Title
Characterizing Real World Neural Systems Using Variational Methods of Data Assimilation

Permalink
https://escholarship.org/uc/item/8dz7c69k

Author
Breen, Daniel

Publication Date
2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Characterizing Real World Neural Systems Using Variational Methods of Data Assimilation

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

in

Physics

by

Daniel Breen

Committee in charge:

Professor Henry Abarbanel, Chair
Professor Daniel Arovas
Professor Seth Leher
Professor Tatyana Sharpee
Professor Stefan Tanaka

2017

The dissertation of Daniel Breen is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2017

DEDICATION

To my family and everyone else who helped make this work and dissertation possible.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

Thanks to other members of my group. Your hard work and curiosity helped to set the pace and spark new ideas. Thanks to my advisor, Henry Abarbanel, for his guidance and encouragement.

Chapter 3 is reproduced in part from material as it appears in a preprint being prepared for publication. “Assimilation of Biophysical Neuronal Dynamics in Neuromorphic VLSI.” Jun Wang, Daniel Breen, Abraham Akinin, Frederic Broccard, Henry DI Abarbanel, and Gert Cauwenberghs. The dissertation author was one of two first authors on this preprint.

Chapter 4 is reproduced from material as it appears in a preprint on ArXiV. “HVC Interneuron Properties from Statistical Data Assimilation.” Daniel Breen, Sasha Shirman, Eve Armstrong, Nirag Kadakia, and Henry DI Abarbanel. The dissertation author was the first author on this preprint.

Chapter 5 is reproduced from material as it appears in a preprint being prepared for publication. “Use of Data Assimilation for Inference of CA1 Neuron Pathology in 3xTg Mouse Model of Alzheimers.” Daniel Breen, Clark Briggs, Grace Stutzmann, and Henry DI Abarbanel. The dissertation author was the first author on this preprint.

VITA

2011	B.S. in Physics *with Honors*, New Mexico Institute of Mining and Technology
2011-2017	Graduate Teaching Assistant, University of California, San Diego
2013	M.S. in Physics, University of California, San Diego
2017	Ph.D. in Physics, University of California, San Diego

PUBLICATIONS

Breen, D., Shirman, S., Armstrong, E., Kadakia, N., DI Abarbanel, H. (2016). “HVC Interneuron Properties from Statistical Data Assimilation”, ArXiv Preprint.

Wang, J., Breen, D., Akinin, A., DI Abarbanel, H., Cauwenberghs, G. (2016, October). “Data Assimilation of Membrane Dynamics and Channel Kinetics with a Neuromorphic Integrated Circuit”, In Biomedical Circuits and Systems Conference (BioCAS), 2016 IEEE (pp. 584-587). IEEE.

N. Kadakia, E. Armstrong, D. Breen, U. Morone, A. Daou, D. Margoliash, H. D. I. Abarbanel, “Nonlinear Statistical Data Assimilation for HVCRA Neurons in the Avian Song System”, Biological Cybernetics, (2016).

ABSTRACT OF THE DISSERTATION

Characterizing Real World Neural Systems Using Variational Methods of Data Assimilation

by

Daniel Breen

Doctor of Philosophy in Physics

University of California, San Diego, 2017

Professor Henry Abarbanel, Chair

Traditionally, characterizing many properties of biological or silicon neural systems has been expensive, laborious, or impossible. Conductance models describing how properties of these systems change with time can be used with accessible data, such as measured voltage traces, to help characterize inaccessible properties such as ionic currents or transistor mismatch. This is accomplished using variational methods which formulate an inference problem about these properties as nonlinear optimization. Because measurement noise and model error are inevitable in the study of complex systems, the method is designed to cope with unknown processes. Conductance models are overparameterized, causing the inference problem to remain underdetermined, which can result in a proliferation of widely separated sets of estimated model parameters producing accurate predictions. Additionally, real world data will be approximated by a model in a number of ways, leading to an additional contribution to this model identifiability problem. This dissertation probes and overcomes some of the difficulties encountered in the analysis of real world data in individual biological and silicon neurons. One key result is the characterization of a neuromorphic silicon neuron followed by emulation of a biological neuron on the silicon substrate. Another key result is a data mining approach which discovers statistical differences in estimated model parameters, despite underdeterminacy, in an Alzheimer's strain of neurons in mice compared to healthy controls.

Chapter 1

Data Assimilation

Physicists are interested in finding physical mechanisms underlying observed data. Although a few simple systems such as particle trajectories can be explained by an appeal to first principles, the absence of such explanatory descriptions for the dynamics of large networks confronts the modeler with a challenge. There are a set of tools which allow the modeler to estimate and validate quantitative models positing physical mechanisms underlying complex dynamical systems in a reproducible way. The application of some of these tools to the description of neurons is the main focus of this dissertation.

1.1 The Inference Problem

The quantitative measurement and modeling of complex dynamical systems inevitably involves inaccuracies. Systems in the real world almost always contain processes that the modeler is ignorant about or cannot represent. Another difficulty comes from the fact that measurements are noisy, which limits the ability to infer properties of systems even in the presence of perfect models. Measurement noise and model error suggest a probabilistic approach to reasoning about the system under study, here through the use of conditional probabilities of the state variables given a set of observations within an estimation window [1]. Computing these probability distributions is usually intractable because it requires the evaluation of high dimensional integrals. The strategy here is to instead seek expected values of the probability distribution through a variational approximation [66].

Dynamical systems involve describing the evolution of a system in time. Our system consists of a D dimensional state vector $\mathbf{x}(t)$. In formulating our inference problem, we will approach the specification of the system's trajectory $\mathbf{x}(t)$ by making use of a discrete time description, defining the state of the system at times $\{t_1,t_2,...,t_n\}$. We then define the discrete time system trajectory X as $\{\mathbf{x}(t_1),\mathbf{x}(t_2),...,\mathbf{x}(t_n)\}$. In general, the state of the system $\mathbf{x}$ at a time t_i will depend on the state of the system at previous times. We are interested in systems where the state of the system $\mathbf{x}(t_i)$ depends on the state of the system at the previous time point $\mathbf{x}(t_{i-1})$ only. This simplification, known as the Markov property, is one which we will exploit in the derivation of the following probabilistic formulation of our inference problem.

Modeling a complex system involves positing processes which are not measured. These are included as components within X. There are also L processes we can measure. For our problems, L is typically much smaller than D, the dimension of the state vector for our system. The collection of the L processes $\mathbf{y}$ is defined at times $\{\mathbf{y}(t_1),\mathbf{y}(t_2),...,\mathbf{y}(t_n)\}$ and is our data, which we call Y.

We will unravel the probability $P(X|Y)$ of observing the trajectory X given our observed data Y. To do this we will use Bayes' rule, which is an identity of probability theory, and the Markov property.

$$P(X|Y) = \frac{P(X,Y)}{P(Y)} \tag{1.1}$$

$P(Y)$ is a term independent of X which we can ignore since we are eventually interested in expectation values over X. The reason we can ignore it is that $P(Y)$ will appear in numerator and denominator, so it will cancel. We now first make a double application of Bayes' rule and then utilize the Markov property to create a recursion relationship for the probability of the trajectory of the system $P(X_n, Y_n)$ up to time t_n and the probability of the system $P(X_{n-1}, Y_{n-1})$ up to time t_{n-1}.

$$\begin{aligned}
P(X_n, Y_n) &= P(\mathbf{y}(t_n)|X_n, Y_{n-1})P(X_n, Y_{n-1}) \\
&= P(\mathbf{y}(t_n)|X_n, Y_{n-1})P(\mathbf{x}(t_n)|X_{n-1}Y_{n-1})P(X_{n-1}, Y_{n-1}) \\
P(\mathbf{y}(t_n)|X_n, Y_{n-1}) &= P(\mathbf{y}(t_n)|\mathbf{x}(t_n)) \\
P(\mathbf{x}(t_n)|X_{n-1}Y_{n-1}) &= P(\mathbf{x}(t_n)|\mathbf{x}(t_{n-1})) \\
P(X_n, Y_n) &= P(\mathbf{x}(t_n)|\mathbf{x}(t_{n-1}))P(\mathbf{y}(t_n)|\mathbf{x}(t_n))P(X_{n-1}, Y_{n-1}) \qquad (1.2)
\end{aligned}$$

In the third and fourth lines we used the Markov property, the assumption that our measurements are independent of each other in time, and that measurements at previous times give us no information about the state of the system at the current time. Of course, none of these assumptions are true, but they help to shrink the complexity of the problem into a manageable size. We only have to apply our recursion relationship repeatedly to get our expression for $P(X,Y)$, and by extension, $P(X|Y)$.

$$P(X_n, Y_n) = \prod_{i=2}^{n} P(\mathbf{x}(t_i)|\mathbf{x}(t_{i-1}))P(\mathbf{y}(t_i)|\mathbf{x}(t_i))P(\mathbf{x}(t_1))P(\mathbf{y}(t_1)|\mathbf{x}(t_1)) \qquad (1.3)$$

$P(\mathbf{x}(t_i)|\mathbf{x}(t_{i-1}))$, the transition probability, depends on the dynamical model of the physical system. We take the model to have the explicit form:

$$\mathbf{x}(t_i) = \mathbf{f}(\mathbf{x}(t_{i-1})) + \eta(t_{i-1}) \tag{1.4}$$

$\eta(t_{i-1})$ represents a stochastic component in our otherwise deterministic model. $P(\mathbf{y}(t_i)|\mathbf{x}(t_i))$ is the probability recording observation $\mathbf{y}(t_i)$ if the system is in configuration $\mathbf{x}(t_i)$ at that point in time. We will deal with deterministic models in the rest of the dissertation, so we set $\eta = 0$.

We will define a projection operator with components $h_j(\mathbf{x}_i)$ which maps the system to the components of our measurements at each time point, $y_j(t_i)$. We have assumed of our measurement noise that it is additive and uncorrelated in time. We additionally assume it is distributed according to a Gaussian, and we denote the non-zero elements of the diagonal inverse correlation matrix R_m by $R_m(l)$. We assume the model errors are also Gaussian distributed and denote the non-zero values of the inverse model covariance R_f by $R_f(d)$. While our model does not explicitly contain model error, this error provides a relaxation of the dynamical constraints of the model and has proven useful elsewhere [38, 24].

Lastly, we assume maximal ignorance of the initial state $\mathbf{x}(t_0)$, taking it to be uniformly distributed over the dynamical range of the model, though the assumption of such a non-informative prior is rare in large geophysical and meteorological applications [38], where low background covariances convey relatively strong belief in recent forecasts.

Together, these assumptions lead to the Gaussian error action in discrete time [1]:

$$A_0(X|Y) = \sum_{n=1,l=1}^{N,L} \frac{R_m(l)}{2}(h_l(\mathbf{x}(t_n)) - y_l(t_n))^2 \\ + \sum_{n=1,d=1}^{N-1,D} \frac{R_f(d)}{2}(x_d(t_{n+1}) - f_d(\mathbf{x}(t_n)))^2 \quad (1.5)$$

with

$$P(X,Y) = \exp(-A_0(X|Y)) \quad (1.6)$$

If we are interested in the state of the system at the end of an estimation window, we can use another identity of probability theory and marginalize over all possible configurations X_{n-1}. In that case, we arrive at a path integral formulation of our inference problem, where the analogy to the path integral commonly occurring in physics becomes more apparent.

$$P(\mathbf{x}(t_N)|Y) = \int \prod_{i=1}^{N-1} dx_i \exp(-A_0(X|Y)) \quad (1.7)$$

1.2 From Inference to Optimization

It is computationally expensive to evaluate the path integral expression, equation 1.7, or even its predecessor, equation 1.6. Fortunately, in most applications we are not interested in knowing the full probability distribution. We are interested in expectation values $G(\mathbf{x}_n)$ of the probability distribution.

$$G(\mathbf{x}_n) = \int G(X) \prod_{i=1}^{N-1} dx_i \exp(-A_0(X|Y)) \quad (1.8)$$

with equation 1.7 as the appropriate normalization factor.

Such expectation values are easy to compute when only a small portion of the overall probability distribution of equation 1.7 contributes to the integral. In cases when up to only a few regions of path space X dominate the contribution to the integral, the expectation values can be computed via the method of Laplace. In order to do this, the expression for the action A_0 of equation 1.5 in the exponential of equation 1.7 must contain up to only a few minima.

Since one must then proceed to find these regions of path space X, the challenge of the inference problem is now shifted to one of nonconvex optimization. There are a number of ways that the optimization problem can be formulated in addition to minimizing the cost function of equation 1.5 directly. One way is to impose the dynamical rule evolving the system forward in time as a set of nonlinear constraints on the optimization problem. This is a route taken in the last chapter of this disseration. Another method is to instead iteratively use the information contained in educated initial conditions about where these regions may be located. This method is called annealing and is developed and applied substantially in this dissertation to the application of complex real world neural systems.

The values of $R_m(l)$ and $R_f(d)$ are assigned before the beginning of the optimization procedure, and reflect the relative weighting of evidence from the data versus prior information in the model from a Bayesian perspective.

The Gaussian error action (1.5) uses different notation from, but is otherwise equivalent to, the weak 4D-Var cost function known in the meteorological literature [84, 38]. This is so as long as there is zero background precision in the action (Eq. 1.5), the assumption of a uniform prior.

$$\begin{aligned}J(\mathbf{x}) = \frac{1}{2}(\mathbf{x}_0 - \mathbf{x}_b)\mathbf{B}^{-1}(\mathbf{x}_0 - \mathbf{x}_b) \\ + \frac{1}{2}\sum_{n=0}^{N}(\mathbf{y}_n - H_n(\mathbf{x}_n))\mathbf{R}_n^{-1}(\mathbf{y}_n - H_n(\mathbf{x}_n)) \\ + \frac{1}{2}\sum_{n=0}^{N}(\mathbf{x}_{n+1} - f(\mathbf{x}_n))\mathbf{Q}_n^{-1}(\mathbf{x}_{n+1} - f(\mathbf{x}_n))\end{aligned} \quad (1.9)$$

Here we use the form of the action in (1.5) as it corresponds to the notation used in our previous work on DA in biological and physical settings [93, 94, 40].

We call A_0, the log-likelihood of $P(\mathbf{x}(t_N)|Y)$, the 'action' due to its analogy with Lagrangian mechanics. In the present work, we approximate the integral (1.7) utilizing Laplace's method. In continuous time, the action is a functional, and Laplace's method amounts to finding extrema paths X_q. Finding the extrema of the action is known in classical mechanics as the variational principle, where the extrema are solutions to the Euler-Lagrange equations [27, 50, 46].

In physics, typically the solutions to the Euler-Lagrange equations are constrained as an initial-value problem. In data assimilation, solutions to the Euler-Lagrange equations are constrained instead as a two-point boundary-value problem [8, 9, 94, 38].

1.2.1 Annealing

The implementation of data assimilation algorithms runs into numerical difficulties when the system is nonlinear and of high dimension, and when the measurements are sparse and noisy. Ill conditioning of the Hessian, cliffs in the cost function, as well as a proliferation of critical points such as local minima and saddle points make it difficult to find the 'correct' set of model parameters. In neuronal systems, typically the time course of the membrane $V(t)$ at the soma can be measured, but not the activation of the gating variables or most ionic concentrations. The annealing method that we discuss here is meant to mitigate this challenge somewhat, but because conductance models of neurons are often overparameterized and because it is only ever possible to approximate a real system, there is a proliferation of close to equally valid minima, problems remain. One potential limitation of annealing as it is presently formulated is that it imposes the model constraints in an unflexible, discrete manner. Allowing the optimization algorithm to instead flexibly impose the constraints at every time point in a continuous manner, could be a worthwhile avenue to explore in improving the annealing method.

The way in which R_m and R_f are weighted relative to one another in the cost function influences the result of minimizing it. Manipulating the cost function by varying these values forms the basis of our annealing method, shown to be effective in state and parameter estimation in archetypal chaotic models such as the Lorenz ’96 model [93]. When $R_m \gg R_f$, the model error is a perturbation to the measurement error. Such an assumption causes the cost function to form minima where measured states in the model fit the data closely. The model error is then permitted to be large and so is enforced weakly, effectively decoupling parameters and unmeasured states from the data. Conversely, when $R_m \approx R_f$, both terms contribute equally to the cost function, so minimizing the cost function will tend to satisfy the data while simultaneously enforcing the dynamical map. In contrast, when the model error is forced, by large R_f, to be small, the nonlinearity of the vector field $\mathbf{f(x)}$ manifests itself at the smallest scales in the phase space of the paths X over which we are searching. This results in complicated fine structure seen as multiple local minima [1] in the action, especially when the number of measurements L is too small. It is unlikely that directly minimizing the cost function under this condition will yield good estimates of the system’s parameters and state variables.

The innovation of the annealing method developed in previous work [93] is then to start with $R_m \gg R_f$ initially. The cost function is then minimized, which is typically easy, since the model constraints are in a role of giving a small perturbation to the parabolic, and therefore convex, measurement constraints. Then R_f is increased in magnitude by a factor $\alpha > 1$, and the cost function is minimized again, starting the search for minima at the previous solution. The process is repeated until $R_f \gg R_m$. In this way, the nonlinearities are introduced gradually so that by essentially just using Newton's method, the effect of the computations amounts to carefully crawling towards a minimum which fits the data and the model better than other options.

The implementation of the annealing algorithm as well as other formulations of the optimization problem used in this disseration were accomplished through the use of the open source software package IPOPT (Interior Point OPTimizer) with the linear solver ma57 [90].

