Erds. Theory of the locomotion of nematodes: Control of the somatic motor neurons by interneurons. *Mathematical Biosciences*, 118:51–82, 1993.

[88] SN Niogi, P Mukherjee, J Ghajar, C Johnson, RA Kolster, R Sarkar, H Lee, M Meeker, RD Zimmerman, GT Manley, and BD McCandliss. Extent of microstructural white matter injury in postconcussive syndrome correlates with impaired cognitive reaction time: A 3t diffusion tensor imaging study of mild traumatic brain injury. *American Journal of Neuroradiology*, 29(5):967–973, 2008.

[89] HJ Park and K. Friston. Structural and functional brain networks: From connections to cognition. *Science*, 342:1238411, 2013.

[90] J. L. Proctor, S. L. Brunton, and J. N. Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016.

[91] JL Proctor, SL Brunton, and JN Kutz. Dynamic mode decomposition with control. *SIAM Journal of Applied Dynamical Systems*, 15:142–161, 2016.

[92] Y.B. Qi, M.D. Po, P. Mac, T. Kawano, E.M. Jorgensen, M. Zhen, and Y. Jin. Hyperactivation of b-type motor neurons results in aberrant synchrony of the caenorhabditis elegans motor circuit. *J. Neurosci.*, 33(12):5319–5325, 2013.

[93] M. Rabinovich, R. Huerta, P. Varona, and V.S. Afraimovich. Transient cognitive dynamics, metastability, and decision making. *PLoS Comput. Biol.*, 4:e1000072, 2008.

[94] M. Rabinovich and P. Varona. Robust transient dynamics and brain functions. *Front. Comput. Neurosci.*, 5(24), 2011.

[95] M. Rabinovich, A. Volkovskii, P. Lecanda, R. Huerta, H.D.I. Abarbanel, and G. Laurent. Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lettl.*, 87:068102, 2001.

[96] A Renart, J De la Rocha, P Bartho, L Hollender, N Parga, A Reyes, and K Harris. The asynchronous state in cortical circuits. *Science*, 327:587590, 2010.

[97] DL Riddle, T Blumenthal, BJ Meyer, and JR Priess. *C. elegans II.* Cold Spring Harbor, 2nd edition. edition.

[98] JA Riffell, E Shlizerman, E Sanders, L Abrell, B Medina, AJ Hinterwirth, and JN Kutz. Flower discrimination by pollinators in a dynamic chemical environment. *Science*, 344:1515–1518, 2014.

[99] CW Rowley, I Mezic, S Bagheri, P Schlatter, and DS Henningson. Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 641:115–127, 2009.

[100] V Rubovitch, M Ten-Bosch, O Zohar, CR Harrison, C Tempel-Brami, E Stein, BJ Hoffer, CD Balaban, S Schreiber, WT Chiu, and CG Pick. A mouse model of blast-induced mild traumatic brain injury. *Experimental Neurology*, 232(2):280–289, 2011.

[101] PJ. Schmid. Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.*, 656:5–28, 2010.

[102] P. Sengupta and A. Samuel. Caenorhabditis elegans: a model system for systems neuroscience. *Current Opin. Neuro.*, 19:1–7, 2009.

[103] DJ Sharp, G Scott, and R. Leech. Network dysfunction after traumatic brain injury. *Nature Reviews Neurology*, 10:156–166, 2014.

[104] XN Shen and PE Arratia. Undulatory swimming in viscoelastic fluids. *Phys. Rev. Lett.*, 106:208101, 2011.

[105] E. Shlizerman, J. Riffell, and J. N. Kutz. Data-driven inference of network connectivity for modeling the dynamics of neural codes in the insect antennal lobe. *Frontiers in Neuroscience*, 18:70:1–15, 2014.

[106] Eli Shlizerman, K Schroder, and J Nathan Kutz. Neural activity measures and their dynamics. *SIAM J. Appl. Math.*, 72(4):1260–1291, 2012.

[107] DH Smith, JW Wolf, TA Lusardi, VMY Lee, and DF. Meaney. High tolerance and delayed elastic response of cultured axons to dynamic stretch injury. *The Journal of Neuroscience*, 19(11):4263–4269., 1999.

[108] H Sompolinsky, A Crisanti, and HJ. Sommers. Journal of physiology. *Physical Review Letters*, 61:259–262, 1988.

[109] G. J. Stephens, M. B. de Mesquita, W. S. Ryu, and W. Bialek. Emergence of long timescales and stereotyped behaviors in caenorhabditis elegans. *PNAS*, 108:7286–7289, 2011.

[110] GJ Stephens, B Johnson-Kerner, W Bialek, and WS. Ryu. Dimensionality and dynamics in the behavior of c. elegans. *PLOS Computational Biology*, 4(4), 2008.

[111] J.E. Sulston, E. Schierenberg, J.G. White, and J.N. Thomson. The embryonic cell lineage of the nematode *Caenorhabditis elegans*. *Developmental Biology*, 100:64–119, 1983.

[112] D Sussillo and LF. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63:544–557, 2009.