Chapter 2

Neurobiological Background

The focus of this dissertation is on developing and applying methods of data assimilation for inferring parameters and states in neural systems when confronted with real data. Neurons are complex systems, while the models used to describe them contain only a sliver of the complexity of the real system. While there is an extensive literature focusing on computational modeling of neurons, often this literature focuses on estimating specific facets of neurons such as the properties of individual ionic currents or detecting qualitative differences between populations of neurons, such as snapshots of these altered processes, including the magnitude of calcium release from internal stores or afterhyperpolarization. They are not usually focused on predicting dynamics, and they often average over many different neurons instead of characterisizing the variability over populations of neurons.

Data assimilation provides a means to ask and answer specific questions about complex real physical systems, such as neurons, from limited data. In this chapter, we will review neurobiological background material to set the stage for the next several chapters.

2.1 Characterising Neuron Electrophysiology

2.1.1 Neurobiological Experiments

The data that is the topic of analyses covered in this text consists of time series measurements of the membrane voltage in response to a specified current injected into the neuron. The membrane voltage is the difference between the voltage at the outside and inside of the cell. The cell membrane is a barrier selectively porous to currents and other substances.

Specialized instruments deliver the current into the cell and record the voltage. In broad strokes, a narrow pipette on the order of microns in diameter sucks in a small patch of cell membrane. The membrane forms a bond with the inside of the pipette, and the remaining membrane inside the pipette is broken so that the cell is open to the pipette, but a seal has been formed so that current can be injected directly into the cell while the membrane voltage can be recorded.

A ’current clamp’ consists of ’clamping’, or fixing, the membrane potential at a fixed value, then superimposing an injected current to manipulate the voltage on top of that. The current clamp procedure is designed to manipulate the voltage of the neuron in such a way as to activate selectively permeable current channels in the membrane of the neuron. There is an enormous variety of current channels on a neuron membrane which acts to either push the membrane voltage to more positive or more negative values.

The goal of data assimilation is to characterize these current channels computationally, in some cases to characterize the biophysics of the neuron to predict the response of the neuron to future driving currents, in other cases to identify population differences between classes of neurons for biomedical purposes. We will devote some time to understanding how to model these current channels computationally.

2.1.2 The Charge Conservation Equation

The neuron membrane is basically just a capacitor across which currents can pass and deposit charge on the inside or the outside. The expression for the change in voltage over time is just

$$C_m \frac{dV}{dt} = \sum_i I_i \tag{2.1}$$

where C_m is the membrane capacitance of the neuron and V is the membrane voltage. The I_i are just currents which flow into or out of the membrane.

We could just list all of the many ionic currents on the right hand side, but this would lead to an intractably complicated equation. So we typically choose only a few of the many I_i to list on the right hand side. We will choose only a couple of the most important, I_{Na} and I_K. I_{Na} is the inactivating sodium current, meaning that it is capable of turning itself off, and typically shortly does turn itself off after being activated. In the normal range of physiological behavior, it tends to move the cell to more positive membrane voltages. I_{Na} passes Na$^+$ ions, while I_K passes K$^+$ ions. I_K is the delayed rectifier current. It is not inactivating and tends to move the cell to more negative membrane voltages.

In addition to I_{Na} and I_K, we will also include the leak current I_L. The membrane is somewhat nonselectively porous to ions, but I_L also plays the role of a fudge factor, mostly playing the role of an offset DC current and giving the neuron's membrane a characteristic electrical conductance at baseline voltages.

Finally we include the injected current I_{inj} so that our current balance equation reads

$$C_m \frac{dV}{dt} = I_{Na} + I_K + I_L + I_{inj} \tag{2.2}$$

We will now turn to characterising I_K and I_{Na}.

2.1.3 Conductance Based Models

We have yet to specify the form of the current channels embedded in the neurons membrane are actually proteins which have a complicated structure designed to change in response to a wide variety of factors, including cell signalling ligands and the voltage difference across the cell membrane. Ligands are chemicals which bind to places on the protein and change its structure, thereby changing how it passes currents. The simplest example of a ligand is a calcium ion. Calcium ions are involved in many aspects of cell signalling and a few of those will be touched upon later in this work. I_K and I_{Na} are among those proteins which pass current in a way that is strongly dependent on membrane voltage.

It would be extremely difficult to exactly characterize how individual channels respond to changes in membrane voltage and environmental mileau. The modeler makes a choice here to use a very simplified rule describing how the channels change in response to voltage. Presently, there aren't many good alternatives to accepting model errors in exchange for the simple and effective conductance based *Hodgkin-Huxley* (H-H) description.

For most channels, the H-H framework says that the current dependence on voltage is characterized by an equilibrium value as the channel is held at that voltage for an infinite time. This equilibrium value depends on a few factors.

1. The number of ion channels
2. The max amount of current each channel can pass per unit concentration of a particular ion species

3. The ratio of channels which have undergone the number of conformational changes required to open the channel
4. The ratio of channels which have undergone a conformational change which inactivates the channel
5. The ratio of intracellular and extracellular concentrations of a particular ion species

The first two conditions together constitute the max conductance of an ion current, determined by the maximal amount of current passed per channel times the number of channels.

The proteins which pass current and are embedded in the membrane of the channel undergo conformational changes which allow them to pass currents. These are the third and fourth conditions. In the H-H framework, these conformational changes each occur independently of one another, though this is unlikely to be the case in reality. The probability that channels have undergone such a conformational change, in the H-H framework, depends strictly on the voltage.

The population of proteins relax to this ratio exponentially at a characteristic rate which also depends strictly on membrane voltage. Again, this is unlikely to be the case in reality, but is a necessary simplification in order to shrink the inherent biological complexity to a manageable size for present computational resources.

Finally, the flux term summarizes the electrical and chemical forces pushing ions into and out of a cell. It accounts for relative differences in the amount of an ion species inside and outside of the cell and depends on the membrane voltage.

The equations which summarize the above are the following.

$$I_M(V(t)) = g_M m(t)^a h(t)^b \Phi(M,V) \tag{2.3}$$

$$\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_m(V)} \tag{2.4}$$

$$\frac{dh}{dt} = \frac{h_\infty(V) - h}{\tau_h(V)} \tag{2.5}$$

where I_M is an ionic current for ion species M with maximal conductance g_M. a denotes the number of conformational changes a protein must undergo to open a channel, while b denotes the number of independent conformational changes the same protein must undergo to close a channel. $\Phi(M,V)$ is the flux term acoounting for relative differences of an ion species inside and outside of a cell. m and h denote the probability that proteins of a particular species have undergone an activating or inactivating conformational change, respectively. The $x_\infty(V)$ denote the equilibrium ratio of channels having undergone a conformational change given the voltage V, while the $\tau_x(V)$ are the voltage dependent rates at which the ratio of channels having undergone a conformational change m or h relaxes to the equilibrium ratio.

In the literature, typical forms for x_∞ and τ_x are the following.

$$x_\infty = 0.5(1 + \tanh((V - \theta_x)/\sigma_x)) \tag{2.6}$$

$$\tau_x = t_{x0} + t_{x1}(1 - \tanh((V - \theta_x)/\sigma_x)^2) \tag{2.7}$$

Although these are forms estimated in the literature and which have proven their usefulness in data assimilation, these forms are not set in stone. It is up to the modeler to seek alternative forms.

a and b are also determined on a case by case basis.

The form of $\Phi(M,V)$ derives from the Nernst-Planck equation, equation 2.8, for describing the flux of ion species M under the influence of drift and diffusion.

$$J = J_{diff} + J_{drift} = -D\nabla[M] - \mu z[M]\nabla V \tag{2.8}$$

D is the diffusion coefficient, $[M]$ is the concentration of M (mol/cm^3), μ is the mobility ($cm^2/(Vs)$) and z is the valence of the ion (dimensionless).

When ions are assumed not to interact with one another, the gradient of the potential is constant, and ion movement across the membrane is described by Nernst-Planck, then solving for the ion current across a barrier with constant electric field as a function of ion concentration and membrane potential gives equation 2.9.

$$I([M],V) = PzF\xi\frac{[M]_{in} - [M]_{out}exp(-\xi)}{1 - exp(-\xi)} \tag{2.9}$$

where $\xi = \frac{zVF}{RT}$. F is Faraday's constant (96840 C/mol), and R is the gas constant (1.98 $cal/(Kmol)$). The value of P is determined by the expression $P = \frac{\beta u{*}RT}{lF}$, where β is the dimensionless water membrane partition coefficient for the ion, l is the width of the membrane, and $u*$ is the molar mobility of the ion within the membrane ($cm^2/(Vsmol)$).

As our expression relates the current I to the membrane voltage and ion concentrations inside and outside of the cell, we can absorb the coefficients PzF into the max conductance g_M and identify the remaining expression with $\Phi([M],V)$.

$$\Phi([M],V) = V\frac{[M]_{in} - [M]_{out}exp(-\xi)}{1 - exp(-\xi)} \tag{2.10}$$

In the case of Na^+ and K^+ currents, the concentrations inside and outside the cell of Na^+ or K^+ do not change that much when Na^+ or K^+ is allowed to flow across the membrane. These are held constant, and the resulting expression is well approximated by a first order taylor expansion around the potential V where the current changes sign, which is known as the reversal potential. This value can be found with a simple application of L'Hospitals rule. Thus, the expressions for the Na^+ or K^+ currents are approximately ohmic with nonlinear conductances depending on voltage, which itself depends on time.

We note in passing that the ion Ca^{2+} is not as well approximated by a first order taylor expansion because the concentrations of Ca^{2+} inside and outside of the cell are highly disparate. Additionally, because the concentration of Ca^{2+} is very small inside the cell and very large outside of the cell, intracellular Ca^{2+} increases dramatically during influx of Ca^{2+} through Ca^{2+} channels. As a result, for some applications one may need to retain the dependence of $\Phi_{Ca^{2+}}$ on Ca^{2+}. One example would be to increase the sensitivity of the calcium current on intracellular Ca^{2+} by immersing the neuron in a low Ca^{2+} concentration external bath, which could be useful in inferring intracellular levels of Ca^{2+} from measurements of voltage alone.

The only current that we have not examined yet is the leak current. This current might be interpreted as a catch all for many different kinds of ions which leak across the cell membrane, but in fact it is probably more accurate to say that it is a fudge factor which enables more accurate modeling. In modeling electrophysiological experiments, it proves useful as the reversal potential E_L can be set near the value at which the neuron is clamped, typically voltages where the neuron is not very active, while the conductance g_L represents the conductance of the neuron near membrane potentials at which the neuron is clamped.

For I_{Na}, $a = 3$ and $b = 1$, while for I_K, $a = 4$ and $b = 0$ is typical. Our H-H model for a neuron is then

$$C_m \frac{dV}{dt} = g_{Na} m^3 h (E_{Na} - V) + g_K n^4 (E_K - V) + g_L (E_L - V) + I_{inj} \tag{2.11}$$

$$\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_m(V)} \tag{2.12}$$

$$\frac{dh}{dt} = \frac{h_\infty(V) - h}{\tau_h(V)} \tag{2.13}$$

$$\frac{dn}{dt} = \frac{n_\infty(V) - n}{\tau_n(V)} \tag{2.14}$$

$$\tag{2.15}$$

with appropriate expressions for the steady state activations/inactivation $x_\infty(V)$ and relaxation times $\tau_x(V)$ such as those given above.

2.1.4 Modeling of Calcium Dynamics

With these preliminaries aside, we set the background for Chapter 6, where we will discuss promising directions for future research. Here we construct a mathematical model of calcium dynamics within a neuron which, when validated, could provide insight about disease mechanisms of Alzheimer's disease and potential therapeutic remedies. The purpose of this section is to give appropriate background for the discussion in the final chapter about the application of data assimilation to questions of biomedical interest. Our model will include terms for the dependence of calcium dynamics on RyR and IP_3 regulation, efflux through plasma membrane calcium ATPase (PMCA) type pumps, refilling of the ER with calcium ions through store-operated channels (SOCs) in the cell membrane, calcium release from cytoplasm to ER through sarco/endoplasmic reticulum calcium ATPase (SERCA) type pump, and passive calcium leak from the ER to the cytoplasm.

Calcium plays a crucial role in a multitude of intracellular signalling pathways. It acts as a second messenger in the cell, mediating a few processes including the following:

- initiating the biochemical cascades that lead to the changes in receptor insertion in the membrane, which underlie synaptic plasticity
- muscle contraction
- secretion of neurotransmitter at nerve terminals
- gene expression

Chemical reactions in single neurons causing the opening and closing of ion channels are often simplified as simple chemical reactions described by kinetic schemes. These models of chemical reactions hide the real complexity when such a level of detail is not required. In reality, there exist specific sequences of reactions leading from a cause, such as modulator release in the synaptic cleft, to an end effect, such as phosphorylation of voltage gated ion channels that change the kinetics of the channel. These sequences of reactions are known as intracellular signalling pathways.

Ignoring the diffusion of the participant molecules simplifies the model, which then focuses solely on modelling the reaction kinetics. Such an assumption can be justified when the diffusion of the participant molecules is much faster than reaction time course, such as signalling pathways involving both binding and enzymatic reactions.

Law of Mass Action

The simplest binding reaction is the law of mass action. With molecules A and B as the substrates, AB as the product, k^+ as the forward reaction rate, and k^- as the backward rate, the differential equation describing the rate of change of species A (or species B) is:

$$\frac{dA}{dt} = k^{-}[AB] - k^{+}AB \tag{2.16}$$

The time evolution of the product AB is the negative of the expression above. In order to calculate the concentrations of A, B, and AB for all time, one needs to know the initial concentrations of A, B and AB and apply the equations above. Modelling the process as a deterministic reaction assumes that there are a lot of reactants A and B and that the solution is well stirred. If this assumption does not hold, stochastic modelling is required.

One can rearrange the terms in the equation above to solve for the equilibrium ratio of reactants to products:

$$K_d = \frac{A_\infty B_\infty}{[AB]_\infty} = \frac{k^{-}}{k^{+}}$$

which has units of concentration.

Enzymatic Reactions

Enzymatic reactions are two-step reactions in which the action of one molecule, the enzyme E, results in a substrate S being converted into a product P via a reversible reaction that produces a complex ES. E itself is not consumed. This sort of reaction was described by Michaelis and Menten as the reaction sequence:

$$E + S \leftrightarrow ES \rightarrow E + P \tag{2.17}$$

The second reaction step leading to product P is assumed to be irreversible and typically the substrate is in excess, so the reaction sequence is limited by the amount of enzyme. Assuming the law of mass action, the production of complex ES and product P are described by the differential equations:

$$\frac{dES}{dt} = k_1[E][S] - (k_1^- + k^c)ES \tag{2.18}$$

$$\frac{dP}{dt} = k^c ES \tag{2.19}$$

$\frac{dP}{dt}$ can be written, assuming the enzyme and subtrate reaction is much faster than the enzyme to product reaction:

$$\frac{dP}{dt} = V_{max}\frac{S}{K_m + S} \tag{2.20}$$

This results from applying $E = E_{tot} - ES$, where E_{tot} is the total enzyme concentration, and rearranging.

This steady state approximation to the full enzymatic reaction sequence is often referred to as Michaelis-Menten kinetics. If the production of P involves the binding of a number n of identical molecules of S simultaneously, the rate of reaction is given by the Hill equation:

$$\frac{dP}{dt} = V_{max}\frac{S^n}{K_m^n + S^n} \tag{2.21}$$

where n is the Hill coefficient.

Calcium Conservation

Faradays constant gives the amount of charge carried by a mole of monovalent ions: 96490 coulombs per mole. Since calcium is divalent, we multiply this by two then divide the current I_{Ca} by the result to get the number of moles flowing per unit area and per unit time. Multiplying this by the total surface area and dividing by the compartment volume gives the rate of change in concentration, also known as the flux.

The flux that results from an ionic calcium current I_{Ca} in units of current per unit area, is given by:

$$J_{cc} = -\frac{aI_{Ca}}{2Fv} \tag{2.22}$$

where a is the surface area across which the current flows, v is the volume of the intracellular compartment and F is Faradays constant. This is how we translate current into a certain number of calcium ions entering the cellular compartment per unit time. v turns the rate of calcium ions entering the cell into a rate of change of calcium concentration in the compartment.

Other fluxes, such as those due to membrane-bound pumps and buffering, act to restore calcium to its resting level following an influx. The pool model captures the phenomenon that the calcium concentration will always return to its resting value. This model describes calcium decay by a single time constant τ_{Ca}, giving:

$$\frac{dCa}{dt} = J_{cc} - \frac{Ca^{2+} - Ca_e^{2+}q}{\tau_{Ca}} \tag{2.23}$$

This model is consistent with the calcium transients in cellular compartments with fluorescent dyes. Clearance of calcium from small compartments, such as spine heads, may be quite rapid, with a time constant as small as 12 ms. Similarly, clearance of calcium by buffering or other pumps from small domains around individual membrane calcium channels may differ from clearance from the bulk cytosol. More complex models include the fluxes that remove calcium from a cellular compartment, such as pumps, buffering and diffusion. The sum of these fluxes may be approximated by this single exponential decay model.

Membrane-bound pumps contribute to restoring resting levels of calcium by extruding calcium ions back through the membrane, moving them against their concentration gradient. A high-affinity pump has a low dissociation constant K_m so that it reaches its half-maximal pump rate at a low calcium concentration. High capacity means that the maximum pump velocity V_{max} is large. Two kinds of pumps act to remove calcium from a cell. The active membrane calcium-ATPase pump (knows as PMCA) is a high-affinity, low capacity mechanism that can switch on rapidly following calcium entry through voltage-gated channels, and so plays a major role shaping dynamic changes in intracellular calcium. This pump is also responsible for uptake of calcium into intraceullar stores. Passive membrane pumps such as the sodium-calcium exchanger are low-affinity, high capacity calcium pumps that are largely responsible for maintaining resting calcium levels.

Sodium Calcium Exchanger

Although complex pump models are available, a simpler approach to modeling the exchanger is to treat extrusion as an enzymatic reaction:

$$Ca + P \leftrightarrow CaP \tag{2.24}$$

$$CaP \rightarrow Ca_{ext} + P \tag{2.25}$$

The first equation refers to the binding of intracellular calcium to the pump. This changes the conformation of the pump, exposing bound calcium to the exterior facing membrane. This calcium is then released outside the cell. The energy causing the conformational change is provided by sodium ions coming down their concentration and electrical gradient into the cell. The pump flux can then be modified as an instantaneous function of the intracellular calcium concentration, described by a Michaelis-Menten relationship:

$$J_{pump} = V_{pump} \frac{Ca}{K_{pump} + Ca} \tag{2.26}$$

Biophysically, $V_{pump} = ak_2^+ P_m / v$, with P_m the number of pump molecules per unit area of membrane. $K_{pump} = (k_1^- + k_2^+)/k_1^+$, with $k_{1,2}^{+;-}$ the forward or backward rate constants for reaction 1 or 2, respectively.

Data is scarce concerning the properties of these pumps, and data assimilation is well positioned to provide information on them, including the value of parameters in the model, once these measurements become available. Estimates of pump velocity for the calcium-ATPase pump in hippocampal pyramidal cells range over several orders of magnitude from around 10^{-13} to 10^{-10} mol cm^{-2}. The surface area of a dendritic spine in rat hippocampal neurons is about 1 μm^2. $K_{pump} = 1 \mu M$ for the calcium-ATPase pump. The sodium-calcium exchanger will have higher values for V_{max} and K_{pump}.

It is worth noting that slightly more complicated models of calcium flux may not be distinguishable from simpler models where sufficient measurements are not available to tease apart the contributions to the net calcium dynamics from the many different terms.

Intracellular calcium can be sequestered into stores in such structures as the endoplasmic reticulum (ER), with release into the cytoplasm being mediated by second-messenger pathways. Stores can have large capacity, with estimates of Ca^{2+} in stores ranging from 100 μM to 5 mM. Uptake and release of calcium from these stores can result in intracellular calcium waves, or oscillations, on the timescale of seconds. It may be desirable to take recordings from neurons over many seconds in order to resolve the processes underlying calcium stores.

Uptake into the stores is via a calcium-ATPase pump in the smooth ER membrane, known as the SERCA pump. Since it binds two calcium ions for each ATP molecule, it can be described by Michaelis-Menten kinetics with a Hill coefficient of 2:

$$J_{SERCA} = V_{SERCA} \frac{Ca^2}{K_{SERCA}^2 + Ca^2}$$

where V_{SERCA} is again the membrane pump velocity and K_{SERCA} is the equilibrium ratio of the product of unbound substrate and enzyme to the bound complex, or equivalently, the pump affinity for calcium. A limitation of this approach is that the uptake flux does not depend on the intrastore calcium concentration, even though empty stores have a higher uptake rate.

Another mechanism for uptake into the ER is through a special set of pumps which are formed when calcium levels within the ER become low and directly connect the ER membrane to the outside of the cell. These are called store operated calcium pumps and at the time of writing, not much is known about them. They would likely have large pump velocity and lower affinity. The form for these pumps, called store operated calcium pumps, is the following:

$$J_{soc} = v_{soc}\frac{K_{soc}}{K_{soc} + Ca_{ER}} \tag{2.27}$$

The ER membrane contains calcium channels that are activated by calcium itself, leading to calcium-induced calcium release (CICR). The two major channels contributing to this effect are the ryanodine receptors and inositol 1,4,5-triphosphate (IP_3). The IP_3 receptors require the binding of both calcium and IP_3, while the ryanodine receptors (RyR) require only calcium.

The two types of receptor contribute to CICR in different ways for different neurons. In dendritic spines, RyR contribute more to the dynamics than IP_3Rs in hippocampal neurons, while the reverse is true of Purkinje cell spines.

A simple formulation of calcium release through the RyRs involves modelling the activation of the RyRs through a hill function of cytoplasmic calcium. Since calcium is not modulated instantaneously by rapid changes in calcium, one can assume that R relaxes to its steady state value. The calcium flux through the ER membrane due to RyRs is then given by:

$$J_R = V_R R(Ca_{cyt})(Ca_{ER} - Ca_{cyt})$$

with

$$\frac{dR}{dt} = \frac{R_\infty - R}{\tau_R}$$

and

$$R(Ca_{cyt}) = \frac{Ca_{cyt}^n}{K_R^n + Ca_{cyt}^n}$$

Calcium buffering

Calcium interacts with buffers, including fluorescent dyes such as Fura-2, which bind calcium and are thus calcium buffers that strongly affect the level of free calcium. It may be desirable for a model of intracellular calcium based on fluorescent dye measurements to account for this buffering effect explicitly.

The interaction of cytosolic calcium and a buffer B in a well mixed pool is given by the kinetic scheme:

$$Ca + B \leftrightarrow CaB$$

As measurements of calcium concentration near SK or BK channels is not yet available, it may only be desirable to explicitly model buffering processes in the bulk cytoplasm. Then for a particular well-mixed cellular compartment this leads to the system of coupled ODEs:

$$\begin{aligned}\frac{dCa}{dt} &= -k^{+}[Ca][B] + k^{-}[CaB] + \ldots \\ \frac{dB}{dt} &= \frac{dCa}{dt} \\ \frac{dCaB}{dt} &= -\frac{dB}{dt} = -\frac{dCa}{dt}\end{aligned}$$

where a single pool has been assumed, eliminating the need to consider diffusion rates of calcium and free/bound buffer.

As buffers bind calcium, their affinity is defined by $K_B = k^{-}/k^{+}$. A slow buffer with a high affinity (smaller K_B) may have less of an effect on calcium transients observed, such as those appearing in independent pools for SK and BK models where buffering processes are modeled implicitly. However, it will speed the return to equilibrium calcium concentration in the cytosol. On the other hand, a fast buffer with a low affinity (high K_B) will greatly reduce transients, perhaps eliminating the effects of SK or even BK currents altogether, while the return to equilibrium will be about as slow in the cytosol, perhaps slower than in the case with no buffer. This is due to the slow equilibration throughout the compartment and return to resting calcium levels determined by the time course of unbinding of calcium from a buffer free to diffuse from locations of high calcium concentration in the submembrane shell to locations of lower calcium concentrations in the bulk cytosol.

As buffers affect the calcium dynamics in the cell, they may be candidates for therapeutic strategies. The advantage of modeling a buffer explicitly is that it can be removed to better reflect the situation for free calcium transients in buffer free media in situations that cannot be explored experimentally.

Rapid Buffer Approximation

The rapid buffer approximation is used when the buffering of calcium is fast so that the calcium and buffer concentrations are essentially always in equilibrium with each other, so that the effects of the buffer can be accounted for easily as:

$$\frac{dCa}{dt} = \frac{J_{Ca}}{1+\kappa}$$

where

$$\kappa = \frac{K_B [B]_{tot}}{(K_B + Ca)^2}$$

which indicates how much calcium entering a cellular compartment becomes bound to the buffer. When calcium concentrations are much less than the dissociation constant K_B, so that the affinity of the buffer is low, this ratio is $\kappa = [B]_{tot}/K_B$.

Relating Fluorescence to Calcium Concentration

It is assumed that the fluorescence f is related linearly to the free calcium indicator dye concentration F and the calcium-bound dye concentration CaF via:

$$f = S_F[F] + S_{FCa}CaF = S_F[F]_{tot} + (S_{FCa} - S_F)CaF$$

where the total dye concentration is $[F]_{tot} = F + CaF$ and the coefficients S_F and S_{FCa} specify the contribution from the dye's unbound and bound forms, respectively. The fluorescent intensity of saturated dye (i.e., maximally bound by calcium so that [F] can be neglected compared to [CaF]) is $f_{max} = S_{FCa}[F]_{tot}$. The intensity in minimal calcium is $f_{min} = S_F[F]_{tot}$.

It is additionally assumed that the rapid buffer approximation holds. Given this approximation and the expressions for f, f_{max}, and f_{min} above, the free calcium concentration can be derived as:

$$Ca = K_F \frac{f - f_{min}}{f_{max} - f}$$

In practice, it is difficult to measure f_{min} as calcium cannot be entirely removed from the experimental tissue slice preparation. Defining $R_f = f_{max}/f_{min}$ of the indicator, the dynamic range, the intracellular calcium concentration is:

$$Ca = K_F \frac{f/f_{max} - 1/R_f}{1 - f/f_{max}}$$

If R_f is large, as it is for dyes such as Fluo-3 and Fluo-4, then $1/R_f$ is much smaller than f/f_{max} and does not much influence the calculated value of Ca^{2+}.

Changes in Ca^{2+} from a baseline can be estimated. With f_{res} as the fluorescence at the resting baseline, $\Delta f/f = (f - f_{res})/f_{res}$, with maximum value $\Delta f_{max}/f = (f_{max} -$

$f_{res})/f_{res}$:

$$\Delta Ca = K_F(1 - 1/R_f)\frac{(\Delta f_{max}/f_{res} + 1)|\delta f}{(\Delta f_{max}/f_{res} - \Delta f/f_{res})\Delta f_{max}}$$

$$Ca_{res} = K_F\frac{(1 - 1/R_f)}{\Delta f_{max}/f_{res}} - \frac{K_F}{R_f}$$

IP$_3$ Model

On relatively simple model for the IP$_3$ dynamics is the one of . A potentially undesirable simplification of is that it neglects the production and degradation of IP$_3$ and how that might contribute to the dynamics. A model from contains these terms. The model is designed for understanding calcium dynamics in glial cells, but the authors claim the model should be applicable to other kinds of cells as well. The variables in the models are the intracellular Ca^{2+} level C, the fraction of inactive IP$_3$ receptors h, and the available IP$_3$ concentration I. A term modeling the contribution of glutamate signals to IP$_3$ production may be included as an additional production term. Calcium-induced Ca^{2+} release in the model is controlled by the interplay of two effects: flow of calcium from the ER to the cytoplasm mediated by Ca^{2+} dependent opening of the IP$_3$ receptor channels and flow of calcium into the ER due to the action of sarco-endoplasmic reticulum Ca^{2+} ATPase (SERCA) pumps. The equilibrium calcium levels in the cytoplasm are set by a balance of Ca^{2+} leak from the ER, SERCA uptake, and plasma membrane Ca^{2+} transport. At these levels, slight increases in cytoplasmic Ca^{2+} levels leads to slight CICR. The opening probability of IP$_3$ receptors depends nonlinearly on the Ca^{2+} concentration. CICR increases dramatically when Ca^{2+} levels rise sufficiently but reverses at high cytoplasmic Ca^{2+} concentrations when inactivation of IP$_3$ channels occurs. SERCA pumps also tend to bring down cytoplasmic Ca^{2+} levels, increasing their sequestering of Ca^{2+} into the ER lumen at elevated Ca^{2+} levels. When the cytoplasmic Ca^{2+} levels return to baseline, IP_3 channels deinactivate.

The equation for cytoplasmic calcium (C) balance is given by:

$$\frac{dC}{dt} = J_{IP_3} + J_{leak} - J_{pump} \tag{2.28}$$

As in the case of the RyR model, SERCA pump rate has a Hill rate expression with exponent 2:

$$J_{pump} = v_{pump} \frac{C^2}{C^2 + K_p^2} \tag{2.29}$$

The nonspecific Ca^{2+} leak current is proportional to the Ca^{2+} gradient across the ER membrane:

$$J_{leak} = v_{leak}(C_{ER} - C) \tag{2.30}$$

IP_3 receptor channels are modeled as ensembles of four independent subunits with three binding sites each: one for IP_3 and two for Ca^{2+}. Ca^{2+} sites include an activation site and a separate site for inactivation. The channel is open when IP_3 and Ca^{2+} are bound to a fixed set of three out of four subunits.

Then the form for the calcium current through the IP 3 receptors (IP_3 Rs) is:

$$J_{IP_3} = v_{IP_3} m_\infty^3 n_\infty^3 h^3 (C_{ER} - C) \tag{2.31}$$

where $m_\infty = \text{Hill}(I, d_1)$, $n_\infty = \text{Hill}(C, d_5)$, and h account for the three gating reactions, IP_3 binding, activating Ca^{2+} binding, and Ca^{2+} dependent inactivation of the receptor. I is the intracellular IP_3 concentration and v_{IP_3} is the maximum channel permeability. The calcium concentration in the ER is:

$C_{ER} = \frac{1}{r}(C_{tot}C)$ where C_{tot} is the total calcium concentration and r is the volume ratio between the ER and the cytosol volumes.

This equation is coupled with an equation for h that accounts for the kinetics of IP_3Rs:

$$\frac{dh}{dt} = \frac{h_\infty - h}{\tau_h}$$

where:

$$h_\infty = \frac{Q_2}{Q_2 + C}, \tau_h = \frac{1}{a_2(Q_2 + C)}, Q_2 = d_2\frac{I + d_1}{I + d_3}$$

According to [3], IP_3 is produced by hydrolysis of PIP_2 by PLCβ and PLCδ. PLCβ pertains to glutamate-dependent IP_3 production, while PLCδ is activated by increased intracellular calcium levels. In [3], it is modeled as follows:

$$v_\delta = \frac{\bar{v}_\delta}{1 + \frac{I}{\kappa_\delta}}$$

where κ_δ is the inhibition constant of PLCδ activity. Two major IP_3 degradation pathways include phosphorylation and dephosphorylation of IP_3. Both are modeled as of Michaelis-Menten type:

$$v_{5P} = \bar{v}_{5P}\mathrm{Hill}(I, K_5) \tag{2.32}$$

$$v_{3K} = \bar{v}_{3K}\mathrm{Hill}(C^4, K_D)\mathrm{Hill}(I, K_3) \tag{2.33}$$

$$\tag{2.34}$$

Then IP_3 balance is modeled as:

$$\frac{dI}{dt} = v_\delta - v_{5P} - v_{3K}$$

The contribution of glutamate signals to IP_3 production can be taken into account as well. It is modeled as:

$$v_\beta = \bar{v}_\beta R(\gamma, C)$$

where γ is the extracellular glutamate concentration. $\bar{v}_\beta$ is the maximal PLCβ rate that depends on the surface density of mGluRs and $R(\gamma, C)$ is the fraction of activated (bound) mGluRs. PLCβ activity also depends on intracellular calcium concentration when [Ca^{2+}] ¿10 μM . However, this is typically outside of the physiological range, so this intracellular [Ca^{2+}] dependence is not included in this model. $R(\gamma, C)$ is modeled with a Hill-binding reaction scheme with an exponent ranging between 0.5 and 1. In [3], the exponent is 0.7.

$$R(\gamma, C) = \mathrm{Hill}(\gamma^0.7, K_\gamma(\gamma, C))$$

$K_\gamma(\gamma, C)$ is a term lumping two pathways underlying termination mechanisms of PLCβ signaling. One is involved in the reconstitution of the inactive G-protein heterotrimer due to intrinsic GTPase activity of activated Gα subunits. The second is PKC phosphorylation of the receptor, or of the G protein, or of PLCβ, or some combination thereof.

The effective Hill midpoint of $R(\gamma, C)$ increases as PLCβ termination takes over:

$$K_\gamma = K_R(1 + \frac{K_p}{K_R}\mathrm{Hill}(\gamma^0.7, K_R)\mathrm{Hill}(C, K_\pi))$$

K_R is the Hill midpoint of glutamate binding with its receptor whereas K_p measures the increment of the apparent affinity of the receptor due to PLCβ terminating signals.

The purpose of including this more complicated term modeling the contribution of glutamate to IP$_3$ production is only for completeness. It isnt clear yet whether such a term will be necessary for the purposes of data assimilation when quantitative measurements of calcium concentration become available, especially in a model which excludes the contributing effects of synapses to neuronal dynamics.