[113] Y. Susuki, I. Mezic, and T. Hikihara. Coherent swing instability of power grids. *Journal of Nonlinear Science*, 3:403–439, 2011.

[114] B Szigeti, P Gleeson, M Vella, S Khayrulin, A Palyanov, J Hokanson, M Currie, M Cantarelli, G Idili, and S Larson. Openworm: an open-science approach to modeling *Caenorhabditis elegans*. *Frontiers in Computational Neuroscience*, 8:137, 2014.

[115] J Sznitman, X Shen, R Sznitman, and PE Arratia. Propulsive force measurements and flow behavior of undulatory swimmers at low reynolds number. *Phys. Fluids*, 22:121901, 2010.

[116] MD Tang-Schomer, VE Johnson, PW Baas, W Stewart, and DH. Smith. Partial interruption of axonal transport due to microtubule breakage accounts for the formation of periodic varicosities after traumatic axonal injury. *Experimental Neurology*, 233:364–372, 2012.

[117] MD Tang-Schomer, AR Patel, PW Bass, and DH. Smith. Mechanical breaking of microtubules in axons during dynamic stretch injury underlies delayed elasticity, microtubule disassembly, and axon degeneration. *The FASEB Journal*, 24(5):1401–1410, 2010.

[118] BD Trapp, J Peterson, RM Ransohoff, R Rudick, S Mork, and L. Bo. Axonal transection in the lesions of multiple sclerosis. *The New England Journal of Medicine*, 338:278–285, 1998.

[119] J Tsai, J Grutzendler, K Duff, and WB. Gan. Fibrillar amyloid deposition leads to local synaptic abnormalities and breakage of neuronal branches. *Nature Neuroscience*, 7:1181–1183, 2004.

[120] JH Tu, DM Luchtenburg, CW Rowley, DM Luchtenburg, SL Brunton, and JN Kutz. On dynamic mode decomposition: theory and applications. *J. Comput. Dyn.*, 1:391–421, 2014.

[121] C van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274:1724–1726, 1996.

[122] L.R. Varshney, B.L. Chen, E. Paniagua, D. H. Hall, and D. B. Chklovski. Structural properties of the *Caenorhabditis elegans* neuronal network. *PLoS Comput. Biol.*, 7(2):e1001066, 2011.

[123] A. Vidal-Gadea, S. Topper, L. Young, A. Crisp, L. Kressin, E. Elbel, T. Maples, M. Brauner, K. Erbguth, A. Axelrod, A. Gottschalk, D. Siegel, and J.T. Pierce-Shimomura. Caenorhabditis elegans selects distinct crawling and swimming gaits via dopamine and serotonin. *Proc. Natl. Aca. Sci. USA*, 108(42):17504–9, 2011.

[124] T Vogels and L Abbott. Signal propagation and logic gating in networks of integrate-and-fire neurons.

[125] T. Wakabayashi, I. Kitagawa, and R. Shingai. Neurons regulating the duration of forward locomotion in caenorhabditis elegans. *Neurosci. Research*, 50:103–111, 2004.

[126] J Wang, RJ Hamm, and JT. Povlishock. Traumatic axonal injury in the optic nerve: evidence for axonal swelling, disconnection, dieback and reorganization. *Journal of Neurotrauma*, 28(7):1185–1198, 2011.

[127] S Ward, N Thomson, JG White, and S Brenner. Electron microscopal reconstruction of the anterior sensory anatomy of the nematode caenorhabditis elegans. *J. Comp. Neurol.*, 160:313–338, 1975.

[128] RW Ware, C Clark, K Crossland, and RL Russell. The nerve ring of the nematode, caenorhabditis elegans. *J. Comp. Neurol.*, 162:71–110, 1975.

[129] DJ Watts and SH. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.

[130] Q. Wen, M. Po, E. Hulme, S. Chen, X. Liu, S. Kwok, M. Gershow, A. Leifer, V. Butler, C. Fang-Yen, T. Kawano, W. Schafer, G. Whitesides, M. Wyart, D. Chlovskii, M. Zhen, and A. Samuel. Proprioceptive coupling within motor neurons drives c. elegans forward locomotion. *Neuron*, 76:750–761, 2012.

[131] J. G. White, E. Southgate, J.N. Thomson, and S. Brenner. The structure of the nervous system of the nematode *Caenorhabditis elegans. Phil. Trans. Roy. Soc. Lond. B*, 314:1–340, 1986.

[132] JG White, E Southgate, JN Thomson, and S Brenner. The structure of the ventral cord of caenorhabditis elegans. *Phil. Trans. R. Soc. Lond.*, 275:327–348, 1976.

[133] S Wicks, C Roehrig, and C. Rankin. A dynamic network simulation of the nematode tap withdrawal circuit: Predictions concerning synaptic function using behavioral criteria. *J. Neurosci.*, 16(12):4017–4031, 1996.

[134] KA Wright. *Nematodes as biological models*, chapter Nematode sense organs. New York: Academic Press, 1980.

[135] Y Xiong, A Mahmood, and M. Chopp. Animal models of traumatic brain injury. *Nature Reviews Neuroscience*, 14(22):128–142, 2013.