Explicit Calcium Dynamics Model

The sections preceding can be combined into a set of equations describing the intracellular calcium dynamics of the bulk cytosol and ER. Such a detailed model requires many more measurements of the calcium dynamics inside a cell than are currently available. We omit the contribution from IP$_3$, as it adds another level of complexity due to the difficulty of obtaining measurements associated with this signalling pathway, but it is straightforward to include this term in the model for completeness as well.

$$J_{soc} = v_{soc}\frac{K_{soc}}{K_{soc} + Ca_{ER}}$$

$$J_{SERCA} = v_{SERCA}\frac{Ca_{cyt}^2}{Ca_{cyt}^2 + K_{SERCA}^2}$$

$$J_{PMCA} = v_{PMCA} \frac{Ca_{cyt}^2}{Ca_{cyt}^2 + K_{PMCA}^2}$$

$$J_{Na-Ca} = v_{Na-Ca} \frac{Ca_{cyt}}{Ca_{cyt} + K_{Na-Ca}}$$

$$J_{leak,PM} = v_{leak,PM}(Ca_{ext} - Ca_{cyt})$$

$$J_{leak,ER} = v_{leak,ER}(Ca_{ER} - Ca_{cyt})$$

$$J_{RyR} = v_{RyR} R(Ca_{ER} - Ca_{cyt})$$

$$\frac{dR}{dt} = \frac{R_\infty - R}{\tau_R}$$

$$R_\infty = \frac{Ca_{cyt}^n}{Ca_{cyt}^n + K_{RyR}^n}$$

$$J_{buf,PM} = k_{B,PM}^+ B_{eq,PM}(Ca_{eq,PM} - Ca_{cyt})$$

$$J_{buf,ER} = k_{B,ER}^+ B_{eq,ER}(Ca_{eq,ER} - Ca_{ER})$$

Here the excess buffering approximation is used to approximate the effect of the endogenous buffers, where the concentration of buffer is very large compared to the concentration of free calcium in the cytosol or ER. The reaction of buffers and calcium is also assumed to be fast, so that calcium and buffer is always in equilibrium. B_{eq} and Ca_{eq} are the equilibrium concentrations of the buffer and free calcium, respectively. k_B^+ is the rate constant for the forward reaction of buffer and free calcium into the state where calcium is bound to the buffer.

$$\kappa = \frac{K_F F_{eq}}{(K_F + Ca_{cyt})^2}$$

$$J_{in,F} = \frac{J_{in}}{1 + \kappa}$$

Here $J_{in,F}$ and the binding ratio κ describes how the rapid buffering of the fluorescence dye attenuates the influence of calcium influx J_{in} due to VGCC, NMDA receptors, and other sources, on the free calcium concentration. The dynamics of calcium in the ER and cytosol are then:

$$\frac{dCa_{cyt}}{dt} = J_{in,F} + J_{buf,PM} - J_{Na-Ca} - J_{PMCA} - J_{SERCA} + J_{RyR} + J_{leak,PM} + J_{leak,ER}$$

$$\frac{dCa_{ER}}{dt} = J_{soc} + J_{buf,ER} + J_{SERCA} - J_{RyR} - J_{leak,ER}$$

This model contains terms accounting for most known contributions to net calcium dynamics, including CICR, thought to be altered in Alzheimer's disease. For the purposes of teasing apart different mechanisms underlying the pathology, it is necessary to include all of them in a detailed model. As measurements of calcium concentration involves measuring fluorescence levels, it is necessary to include terms accounting for the effects of fast dyes used in this process which alter the ordinary calcium dynamics. Until chapter 6, we will not attempt a further investigation involving the detailed model, instead choosing to ask a more modest set of questions which can be answered with simpler neuron models and measurements that are presently available.

Chapter 3

Inferring Membrane Dynamics and Channel Kinetics of a Neuromorphic Integrated Circuit

Emulating the biophysics of neuronal dynamics and behavior offers a principled analysis-by-synthesis approach towards understanding mechanisms of brain function. Here we report on a set of procedures assimilating and emulating neurobiological data on a neuromorphic very large-scale integrated (VLSI) circuit. The analog VLSI chip, Neurodyn, features 384 digitally programmable parameters specifying for 4 generalized Hodgkin-Huxley neurons coupled through 12 conductance based chemical synapses, reversal potentials, conductances, and spline regressed gating variables. In one set of experiments, we assimilated membrane potential recorded from one of the neurons on the chip to the model structure upon which NeuroDyn was designed and the known current input sequence, arriving at the programmed parameters except for model errors due to analog imperfections in the chip fabrication. In a related set of experiments, we replicated songbird individual neuron dynamics on NeuroDyn by estimating and configuring parameters extracted using data assimilation (DA) from intracellular neural recordings. Faithful emulation of detailed biophysical neural dynamics will enable the use of NeuroDyn as a tool to probe electrical and molecular properties of functional neural circuits, with neuroscience applications in studying the relationship between molecular properties of neurons and the emergence of different spike patterns or different brain behaviors, and clinical applications in studying and predicting effects of neuromodulators or neurodegenerative diseases on ion channel kinetics.

The analogy that exists between the kinetics of biological channels and semiconductors can be exploited to design electronics which emulate neural circuits. The analogy is in the voltage-dependent electron/hole channels in silicon transistors and voltage dependent ion channels in biology. The Boltzmann distribution describing both the hole/electron energy in silicon transistors and the channel activation in biological ion channels scales exponentially in response to the gate and membrane voltage respectively. This analogy can be used to construct electrical circuits which have the exponential current-voltage characteristic.

Neuromorphic engineering [55] pursues the design of integrated electronic systems that physically emulate the function and structure of biological neural systems driven by two complementary but synergistic objectives: the engineering of naturally intelligent systems for perception and computation that approach the robustness, noise resilience and energy efficiency of their counterparts in biology; and the science of progressing towards a more fundamental understanding of the cognitive function of the brain [13]. These two objectives are jointly pursued through *analysis-by-synthesis* as the combination of top-down deconstruction and bottom-up construction of physical models of brain function, in the spirit of Richard Feynman's famous words: "*What I cannot build, I do not understand.*"

Despite formidable advances in the engineering of neuromorphic silicon models of perception and cognition, the “morphing” of form and structure from neurobiology to silicon integrated circuits has mostly been a qualitative analysis-by-synthesis endeavor, with systematic quantitative methods for their precise alignment pursued only in a few instances, *e.g.*, [71]. This work pursues the application of a proven systematic quantitative method, *data assimilation* (DA) [1], to analysis-by-synthesis in neuromorphic engineering by aligning the dynamics of biological and model neuronal state variables in mapping the biophysics onto finely tuned equivalent physics in the silicon emulation medium, illustrated in Fig. 3.1.

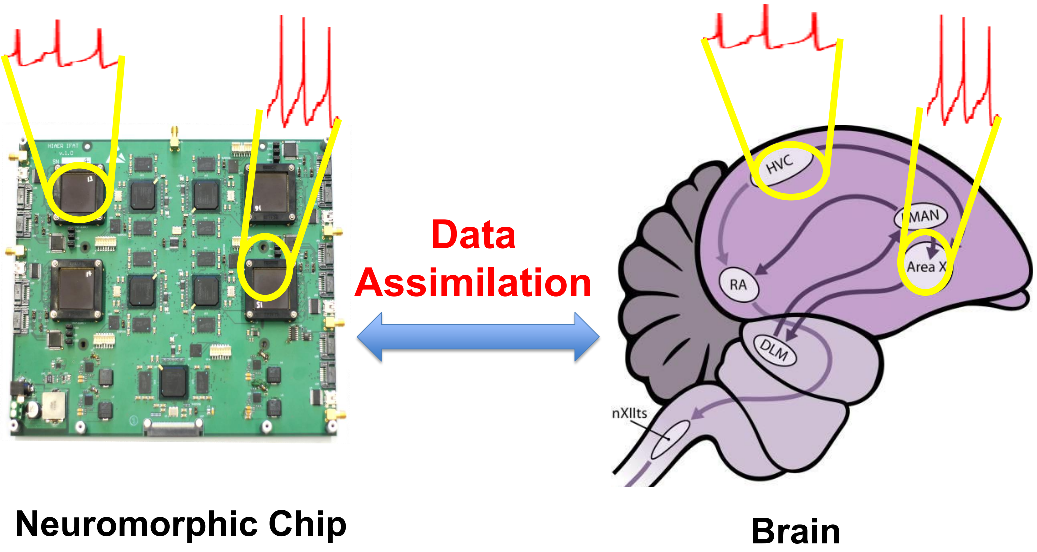

Figure 3.1: Motivation: using the physical medium of silicon neurons to emulate biological neuronal dynamics.

Formulating a proper model to emulate multiple types of neurons is a critical step in the synthesis. However, realizing the complex functional form of membrane currents and channel variables is difficult, especially in analog circuits. This has motivated alternative realizations by simplifications in the model. The prevailing approach has been to abstract the neuron membrane action potential to discrete-time spike events in simplified models that capture the essence of integrate-and-fire dynamics and synaptic coupling between large numbers of neurons in an address-event representation. These approaches may lead to highly efficient and densely integrated implementations in analog very-large scale integrated (VLSI) silicon, *e.g.*, [65]. However, to examine effects of neuromodulators, neurotoxins, and neurodegenerative diseases on ion channel kinetics, and to accurately emulate behavior of different types of neurons, a more sophisticated and flexible model is necessary. This has motivated the custom design of specialized and yet highly flexible neuromimetic analog integrated circuits capable of emulating the detailed and parameterized continuous-time dynamics of neuronal and synaptic state variables, *e.g.*, [71].

NeuroDyn [96, 97] is such an analog very large-scale integrated (VLSI) circuit instantiation of a general continuous-time model of biophysical neuronal dynamics on a small-scale, 4-neuron 12-synapse network. NeuroDyn features 384 digitally programmable parameters, specifying for each neuron and synapse the reversal potentials, conductances, and spline-regressed voltage dependence profile of opening and closing rates of the gating variables. These parameterized characteristics in NeuroDyn provide the capacity to emulate a large variety of neuron and synapse behaviors, which in turn requires the complex task of adjusting a large number of parameters. A relatively simple calibration and parameter fitting procedure proved adequate to set parameters in the biophysical model approximately to desired values[96] and even to generate phasic and tonic bursting in an extended Hodgkin-Huxley model formalism[97]. This was only possible by tuning each of the internal variables in the dynamics in isolation based on detailed model knowledge. In contrast, this luxury cannot be afforded in the more complex settings typical in experimental neuroscience that require inferring neural form and structure from very limited data recorded from very sparse locations in the brain. Thus, a more systematic and powerful method for arriving at values for parameters in such complex models and porting them to highly parameterized neuromorphic emulation platforms, accounting for substantial uncertainties in the modeling and noise in the observations as well as sources of imprecision and transistor mismatch in the physical emulation platform, is highly desirable.

Data assimilation (DA) methods have been applied to model estimation and time series prediction in biological neural systems [56, 61]. Extending the model estimation to silicon neural systems [91], DA accomplishes the alignment of biological and silicon physics, inspiring confidence in the implementation of functional neural circuits in the silicon medium. In this work we map biophysical neural function onto NeuroDyn, enabling the task of programming its parameters; then use NeuroDyn to predict the behavior of a biological neuron. The scheme is illustrated in Fig. 3.2.

The inference problem of DA is formulated as nonlinear optimization over a high-dimensional path integral and has been explored both in its exact and approximate form on various chaotic and neural models [88, 45, 2, 60, 56, 44, 40, 7]. We find that DA is capable of estimating parameters in a model of biophysical neural dynamics in data recorded from a songbird HVC neuron, mapped itself onto a model describing the dynamics of physical state variables in the chip. When these parameters are programmed onto the chip, it successfully emulates the voltage time series recorded from the songbird HVC neuron.

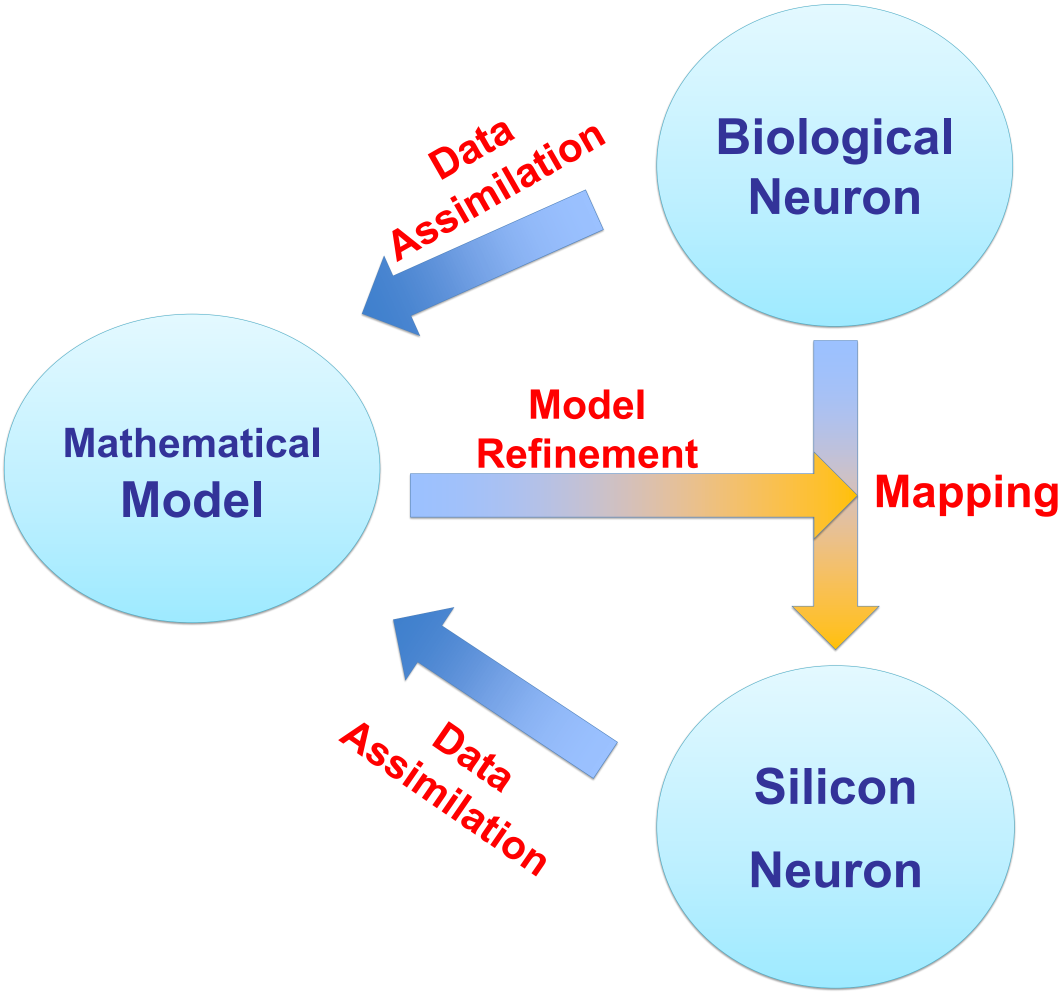

Figure 3.2: Scheme of this work. Harnessing DA to map a biological system to a neuromorphic silicon chip, through independent characterization of each system [91] and model integration to complete the mapping. Previous notable work in fitting neural data onto a mathematical model [2] and characterization of neuromorphic hardware [60] are provided in context.

The analog VLSI design of the NeuroDyn system was presented in [96]. Initial results of applying DA of songbird HVC neural data to the NeuroDyn chip were presented in [91]. This chapter extends this work and contributes a systematic method to correct for analog mismatch in the NeuroDyn chip leading to demonstration of precise mapping from recorded HVC dynamics to equivalent neuronal state variable dynamics recorded on the chip. This chapter is organized as follows: the NeuroDyn chip model and a new method for correcting analog mismatch based on DA are presented in Section 3.1; results of applying the proposed DA methods on synthetic benchmark data, measured NeuroDyn data, and recorded HVC neuron data in mapping the HVC neuron dynamics on the NeuroDyn chip are presented in Section 3.2; extensions of the DA methods and application to dynamically interactive neural prostheses are discussed in Section 3.2.6; and concluding remarks are offered in the last Section.

3.1 NeuroDyn Model

3.1.1 Generalized Model of Biophysical Neural Dynamics

The Hodgkin-Huxley (H-H) model has been a *de facto* standard in biophysical modeling of single-unit neural dynamics, described by a system of four differential equations in the membrane potential V and three gating variables h, m, and n [32]. These four variables interact to give the neuron its excitable dynamics, generating action potentials in response to external or synaptic current stimuli I_{Inj}. The H-H model specifies precise equations governing the kinetics of voltage-gated channel opening and closing in the membrane conductances, which Hodgkin and Huxley derived by curve-fitting detailed measurements on the giant squid axon with astounding accuracy [32]. However, these kinetics depend on properties of membrane ion channels that are highly variable across neuronal types and species, calling for greater flexibility in their functional form.

NeuroDyn [96, 97] implements an extended form of H-H dynamics and rate-based kinetics with general parameterized voltage dependence of the opening and closing rates in the ion channel gating variables. It also provides for general parameterized rate-based synaptic coupling between neurons [96], which is not investigated here. Each neuron $i = 1, \ldots 4$ undergoes membrane dynamics of the form:

$$C_{mem} \frac{dV_i}{dt} = -I_{Na,i} - I_{K,i} - I_{L,i} + I_{Inj,i} \tag{3.1}$$

with membrane capacitance C_{mem}, injected current $I_{Inj,i}$, and sodium, potassium and leak conductance-based currents of the approximate form:

$$\begin{aligned} I_{Na,i} &= m^3 h \, g_{Na,i} \, (V_i - E_{Na,i}), & (3.2) \\ I_{K,i} &= n^4 \, g_{K,i} \, (V_i - E_{K,i}), & (3.3) \\ I_{L,i} &= g_{L,i} \, (V_i - E_{L,i}). & (3.4) \end{aligned}$$

In turn, the dynamics of the gating variables h_i, m_i and n_i are described by rate-based kinetics of the form:

$$\frac{dx_i}{dt} = \alpha_{x,i}(V_i) \, (1 - x_i) - \beta_{x,i}(V_i) \, x_i \tag{3.5}$$

where the notation x stands for each of h, m, and n. Unlike the specific voltage dependence of the rate kinetics in the standard H-H formulation, the opening rates $\alpha_{x,i}$ and closing rates $\beta_{x,i}$ in NeuroDyn depend on membrane voltage V_i in general parameterized form regressed as 7-point additive spline sigmoidal functions:

$$\begin{aligned} \alpha_{x,i}(V_i) &= \sum_{k=1}^{7} \alpha_{x,i,k} \, \sigma_k(V_i) & (3.6) \\ \beta_{x,i}(V_i) &= \sum_{k=1}^{7} \beta_{x,i,k} \, \sigma_k(V_i) & (3.7) \end{aligned}$$

with predetermined sigmoids[1]:

$$\sigma_k(V) = \frac{1}{1 + e^{\pm\kappa(V_{b,k} - V)/V_T}} \tag{3.8}$$

at uniformly spaced centers spanning the voltage range:

$$V_{b,k} = V_{b,min} + \frac{k-1}{6}(V_{b,max} - V_{b,min}). \tag{3.9}$$

3.1.2 Mixed-Signal VLSI Circuit Implementation

$$\begin{aligned} I_{Na} &= 2I_{g_{Na}m^3h} \tanh(\frac{\kappa}{2V_T}(V_m - E_{Na})) \qquad (3.10) \\ I_{g_{Na}m^3h} &= I_{g,Na}(\frac{I_m}{I_{ref,m}})^3 \frac{I_h}{I_{ref,h}} \qquad (3.11) \end{aligned}$$

$$\begin{aligned} I_K &= 2I_{g_K n^4} \tanh(\frac{\kappa}{2V_T}(V_m - E_K)) \qquad (3.12) \\ I_{g_K n^4} &= I_{g,K}(\frac{I_n}{I_{ref,n}})^4 \qquad (3.13) \end{aligned}$$

$$\begin{aligned} I_L &= 2I_{g_L} \tanh(\frac{\kappa}{2V_T}(V_m - E_L)) \qquad (3.14) \\ I_{g_L} &= I_{g,L} \qquad (3.15) \end{aligned}$$

---

[1] The polarity (±) in the exponent is programmed as either $+1$ or -1 through an additional binary parameter for each $\alpha_{x,i}$ and $\beta_{x,i}$. This supports either a monotonically increasing or a monotonically decreasing voltage profile for each of the opening and closing rates.

Reversal potentials E_{Na}, E_K and E_L are described by

$$E_{rev} = I_{vol,rev}(e_{rev}/1024)R_{rev} + V_{ref} \tag{3.16}$$

where e_{rev} is a programmable integer which can be programmed from 0 to 1023 to configure reversal potentials.

The general form of $I_{g,Na}$, $I_{g,K}$ and $I_{g,L}$ is represented by

$$I_{g,gate} = I_{mas,gate}(g_{gate}/1024) \tag{3.17}$$

where g_{gate} is also a programmable integer from 0 to 1023. Thus,

$$\frac{d}{dt}X = \frac{I_\alpha}{C_{gate}V_T}(1-X) - \frac{I_\beta}{C_{gate}V_T}X \tag{3.18}$$

$$X = \frac{I_{gate}}{I_{ref,gate}} \tag{3.19}$$

where X is a placeholder for the gating variables n, m and h. $V_T = kT/q$ is the thermal voltage. $\frac{I_\alpha}{C_{gate}V_T}$ and $\frac{I_\beta}{C_{gate}V_T}$ function as opening and closing rate in the standard H-H model.

In contrast to the specific heuristic equations for the opening and closing rates of the kinetics in the gating variables, as arrived at through intricate curve fitting in the original H-H model [32], NeuroDyn permits a more general form regressed as 7-point additive spline sigmoidal functions:

$$I_{rate}(V) = \sum_{k=1}^{7} I_{bk}\, I_{\sigma,k}(V) = \sum_{k=1}^{7} \frac{I_{bk}}{1+e^{\pm\kappa(V_{bk}-V)/V_T}} \tag{3.20}$$

where the output current I_{rate} denotes either one of the I_α and I_β rates, and V_{bk} are fixed uniformly spaced voltages described by Eqn.3.22. I_{bk} is describe in Eqn.3.23.

$$V_{b1} = V_{ref} - R_0 I_{vol,bias} \tag{3.21}$$

$$V_{bk} = V_{b1} + R_{bias,k} I_{vol,bias} \tag{3.22}$$

Here, $k = 2...7$.

$$I_{bk} = I_{mas,rate_{gate,k}}(rate_{gate,k}/1024) \tag{3.23}$$

$rate_{gate,k}$ can be $\alpha_{gate,k}$ or $\beta_{gate,k}$ which are digital values for the 7-point additive spline sigmoidal functions. There are seven digital values for each of 3 opening and 3 closing rates, namely 42 configurable variables in total.

All the digital values are converted to analog values by a current DAC. V_{ref} is a reference voltage. All the reversal potentials are shifted by this V_{ref}. So are the V_{bk}.

Therefore, take 42 $I_{mas,rate}$, 3 $I_{mas,gate}$, 3 $I_{vol,rev}$, 1 $I_{vol,bias}$, 3 $I_{ref,gate}$, 1 C_{mem}, 3 C_{gate}, 7 R_{bias} and 1 injected current scaling factor K into consideration, there are 64 parameters to estimate if we ignore the 45 κ and 45 thermal voltage V_T. All the $I_{mas,X}$, $I_{vol,X}$, and $I_{ref,X}$ are mirrored from one of three current sources respectively.

3.1.3 Twin Experiments

To build confidence in the ability of our algorithm to return the correct values of unknown parameters and states on a given system with sparse measurements, we attempt experiments in which we have as much knowledge and control of the system and experimental data as possible. We generate synthetic data so that the system, including all unknown states and parameters, are known to the experimenter. The experimenter compares the output of the algorithm with the true underlying dynamics. Because the model is exact and mirrors the true system, we call such an experiment a “twin experiment” [1].

In a real experiment, the experimenter will not know the value of most parameters and states. The experimenter will usually have data for measured states for times greater than the end of the estimation window. When this is the case, the experimenter can test the algorithm’s estimates by integrating the model forward using the estimated dynamical map from the configuration of the system at the end of the estimation window. This is called the model prediction. The prediction is compared with the measurements to evaluate the quality of state and parameter estimates within the estimation window. Good ’fits’ to the data are easy to achieve, and are not considered a good measure of the quality of the estimated model.

In neural systems, injected currents that are high enough in amplitude to drive spiking behavior, long enough in duration to sample all the degrees of freedom of a model neuron, and low enough in frequency to not be absorbed into the *RC* time constant of the membrane are necessary for the algorithm to succeed in a twin or real experiment. This is a result of the fact that nothing can be inferred about a component of a process if it is not influencing the behavior of the system in a data set. In order for the dynamical map to be inferred from the data, all of its degrees of freedom must be activated for the algorithm to have a chance to infer values of parameters and unknown states that generalize outside of the training set [31]. A chaotic stimulus waveform of sufficiently low frequency, such as the trajectory of one of the states generated from the Lorenz ‘63 dynamical equations, will satisfy the above conditions.

3.2 Experiments and Results

3.2.1 Data Assimilation with Synthetic Data

The experiments we conducted confirm the power of the DA method to perform parameter estimation and predict neuronal dynamics from limited data, both in computer simulation and in measurements on the NeuroDyn chip. To establish a baseline for the DA experiments with the NeuroDyn chip, a twin experiment of DA in software with synthetic data was conducted first, in which the data was generated to fit the model of the NeuroDyn chip perfectly. We used the theoretic model, (3.5)-(3.20), extracted from the chip, and a chaotic time-series current waveform I_{Inj} to generate the time-series of membrane potential $V(t)$. We then used the DA method of Section II to estimate the model parameters. The results shown in Fig. 3.3 suggest that when the measurement functions h_l are known and the model of the physics on the chip is without error, sufficient measurements can be obtained from the chip for the DA algorithm to correctly estimate all unknown parameters.

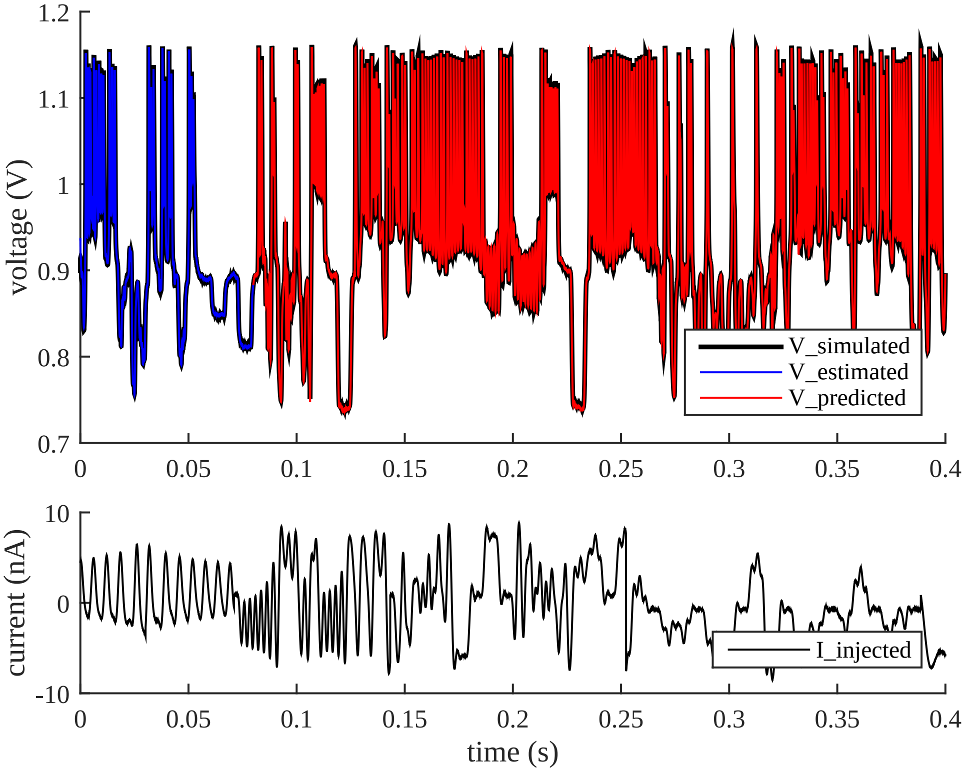

Figure 3.3: Estimated (blue), predicted (red), and simulated (black) data of the NeuroDyn model in software. The prediction is obtained by integrating the equations of motion forward using the estimated parameters and configuration of the system at the end of the estimation window. The injected stimulating current protocol I_{Inj} (bottom graph) is also displayed.

3.2.2 Data Assimilation with NeuroDyn

In subsequent experiments DA was applied to estimate parameters in the NeuroDyn model from measurements on the NeuroDyn chip. A similar complex stimulating current waveform I_{Inj} was applied to the Neurodyn chip to elicit dynamical waveforms, which where then used as the data for DA of the NeuroDyn model. Data obtained from the chip yielded time series measurements of $[V(t), m(t), h(t), n(t)]$ at a time resolution of 500 kHz. Details on the experimental prototols used to obtain the recordings of these time series are presented in [91].

The results of DA are displayed in Fig. 3.4. Although the estimates and predictions are quite good, there is some disagreement between the theoretical values given the model of the chip and the configured values, shown in Table 3.1. It is desirable to check for discrepancies between configured parameter values and theoretical parameter values obtained with DA. Two factors mainly result in the discrepancies: inaccuracy of the model describing circuits implemented in NeuroDyn, and mismatch during fabrication. In the next section, a systematic means to correct for model error and mismatch using DA is introduced.

3.2.3 Model Error and Mismatch Correction

As all the $I_{mas,X}$, $I_{vol,X}$, and $I_{ref,X}$ are mirrored from corresponding current sources by current mirrors, to predict the membrane dynamics and channel kinetics accurately, the mismatch between current mirrors should be estimated first. Thus a general first order linear equation is proposed as (3.24) to describe the mismatch of the current mirrors.

$$I_{X,real} = I_{X,ideal}(1 + \varepsilon_X) \tag{3.24}$$

$I_{X,ideal}$ is the theoretical current mirrored from current source, and $I_{X,real}$ is the real current which takes mismatch ε_X of current mirrors into consideration. To avoid degeneracy, if more than one independent parameters multiply together in the model, they would be estimated as one parameter. For instance, $A * B$ will be replaced with C if A and B are independent parameters. Take estimation of reversal potential as an example, when mismatch is considered, (3.16) is modified as (3.25). Instead to estimate all the three parameters $I_{vol,rev}$, R_{rev} and ε_{I_vol}, we only estimate ε_{I_vol} and assign theoretical value to $I_{vol,Rev}$ and R_{rev}. In this way, degeneracy is avoided and fewer parameters are required to be estimated. This also speeds the process of estimation.

$$E_{rev} = I_{vol,rev}(1+\varepsilon_{I_vol})R_{rev}(e_{rev}/1024)+V_{ref} \tag{3.25}$$

As a result, we need to estimate two distinct sets of parameters. The first set of parameters are represented as ε_{I_vol} in (3.25). The other set are digitally programmed onto the NeurDyn chip, as e_{rev} in (3.25).

Fundamentally, the first set is specific to the chip and its analog mismatch, whereas the second set is specific to the neural data and the model. In principle, the two DA estimations can be completely decoupled: the estimation of the first set should be done separately for each chip independent of the particular neural data that it is tasked to emulate; and the estimation of the second set can be done for each new data set independent of the chip that will target its implementation.

The result of estimating first set of parameters is shown in table 3.2. The 'remaining' parameters noted in the table could not be estimated, as DA results are not sensitive to them. To configure NeuroDyn chip, they were assigned zero. Parameter set one then were incorporated to refine the mathematical model. This refined model was used to conduct DA on neurobiological data to estimate the programmable parameters. The results are presented in sections 3.2.4 and 3.2.5.

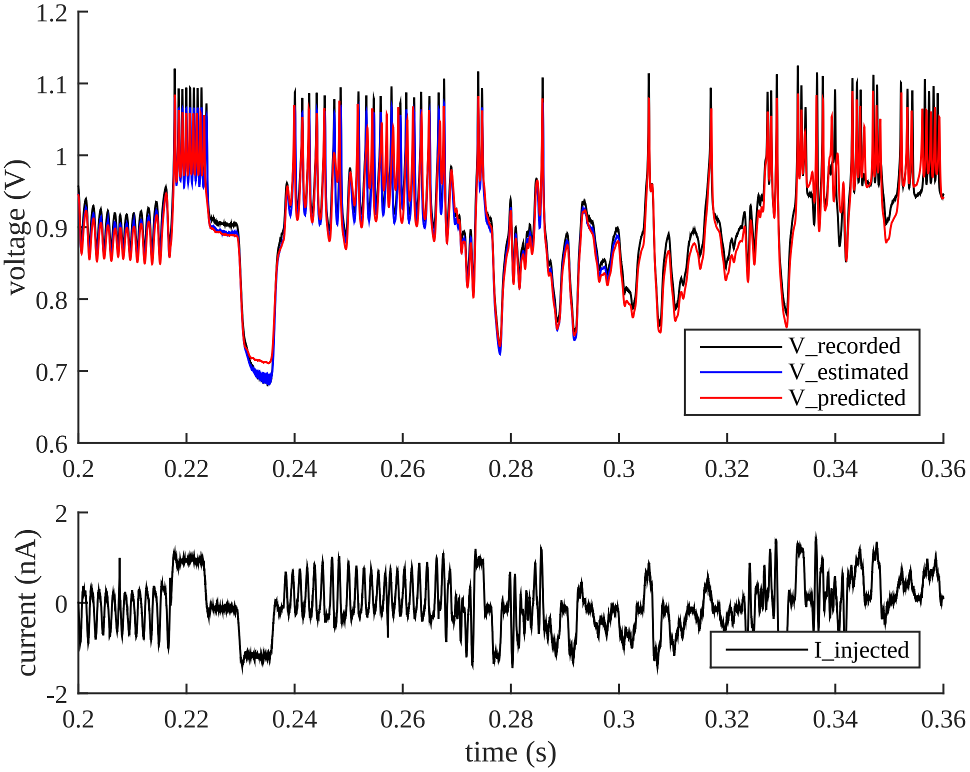

Figure 3.4: Estimated (blue), predicted (red), and measured (black) data recorded from the NeuroDyn chip. The estimation and prediction were obtained, and the curent injection (bottom graph) was applied, under identical conditions as in Fig. 3.3.

3.2.4 Data Assimilation with Biological Neuron Data

Using the above DA method, we proceeded to assimilate voltage data obtained from zebra finch HVC_I neurons[45] to the model of the physics on the chip. Integrating forward the state of the model with the obtained parameter set, the resulting waveform matches the recorded voltage data within the intrinsic variability of the neuron. This is displayed in Fig. 3.5. The activation and inactivation of channels are shown in Fig. 3.6.

This result is notable because to our knowledge such a simple HH model with only I_{Na}, I_K and I_L has previously been insufficient to accurately describe features including spike timing and amplitude, AP shape, and subthreshold variations. Typically a number of sodium, potassium, and calcium currents are included in conductance based models as

Table 3.1: Reference and Estimated Values

Params	Ref	Estimated	Params	Ref	Estimated
g_{Na}	600	294.81	α_{h5}	0	4.3293e-6
e_{Na}	450	408.09	α_{h6}	0	3.0360e-6
g_K	160	100	α_{h7}	0	3.0332e-6
e_K	200	234.40	β_{h1}	0	10
g_L	12	2.61	β_{h2}	0	2.8240e6
e_L	250	235.88	β_{h3}	0	1.0368
k_{inj}	1	0.08	β_{h4}	0	3.0750e-5
σ	13.5	15	β_{h5}	41	79.9999
V_{b1}	0.61	0.61	β_{h6}	25	49.9737
V_{step}	0.123	0.13	β_{h7}	8	48.7581
α_{m1}	0	0.0003	α_{n1}	0	4.9057e-5
α_{m2}	0	0.0001	α_{n2}	0	0.0374
α_{m3}	120	222.943	α_{n3}	0	1.3169e-6
α_{m4}	400	306.742	α_{n4}	0	1.2238
α_{m5}	800	600.001	α_{n5}	18	49.9987
α_{m6}	1023	800.211	α_{n6}	5	1.5156
α_{m7}	1023	809.773	α_{n7}	43	18.9959
β_{m1}	1023	1000.14	β_{n1}	1	0.0414
β_{m2}	1023	1022.93	β_{n2}	0	0.0009
β_{m3}	1023	1023	β_{n3}	0	2.3919
β_{m4}	1023	800	β_{n4}	1	8.3258e-6
β_{m5}	0	9.9998	β_{n5}	0	2.2087e-5
β_{m6}	0	9.9999	β_{n6}	0	0.0067
β_{m7}	0	9.9999	β_{n7}	1	2.6981
α_{h1}	237	120	I_{master}	100	83.5093
α_{h2}	5	1	I_{volt}	232	228.9760
α_{h3}	7	1.0000	V_{offset}	1	0.9913
α_{h4}	6	14.4858	R_{gate}	1.53e+06	1.5373e+6

well as a number of spatial compartments. These have been tuned by hand [21, 59, 89] or by using other approaches including exhaustive grid, stochastic, and evolutionary search algorithms [25, 68, 22, 70, 31, 11, 56]. The ability to successfully assimilate data to such a simple model may be attributable to our methods of DA and/or to the flexible implementation of opening rates α and closing rates β in NeuroDyn's highly parameterized model.

Table 3.2: Intrinsic parameters and mismatch

Intri + Mismatch	Estimated	Intri + Mismatch	Estimated
e_{ImNa}	4.95	e_{Iah3}	1.12
e_{IvNa}	-0.34	e_{Ibh4}	-0.99
e_{ImK}	4.95	e_{Ibh5}	1.65
e_{IvK}	3.00	e_{Ibh6}	0.44
e_{ImL}	1.62	e_{Ian4}	-0.98
e_{IvL}	4.23	e_{Ian5}	-0.99
$offsetI$	5.45	e_{Ian6}	-0.98
e_{Iam4}	-0.72	e_{Ian7}	2.35
e_{Iam5}	-0.98	e_{Ibn1}	2.78
e_{Ibm3}	4.55	e_{Ibn2}	1.93
e_{Ibm4}	-0.97	e_{Ibn3}	2.75
e_{Ibm5}	-0.48	$remaining$	N/A

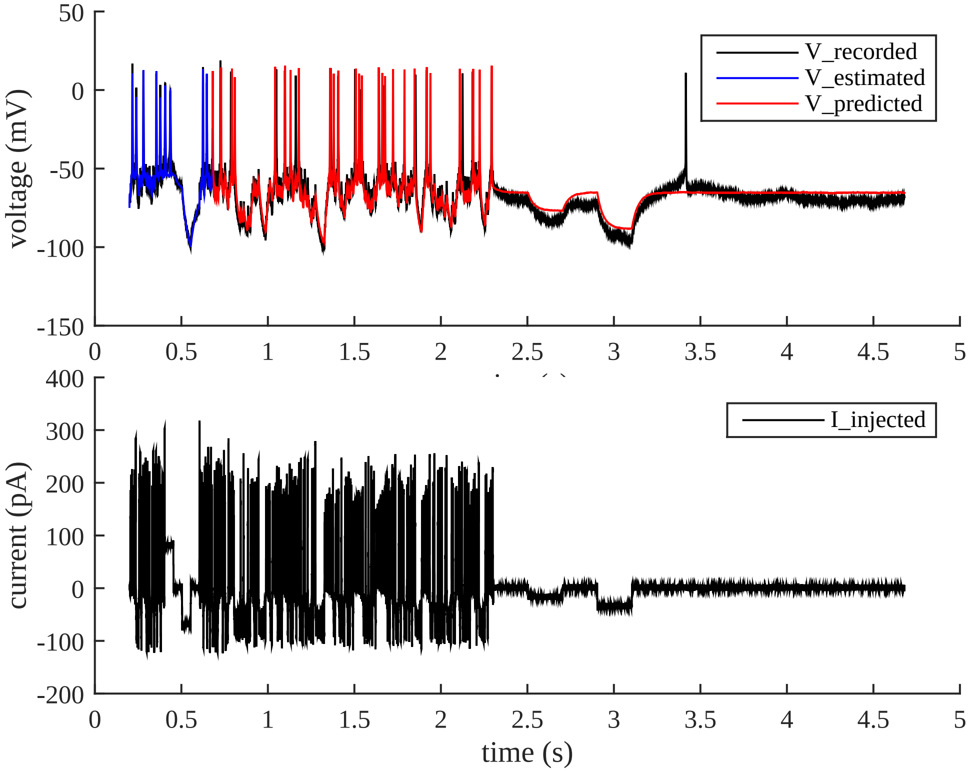

Figure 3.5: Estimated (blue), predicted (red), and electrophysiological (black) data recorded *in vivo* from an HVC interneuron (HVC$_\mathrm{I}$) [45] and instantiated onto the NeuroDyn hardware model. The input current I_{Inj} injected into the systems (bottom graph) elicits an almost identical response.

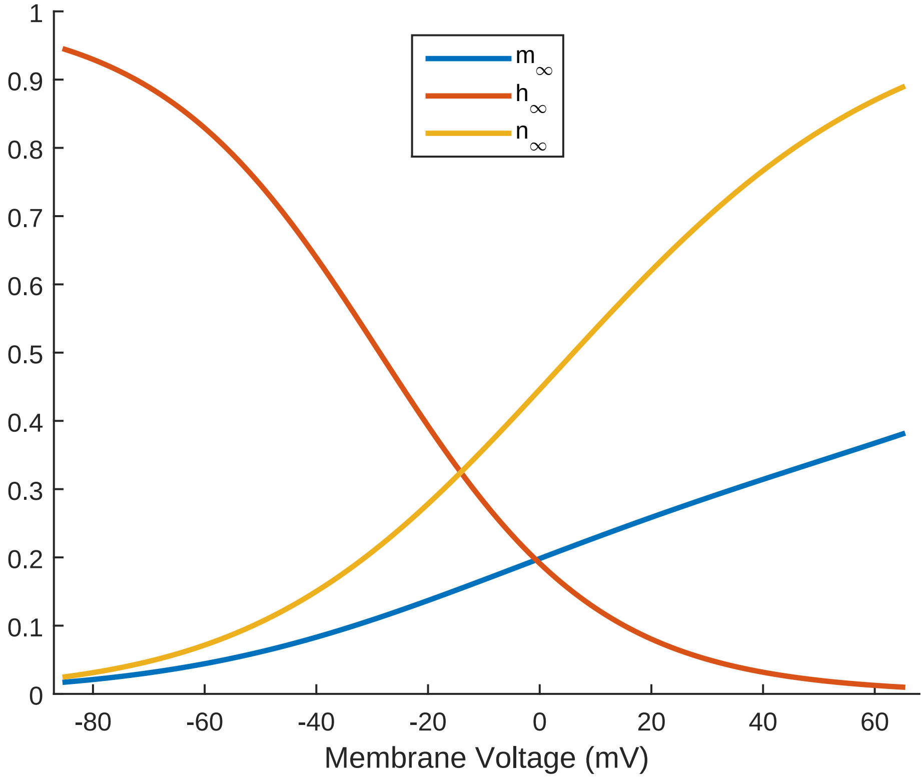

Figure 3.6: Estimated Activation and Inactivation of Channels.

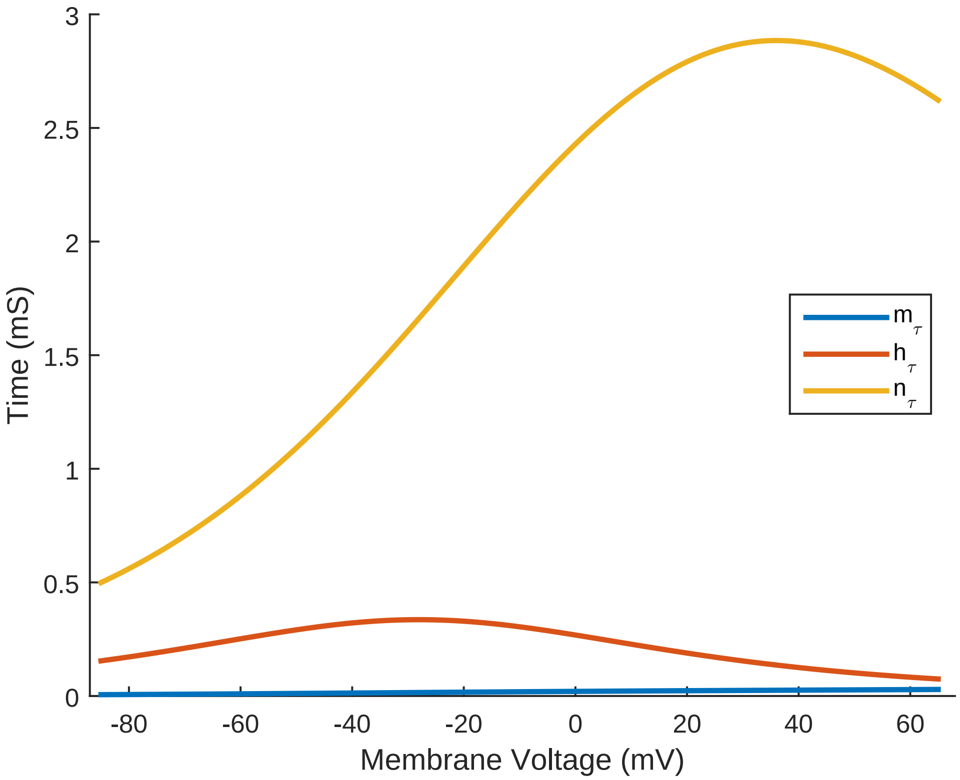

Figure 3.7: Estimated Voltage-dependent time constants.

3.2.5 Biological Neuron Emulation with NeuroDyn

Finally, the estimated parameters for HVC neuron were used to configure NeuroDyn chip. However, we still need to tune the estimated parameters to make NeuroDyn behavior close to HVC neuron. The value of estimated and actually used parameters is shown in table 3.3. The behavior of NeuroDyn chip is more sensitive to g_{Na}, e_{Na}, g_K, e_K, g_L and e_L, compared with αs and βs which are the amplitude of spline sigmoidal functions.

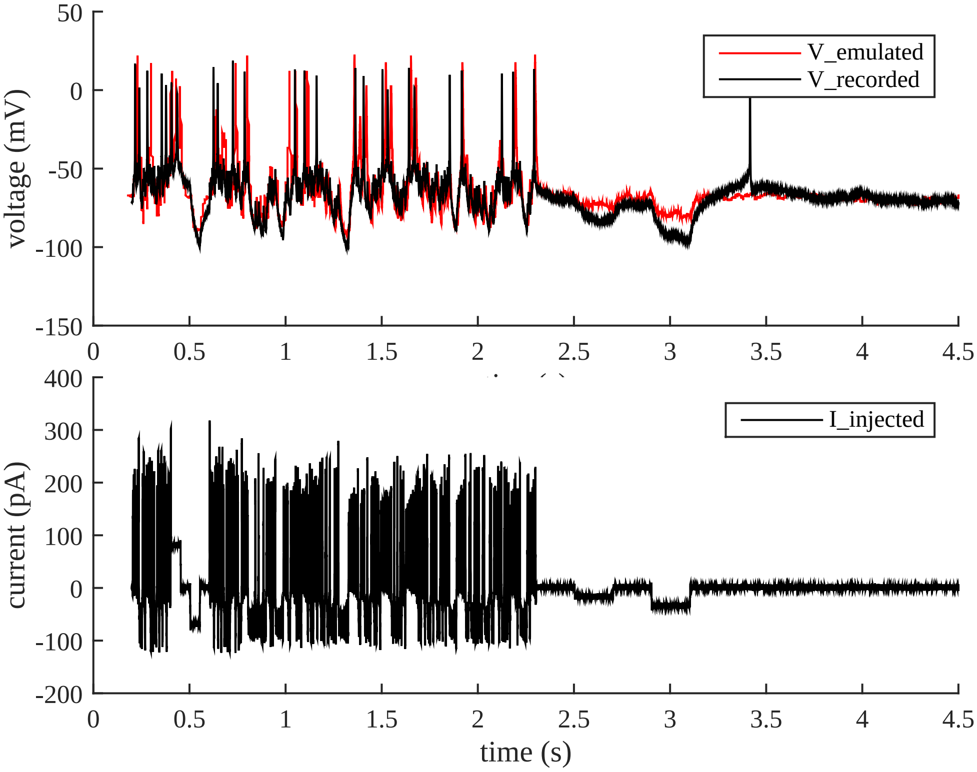

Figure 3.8: Emulated (red) data from an HVC interneuron (HVC_I) [45] and instantiated onto the NeuroDyn hardware model.

Table 3.3: Estimated Values and used values

Params	Estimated	Used	Params	Estimated	Used
g_{Na}	600.04	690	α_{h3}	111.38	100
e_{Na}	855.91	800	β_{h4}	29.96	20
g_K	701.26	800	β_{h5}	59.99	60
e_K	235.74	200	β_{h6}	42.16	40
g_L	37.05	40	α_{n4}	30.00	40
e_L	297.42	310	α_{n5}	100.03	100
α_{m4}	128.13	160	α_{n6}	300.53	300
α_{m5}	678.50	680	α_{n7}	350.83	350
β_{m3}	982.53	1000	β_{n1}	14.59	14
β_{m4}	302.97	300	β_{n2}	8.36	10
β_{m5}	119.99	98	β_{n3}	20.00	15
$offsetI$	5.45	N/a	$other\_\alpha, \beta$	N/A	0

3.2.6 Discussion

The formation of biohybrid circuits between biological and silicon neurons has enabled neuroscientists to investigate the interactions of cellular and network properties with an unprecedented level of experimental control inaccessible with traditional pharmacological and electrophysiological techniques. Over the last two decades, the coupling of biological and silicon neurons through an artificial synapse has allowed to dissect and reveal the role of various individual cellular and synaptic conductances in the activity of a single neuron or neuronal network in invertebrates [83, 77] and vertebrates [35, 48]. Analog silicon neuron designs and implementations ensure that the model neuron runs in real-time independently of the model complexity while offering a great level of flexibility for the configuration of different neuron types [97, 49]. While a high level of programmability is desirable, the configuration of analog silicon neurons can become problematic due to the inherent nonlinearities of the model neuron and the intrinsic VLSI process variability of the hardware implementation. Thus, automated parameter estimation and configuration of silicon neurons are needed, especially for extended dynamic clamp applications where more than one silicon neurons are couped with biological neurons [67].

Here, we have presented a DA procedure capable of tuning the parameters of a model of an analog VLSI chip emulating membrane dynamics and channel kinetics of generalized H-H neurons, and have shown that if all measurement functions h_l are known and the dynamical model of the chip is accurate, we can correctly estimate all of the parameters, many of which enter the equations nonlinearly. We have also shown that given noisy voltage data recorded from zebra finch $\mathrm{HVC_I}$ neurons and the relatively simple extended H-H model of the Neurodyn chip with only two ionic currents, I_{Na} and I_K, the time evolution of an $\mathrm{HVC_I}$ neuron can be accurately predicted. Finally, we have shown that NeuroDyn's model is approximately correct. With some additional manual adjustment, NeuroDyn's parameters can be tuned, starting from parameter estimates using DA, to achieve emulation of biological data. Potential discrepancies between theoretical values of parameters estimated using DA and configured parameters on Neurodyn have also been identified. Further work is needed to ascertain whether these discrepancies caused by mismatch of transistors can be resolved with improvements to the model, measurement functions h_l, and/or refinements to the DA procedure. We also found during DA on synthetic data and chip data that without prior constraints about the range of parameter values, many different parameter sets could be found which produced accurate estimations and predictions of all state variables which could not be distinguished from each other.

In the future, emulation of multiple biological neurons can be done in three steps. First, inject a standard current waveform to elicit a voltage response on each silicon neuron on the neuromorphic chips. Run the assimilation filter to estimate parameter set 1, parameterizing model error due to the fabrication mismatch. Second, run DA using neurobiological data using the corrected mathematical model describing silicon neurons on the neuromorphic chips. This yields the estimate of parameter set 2, the programmable parameters. Third, based on parameter set 1, parameter set 2 will be mapped onto neuromorphic chips. This will achieve emulation of the biological dynamics on silicon. The combination of DA and Neuromorphic chip will be a promising computational tool to emulate, communicate, and even control biological neurons in real-time. More accurate mathematical model and fast DA are necessary.

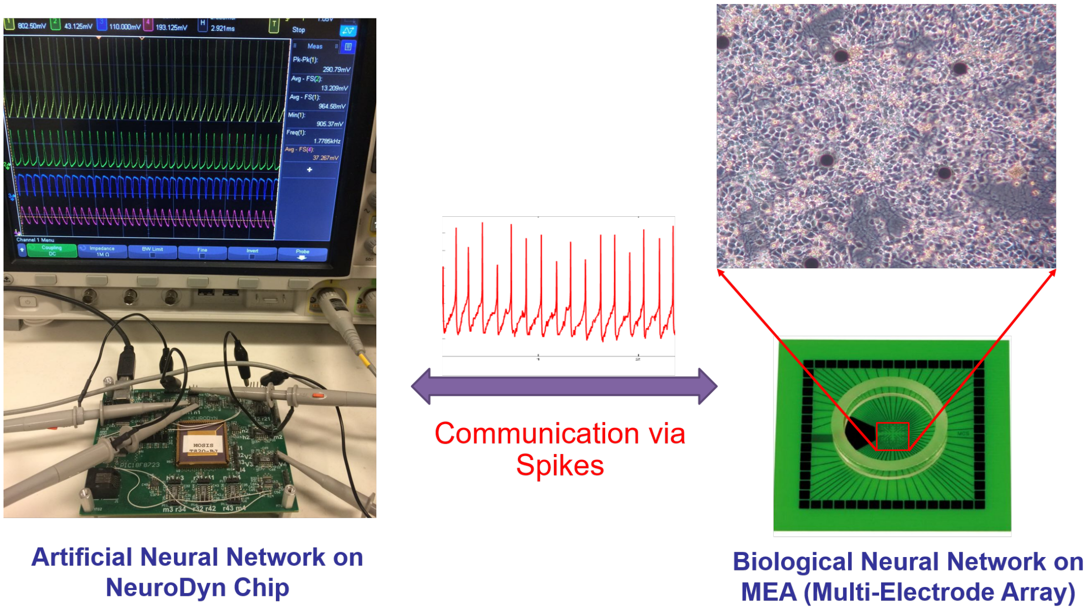

Figure 3.9: Dynamic clamp. (Left) Neuromorphic NeuroDyn chip and measurement PCB setup, (Right) Biological neurons on a multi-electrode array capable of recording and stimulation.

3.3 Conclusions

Importantly, we have successfully mapped biological dynamics to a silicon substrate leveraging the exponential voltage-current relationship present in biological ionic current kinetics and transistors. We have surmounted the fabrication mismatch in the NeuroDyn chip and estimated parameters in the NeuroDyn model, enabling emulation of biological data. We estimated the in/activation and voltage-dependent time constants of ion channels. These results demonstrate the analysis of biological neurons by synthesis and provide a basis for building biologically realistic network models in an integrated analog circuit chip which has the potential of forming closed-loop interaction between artificial and biological neural networks. The results may also help to understand the effects of neuromodulators or neurodegenerative diseases on ion channel kinetics, and may further provide insights into the relationship between molecular properties of neurons and the emergence of different spike patterns or different brain behaviors.

We have shown that the extended H-H model implemented on NeuroDyn is sufficient for describing one class of interneurons within the nucleus HVC, however there is a tremendous diversity of mechanisms in nervous systems. Such mechanisms include other ion channels such as calcium channels, neuromodulators, multi-compartmental dynamics through linear and nonlinear dendritic coupling, and intracellular signaling pathways. This motivates the design of neuromorphic chips containing more neurons which have more channels and each channel is flexible enough to be programmed to describe different ion kinetics.

Chapter 3 is reproduced in part from material as it appears in “Assimilation of Biophysical Neuronal Dynamics in Neuromorphic VLSI. Jun Wang, Daniel Breen, Abraham Akinin, Frederic Broccard, Henry DI Abarbanel, and Gert Cauwenberghs. The dissertation author was one of two first authors on this paper.

Chapter 4

An Underdetermined Problem: Inference on HVC Interneurons Using Data Assimilation

Male zebra finches sing a short, stereotyped motif for the span of their adult lives, a behavior that is learned over several months while they are juveniles. An aspirational goal of the present research is to find the single cell and network mechanisms underlying learned sequential behaviors with a biophysically grounded simulation of the song motor pathway. Conductance based models represent membrane dynamics in terms of ion currents through passive, voltage-gated, and ligand-gated conductances. The dynamics of these conductances can be expressed mathematically in terms of biophysical descriptions of specific channel gating kinetics. This provides a mechanistic link between molecular cell properties and behavior [4, 52]. Previous research strongly suggests that the membrane dynamics of single neurons within nucleus HVC are crucial mechanistic components encoding the zebra finch song motif [29, 34, 51]. Two classes of neuron constitute the building blocks of the motor pathway responsible for song production. The first is a class of neurons called $\mathrm{HVC_{RA}}$ neurons which send excitatory projections to the nucleus RA from HVC. These neurons conspicuously feature at most one burst per song motif. The other kind of neuron is called an HVC Interneuron ($\mathrm{HVC_I}$ neuron) and does not project to other nuclei, but instead inhibits neurons local to HVC.

In this chapter we implement our methods of statistical data assimilation (DA) to construct a model of the $\mathrm{HVC_I}$ neuron as an isolated building block. We demonstrate the power of DA to extract information about $\mathrm{HVC_I}$ neurons given a biophysical model and data. We discover a minimal set of necessary ionic currents and sets of parameters reproducing essential characteristic electrophysiological features. We also find evidence for a manifold of degenerate solutions to minimizing a cost function containing measurement noise and model error terms. This may also represent present limitations of our DA procedure due to model errors including missing currents, insufficient number of measurements, or mundane numerical issues such as ill conditioning. We do not attempt to study models of these cells within the context of a simulated network of neurons. Future work will be aimed at placing these model $\mathrm{HVC_I}$ cells as nodes within a network.

The mathematical form of our model is a Hodgkin Huxley (HH) conductance based model. The ion currents included are biologically grounded in pharmacological experiments and computer modeling . Our data assimilation algorithm estimates all of the unobserved states and parameters of the model conditioned on measurements recorded from $\mathrm{HVC_I}$ neurons in vitro. The present work is significant because it extends previous work [88, 45, 44, 56], demonstrating that the time course of several unobservable variables and dozens of parameters that enter the dynamical equations nonlinearly can be accurately estimated, using recently developed annealing methods of data assimilation [93]. It demonstrates the capability of the data assimilation methods to estimate *all* the parameters of conductance-based Hodgkin Huxley models, including those which enter nonlinearly such as the parameters describing the gating kinetics, given voltage recordings of real $\mathrm{HVC_I}$ neurons. Neuronal properties described by these parameters change over time [14], have a spatial dependence over the different parts of the neuron[4], while also varying among cells of the same class [73]. Because our method is capable of estimating all the parameters which enter into HH models, we can better characterize variability, such as the shape of IV curves, across neurons. Data assimilation also allows one to characterize properties of a model that would not otherwise be readily apparent, such as the existence of a manifold including multiple distinct model parameter sets with indistinguishable predictive power. We discuss such parameter set degeneracies for our $\mathrm{HVC_I}$ neuron model.

Our data assimilation algorithm formulates the problem as one of nonlinear optimization over a high dimensional path integral and has been explored both in its exact and approximate form on various chaotic and neural models [88, 45, 1, 44, 56, 95, 93, 37]. Experiments done on simulated systems in which all of the parameters and states are known to the experimenter we call *twin experiments*. The purpose of such experiments is to inform the use of the algorithm on real neural systems in order to estimate model parameters and unknown states which produce accurate predictions. In real physical systems, only sparse measurements are typically available and errors in the model and measurements are inevitable. If the data assimilation algorithm can recover unknown states and parameters in controlled conditions where the parameters and state of the system at all times are known, then we are more confident of its ability to do so with real systems.

First we motivate and validate our model with results of data assimilation on voltage recordings of $\mathrm{HVC_I}$ neurons *in vitro*, showing that it exhibits key qualitative features and biophysical mechanisms found in other work [47, 19]. We then show results of applying our data assimilation procedure on synthetic and real voltage recordings of single $\mathrm{HVC_I}$ neurons.

4.1 Methods

4.1.1 HVC$_\mathrm{I}$ Neuron Model

Our Hodgkin Huxley model for an isolated HVC$_\mathrm{I}$ neuron derives from the model of [19]. The current balance equation for the membrane voltage is:

$$\begin{aligned} C_m \frac{dV(t)}{dt} =& I_K(V(t)) + I_{Na}(V(t)) + I_{CaL}(V(t),[Ca](t)) \\ & + I_{CaT}(V(t),[Ca](t)) + I_A(V(t)) \\ & + I_{SK}(V(t),[Ca](t)) + I_{KNa}(V(t)) + I_H(V(t)) \\ & + I_{Nap} + I_L(V(t)) + I_{inj}(t) \end{aligned}$$

C_m is the capacitance of the membrane and I_{inj} is a custom built drive current. The ion currents are I_K, the delayed rectifier potassium current; I_{Na}, the inactivating sodium current; I_{CaL}, a high threshold L-type calcium current; I_{CaT}, a low threshold voltage gated calcium current; I_A, an A-type potassium current; I_{SK}, a small-conductance calcium activated potassium current; I_{KNa}, a sodium dependent potassium current; I_H, a hyperpolarization activated cation current; I_{Nap}, a persistent sodium current; and I_L, the leak current. Our analysis is restricted to isolated single HVC$_\mathrm{I}$ neurons, so we do not include a synaptic current term I_{syn}.

We reduce this model to a simplified form with suitability for insertion into nodes within a simulation of HVC in mind. A simpler model is also easier to understand, reduces the computational difficulty of data assimilation and the number of parameter set degeneracies, or model symmetries (to be discussed below), present given the limitations of our present computational resources and methods to transfer the information in the data to our model. We eliminate I_A, I_{SK}, I_{KNa}, and I_{Nap} because these currents have been shown to be small in HVC_I neurons in other work [19]. HVC_I neurons are highly excitable with high resting membrane potentials and very little afterhyperpolarization (AHP), so the absence of these currents can be partially understood as a consequence of I_A, I_{SK}, I_{KNa} being currents which depress the resting membrane voltage of neurons and contribute to AHP. I_{CaL} may contribute to the dynamics of HVC_I neurons, but we did not find a significant difference in the quality of data assimilation results when I_{CaL} was and was not present in addition to the remaining currents. We conclude that it can be justified for removal for our purposes, modeling the time evolution of the voltage dynamics for durations not exceeding $\approx$ 10 seconds, as not important in causing the most conspicuous features of HVC_I neuron electrophysiology: a hyperpolarization induced sag upon hyperpolarizing current injection and rebound spiking upon release from hyperpolarizing current injection. I_{CaT} and I_H are both critical mechanisms in this behavior as discussed in other work [19], so we retain them in our model.

Upon simplification our Hodgkin Huxley model for an isolated HVC_I neuron becomes:

$$
\begin{aligned}
C_m \frac{dV(t)}{dt} &= I_K(V(t)) + I_{Na}(V(t)) + I_H(V(t)) \\
&\quad + I_{CaT}(V(t),[Ca](t)) + I_L(V(t)) + I_{inj}(t) \\
\frac{d[Ca]}{dt} &= \phi I_{CaT}(V(t),[Ca](t)) + \frac{[Ca]_{eq} - [Ca](t)}{\tau_{[Ca]}} \\
\frac{dx}{dt} &= \frac{x_\infty(V) - x(t)}{\tau_x(V)}
\end{aligned}
\tag{4.1}
$$

The variables $n(t), m(t), h(t), H(t), a(t), b(t) \in x$ are gating variables which regulate the conductance of ions through the membrane of a neuron, described by the following equations:

$$
\begin{aligned}
x_\infty(V) &= 0.5(1 + \tanh(\frac{V - \theta_x}{\sigma_x})) \\
\tau_x(V) &= t_1 + t_2(1 - \tanh^2(\frac{V - \theta_{x_t}}{\sigma_{x_t}}))
\end{aligned}
\tag{4.2}
$$

The θ, σ, and t_i are the kinetic parameters, properties of individual neurons.

The ion currents have the following form:

$$\begin{aligned}
I_K(V(t)) &= g_K n(t)^4 (E_K - V(t)) \\
I_{Na}(V(t)) &= g_{Na} m(t)^3 h(t) (E_{Na} - V(t)) \\
I_H(V(t)) &= g_H H(t)^2 (E_H - V(t)) \\
I_{CaT}(V(t), [Ca](t)) &= g_{Ca_T} a(t)^3 b(t)^3 \Phi_{GHK}(V(t), [Ca](t)) \\
\Phi_{GHK}(V(t), [Ca](t)) &= V(t) \frac{[Ca]_{ext} \exp(\frac{-V(t)}{V_T}) - [Ca](t)}{1 - \exp(\frac{-V(t)}{V_T})}
\end{aligned} \tag{4.3}$$

The various g_X and E_X are the maximal conductances and reversal potentials of the ion currents.

$\Phi_{GHK}(V(t), [Ca](t))$ is the Goldman-Hodgkin-Katz equation for the ionic flux through the neuron membrane. Here it is used instead of the ohmic form to accurately model the calcium current.

Calcium appears as a dynamical variable in the model. The equations describing the calcium dynamics are informed by calcium ion conservation. These equations balance the change caused by calcium ion current influx and decay to equilibium calcium concentration.

$$\frac{d[Ca](t)}{dt} = \phi I_{CaT}(V(t), [Ca](t)) + \frac{[Ca]_{eq} - [Ca](t)}{\tau_{[Ca]}}$$

$[Ca](t)$ is the cytosolic, or internal, calcium concentration, $[Ca]_{eq}$ is the equilibrium cytosolic calcium concentration, $\tau_{[Ca]}$ is the time constant describing the rate at which the internal concentration of calcium tends towards its equilibrium concentration, $[Ca]_{ext}$ is the concentration of calcium outside the cell membrane, and V_T is the thermal voltage. With our HVC$_\mathrm{I}$ model, parameters governing only the time evolution of $[Ca](t)$ cannot be determined from data assimilation using a measurement of $V(t)$ alone, as $[Ca](t)$ is only weakly coupled to the dynamics of membrane voltage. Virtually no calcium is available inside the cell to flow outwards at physiological concentrations, so $\Phi_{GHK}(V(t),[Ca](t)) \approx \Phi_{GHK}(V(t))$.

HVC$_\mathrm{I}$ neurons recorded in vitro are highly excitable, have a slow increase in resting membrane potential in response to hyperpolarizing current injection, and fire rebound action potentials in response to release from hyperpolarizing current injection [47, 19]. With the support of pharmacological manipulation and computational modeling, [19] suggests the underlying mechanisms are two voltage gated ion currents, a hyperpolarization activated current I_H and a low threshold T type calcium current I_{CaT}. With these currents, our model reproduces the above qualitative features as shown in Figure 4.2. The parameters of the model in Table 4.1 exhibiting these features are derived from data assimilation on a real HVC$_\mathrm{I}$ neuron. Results of estimation and prediction on the voltage trace, with corresponding stimulating current, are plotted in Figure 4.1. As shown in Figure 4.2, when I_H alone is blocked, the sag is eliminated, but a delayed rebound spike is preserved. When both I_H and I_{CaT} are blocked, the sag and rebound spiking are eliminated. Additionally, when I_{CaT} is blocked, the membrane potential is depressed by about 10 mV, in agreement with experimental observations[19]. This demonstrates that an HH model with currents defined in Equation 4.3 is sufficient to describe experimentally obtained voltage recordings of HVC$_\mathrm{I}$ neurons.

4.1.2 Twin Experiments

In our twin experiments on the HVC$_\mathrm{I}$ model neuron, the voltage is the only measured variable, an expected experimental limitation on what can be measured in real neurons. To our synthetic data, Gaussian noise is added at each time point. We then run the data assimilation algorithm and compare the estimations of unknown parameters and unobserved states to their true values. Following this, the true values or trajectories are compared to the model predicted trajectories outside of the estimation window.

4.1.3 Data Assimilation on Real Data

A twin experiment informs us about what measurements are necessary in order for data assimilation on a real system to succeed. Our twin experiments on a single compartment HVC_I neuron, to be described in a later section, demonstrate the conditions that a time series measurement of only the voltage will suffice to produce accurate predictions. Bounds, such as those on kinetic parameters associated with I_{CaT}, were chosen to be around values reviewed in the literature [33]. Experimental voltage data was sampled at either 50,000 or 10,000 Hz. Lower time resolution has the disadvantage that less subthreshold information is available. However, because simulations are run on single computing nodes and therefore larger problems take a long time to run, an advantage of coarse grained time resolution is that currents stimulating more degrees of freedom in the model can be chosen. This also reduces the numerical difficulty of the optimization problem substantially. Bounds and estimates of the parameters are given in table 4.1. E_{Na}, E_K, E_H, V_T, and $[Ca]_{ext}$ are known for these neurons, so were fixed during data assimilation.

Table 4.1: Parameter estimates after 28 annealing steps using a voltage recording of an actual HVC_I neuron as input into the data assimilation algorithm. The predictions are good, and a few of the parameters are at or near the bounds, but most are somewhere in between. A '-' here denotes that a parameter was set at the indicated value because it is either fixed by experiment or well known (The various E_x's and the extracellular calcium concentration $[Ca]_{ext}$), for example.

Parameter	Lower Bound	Upper Bound	Estimate
$C_m(nF)$	0.01	0.033	0.0317

Table 4.1 continued.

Parameter	Lower Bound	Upper Bound	Estimate
$g_{Na}(\mu\text{S})$	0.01	10.0	0.63
$g_K(\mu\text{S})$	0.01	15.0	2.15
$g_H(\mu\text{S})$	0.0001	0.01	0.0032
$g_{Ca_T}(\mu\text{S}\mu\text{M}^{-1})$	0.00001	0.01	0.0064
$g_L(\mu\text{S})$	0.0001	0.01	0.0052
E_{Na} (mV)	-	-	55.0
E_K (mV)	-	-	-90.0
E_H (mV)	-	-	-40.0
E_L (mV)	-90.0	-30.0	-66.32
θ_H (mV)	-85	-55	-81.62
σ_H (mV)	-62.5	-5.0	-9.80
t_{H_1} (ms)	1.0	1000.0	214.39
t_{H_2} (ms)	10.0	2000.0	157.80
θ_{H_t} (mV)	-80.0	-40.0	-59.70
σ_{H_t} (mV)	-62.5	-5.0	-5.52
θ_a (mV)	-80.0	-30.0	-30.0
σ_a (mV)	5.0	62.5	32.9
t_{a_1} (ms)	0.01	5.0	4.44
t_{a_2} (ms)	1.0	20.0	4.24
θ_{a_t} (mV)	-80.0	-40.0	-55.12
σ_{a_t} (mV)	5.0	62.5	5.0
θ_b (mV)	-90.0	-60.0	-61.98
σ_b (mV)	-62.5	-5.0	-62.5

Table 4.1 continued.

Parameter	Lower Bound	Upper Bound	Estimate
t_{b_1} (ms)	0.01	10.0	2.90
t_{b_2} (ms)	1.0	100.0	7.57
θ_{b_t} (mV)	-90.0	-50.0	-59.6
σ_{b_t} (mV)	-62.5	-5.0	-15.1
θ_m (mV)	-50.0	-30.0	-32.304
σ_m (mV)	5.0	62.5	32.4
t_{m_1} (ms)	0.001	1.0	0.001
θ_h (mV)	-60.0	-20.0	-58.54
σ_h (mV)	-62.5	-5.0	-59.2
t_{h_1} (ms)	0.01	1.0	0.42
t_{h_2} (ms)	1.0	10.0	4.44
θ_{h_t} (mV)	-60.0	-20.0	-60.0
σ_{h_t} (mV)	-100.0	-5.0	-12.5
θ_n (mV)	-60.0	-20.0	-30.01
σ_n (mV)	5.0	62.5	62.5
t_{n_1} (ms)	0.01	1.0	0.01
t_{n_2} (ms)	0.1	10.0	10.0
θ_{n_t} (mV)	-60	-20	-30.79
σ_{n_t} (mV)	-100.0	-5.0	-37.7
$\phi(\mu\text{MnA}^{-1})$	0.01	10.0	3.88
τ_{Ca} (ms)	0.1	100.0	0.143
V_T (mV)	-	-	12.5
$[Ca]_{ext}(\mu\text{M})$	-	-	2500.0

Table 4.1 continued.

Parameter	Lower Bound	Upper Bound	Estimate
$[Ca]_0(\mu\text{M})$	0.01	5.0	1.11

4.2 Results

4.2.1 Data Assimilation for a Real HVC_I Neuron

As a preliminary analysis we attempted data assimilation with voltage recordings obtained from real HVC_I neurons *in vitro*. The estimates for parameters during this analysis were used to later generate synthetic data used in twin experiments to better understand the extent that our method of data assimilation is able to recover dynamical properties of single neurons. Our estimation window consisted of 16001 time points sampled at 10000 Hz. These data - the resulting estimation of the voltage, the comparison of the obtained prediction of the voltage with the actual recording outside of the training window, and the stimulating current used - are plotted in Figure 4.1. The parameter set obtained from data assimilation is included in Table 4.1. The results of prediction agree well with the data. As discussed in Section 4.2.2, qualitative features of HVC_I neurons including a sag in the voltage in response to hyperpolarizing current injection and rebound spiking upon termination of a hyperpolarizing stimulus were reproduced by this estimated parameter set. When I_H is blocked by setting $g_H = 0$ in the model and simulating the behavior of the neuron, the sag is eliminated and weaker rebound spiking is observed. When I_{CaT} is blocked in addition to I_H, neither the sag nor rebound spiking is observed, and the resting membrane voltage is depressed by about 10 mV, all in accordance with experimental observation [47, 19].

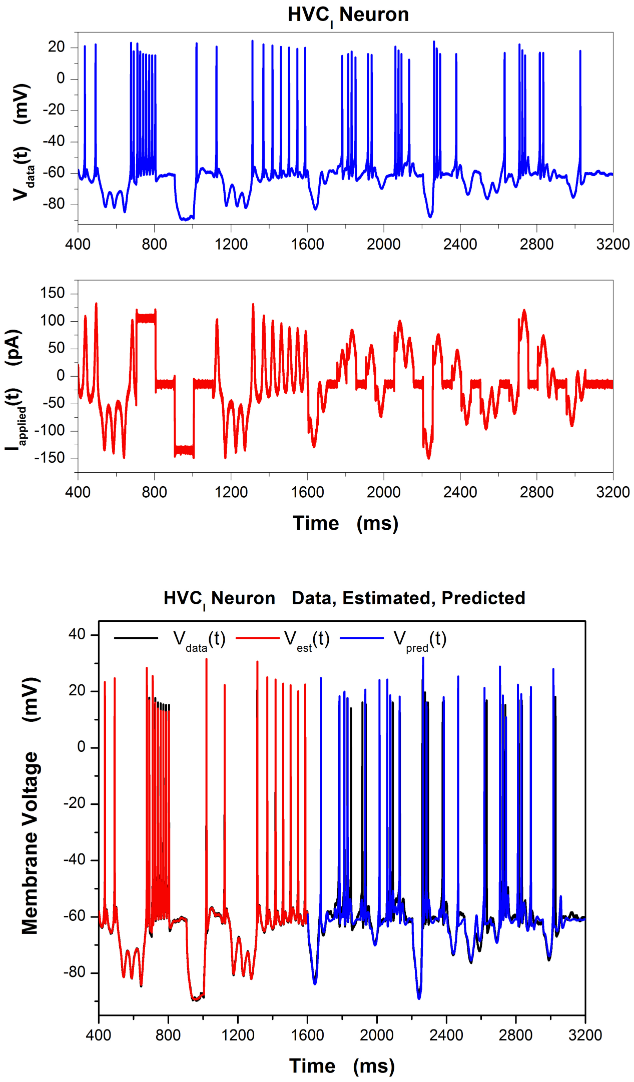

Figure 4.1: Top: Voltage trace (blue) and injected current (red) used in data assimilation, obtained from an HVC$_I$ recorded *in vitro*.
Bottom: Comparison of estimated voltage (red) and predicted voltage for times $t > t_M$ with the recording (black). The prediction is obtained by integrating the configuration of the system at the end of the estimation window $\mathbf{x}(t_M)$ forward with the estimated parameters. The agreement between prediction and observed data is excellent.

4.2.2 Qualitative Behavior of $\mathrm{HVC_I}$ Model

Previously performed experiments on $\mathrm{HVC_I}$ neurons in vitro show that a sag in the voltage appears when negative DC injected current is applied. When this negative DC current is removed, rebound spiking is observed [47, 19]. [19] also shows through pharmacological manipulation and computer modeling that the biological mechanisms responsible for this are a T-type calcium current and a hyperpolarization activated cation current. [33] describes features of T-type calcium currents in the central nervous system. Channels gating these currents are open only when their simultaneous activation and deinactivation are achieved. When the neuron is depolarized, I_{CaT} becomes increasingly inactivated and deinactivation requires a duration of quiescence or hyperpolarization. During hyperpolarization, I_H causes a sag in the voltage waveform to appear. Deinactivated I_{CaT} acts in conjunction with I_H when a neuron is released from hyperpolarization to cause rebound spiking.

$\mathrm{HVC_I}$ neurons are highly excitable, firing with high frequency in response to depolarizing current pulses. This may be in part due to the depolarizing influence of I_H and I_{CaT} which bring the neuron close to firing threshold. Large maximal conductances g_{Na} and g_K corresponding to the currents I_{Na} and I_K may also contribute to high excitability.

Table 4.1 shows the set of parameters that were used in generating Fig. 4.2. When developing a model, a central idea to data assimilation is that the form of the model must be determined by the modeler. However, once the model is specified, the values of the unknown parameters are to be determined by the algorithm.

I_H and I_{CaT} have the following form:

$$\begin{aligned} I_H(V(t)) =& g_H H(t)^2 (E_H - V(t)) \\ I_{CaT}(V(t),[Ca](t)) =& g_{Ca_T} a(t)^3 b(t)^3 \Phi_{GHK}(V(t),[Ca](t)) \end{aligned} \tag{4.4}$$

In Figure 4.2 we verified that our model produces the expected results: I_H induces

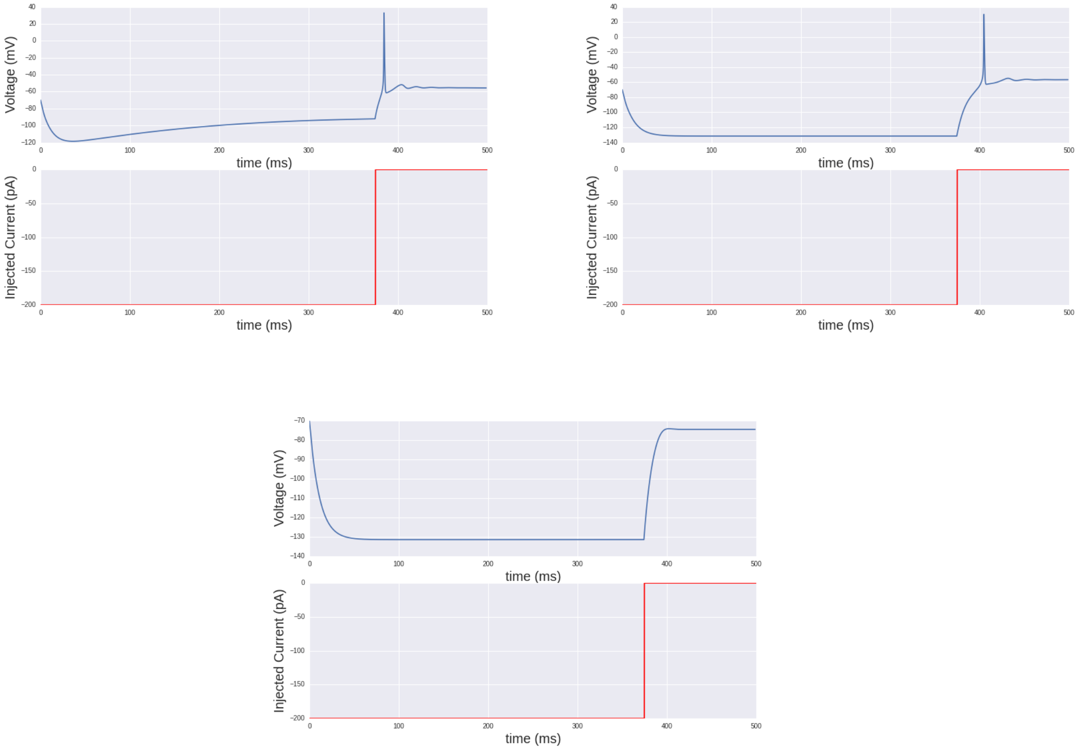

Figure 4.2: Qualitative features of the HVC_I model neuron with a parameter set (Table 4.1) obtained from data assimilation on a real HVC_I neuron. When I_H and I_{CaT} are not blocked, the model produces rebound spiking when released from hyperpolarizing current injection. When I_H is blocked but I_{CaT} is not blocked, the model produces weak rebound spiking while the sag disappears (top). When I_H and I_{CaT} are both blocked, neither rebound spiking nor a sag are seen (bottom). Additionally, the resting membrane voltage is depressed by about 10 mV, in accordance with experimental observations [19].

sag, while I_H and I_{CaT} both contribute to the rebound spiking. When g_H is set to zero, the sag is eliminated while weaker rebound spiking remains. When g_{Ca_T} is also set to zero, both the sag and rebound spiking are eliminated (Figure 4.2). This verifies the results of pharmacologically blocking I_H and I_{CaT} [19].

4.2.3 Real Data: Dependence of Data Assimilation Results on Stimulating Currents and Sample Rate

We contrasted the effects of different stimulating current protocols on the results of data assimilation and examined the structure of model ’symmetries’ given assimilated data collected from $\mathrm{HVC_I}$ neurons in vitro. A ’symmetry’ is the presence of multiple sets of parameters and initial conditions that produce similar voltage waveform behavior given a stimulating current. More precisely, integrating forward our $\mathrm{HVC_I}$ model using a number of estimated parameter sets produces indistinguishable (or nearly so) time evolution in the voltage given a stimulating current protocol. A number of measures, to be enumerated below, were used to analyze the structure in the model producing these symmetries given the assimilated data. To explore the tradeoffs given finite computational power between stimulating extra degrees of freedom and losing resolution of the measurements in time by downsizing the measured data, the same voltage traces were sampled at different frequencies. One set was sampled at 50 kHz, while the other set was sampled at 10 kHz. We used 24001 data points (480.02 ms) when analyzing the 50 kHz voltage trace and 10001 data points (1000.1) data points when analyzing the 10 kHz voltage trace. For either the 10 kHz or 50 kHz sampling rate conditions, 3 different stimulating current protocols were used to drive the voltage; a step current, a high frequency chaotic current, and a low frequency chaotic current. The quality of the resulting estimates were based on the predictive capabilities of the model estimations. In each case the lower frequency chaotic current protocol producing a voltage waveform with the highest quality predictions. The step current produced the lowest quality predictions, while the high frequency chaotic current performed in the middle. This is in contrast to the situation in twin experiments in section 4.2.5, where the step current and low frequency chaotic current stimulus fared the best, with the high frequency chaotic current giving the most mediocre results. This discrepency can be accounted for by the fact that the voltage waveform of the model neuron in twin experiments did not become driven to the subthreshold regime, while the voltage of the *in vitro* HVC$_\mathrm{I}$ neuron did become driven to this regime given the same stimulating current. Since I_{CaT} and I_H are both active in the subthreshold regime, the assimilation failed to produce a model which captured the effects of these currents when the synthetic voltage waveform was presented as data. Another possible factor is that due to the low pass filter properties of the differential equation giving the time evolution of the voltage, the added Gaussian noise in the synthetic voltage waveform in section 4.2.5 was so large that the signal to noise ratio was too small to inform the model in the case of a high frequency stimulating protocol. The step current protocol likely performed better with synthetic data due to the fact that the system generating the data is identical to the assimilated model, a simplification not available in the case of real physical systems. We now turn to the use of data assimilation in the analysis of real neurons.

When multiple excitatory ionic currents such as I_{CaT} and I_{Na} with similarly fast activation times are simultaneously present in a model, different combinations of I_{CaT} and I_{Na} can lead to voltage behavior which appears the same. There are a few other ways to characterize how estimates which produce similar and accurate predictions might be structured.

1. Plot the max attained amplitude and time averaged magnitude of individual ionic currents for each set of parameters as compared to the true value.
2. Compare the relative probability of each estimate producing accurate predictions by examining the distribution of the cost function.

Analysis of data will be provided in the same format for each stimulating current protocol and each of the sampling rate conditions as in experiments with synthetic data in section 4.1.2 below. In each section, the stimulating current and driven voltage waveform, the action level plot, box plots showing the distribution of the estimated average value and maximum magnitude of each of the theoretical ionic currents, and exemplar plots of the estimation and prediction beyond $t = t_M$ of the time evolution of the voltage waveform and theoretical ionic currents will be presented for each stimulating current protocol.

4.2.4 Important Comments

In the rest of the chapter, we present only the results of twin experiments and analyses done on synthetic data. The patterns we found and discuss below in the analyses of synthetic data were replicated on real data. We note here that due to the fact that real world neurons are much more complicated than our models, the interpretation of the meaning of estimated parameters and what constitutes accurate model predictions is less clear. For those results, the interested reader may refer to the archive preprint [7] which this chapter is based on.

We were also interested in what would happen if I_{CaL} were added to the model in addition to the other four ionic currents, I_{Na}, I_K, I_{CaT}, and I_H. The addition of I_{CaL} did not improve the quality of the estimations and predictions. Data assimilation with addition of I_{CaL} produced a variety of distinct parameter sets, but none of these parameter sets produced estimates differing significantly enough in the value of the cost function to dominate the contribution of the path integral. However, dropping I_{CaT} and retaining I_{CaL} eliminated the ability of the assimilated model to reproduce rebound spiking. This suggests that I_{CaL} is not important for reproducing observed behavior and therefore can be dropped from the model.

We also attempted data assimilation at a variety of sampling rates higher and lower than 10 kHz. For conditions in which the sampling rate was higher than 10 kHz, we increased the dimension D of the problem so that the estimation window was the same length in time for a 50 kHz and 25 kHz condition and allowed the algorithm to run as long as necessary to complete. Even with the condition of increasing the computational resources available for the problem, we did not find any improvement in the quality of predictions or undesirable multimodality of the probability distribution.

4.2.5 Twin Experiments: Dependence of Data Assimilation Results on Stimulating Currents and Sample Rate

To continue our investigation into parameter set degeneracies and how successfully our data assimilation methods produced accurate predictions when the sampling rate and stimulating currents varied, we compared predictions from 6 different conditions using synthetic data generated from our model with the estimated parameters of Table 4.1. For this analysis, 3 different current protocols from real experiments were used for data assimilation. As before, to explore the tradeoffs given finite computational power between stimulating extra degrees of freedom and losing resolution of the measurements in time by downsizing the measured data, the same voltage traces were sampled at different frequencies. One set was sampled at 50 kHz, while the other set was sampled at 10 kHz. The same stimulating current waveforms were used here as with DA on real data to demonstrate the presence of multiple sets of parameters and initial conditions that produce similar voltage waveform behavior when the system producing the data was known. Similar measures of these degeneracies were used here as were in the analysis of real data, to be enumerated again below. For all conditions, the length of the estimation window was 16001 time points. This is around the maximum possible number of time points which allows the estimation procedure to complete in 24 hours given constraints on our computational resources. For 10 kHz, this is a 1600.1 ms time window, while for 50 kHz, this is a 320.02 ms time window. For annealing, $\alpha = 2$ and β is incremented by 1 at each annealing step.

The stimulating current protocols, along with the response voltage and representative estimates and predictions for each protocol and sampling rate condition are shown in Figures 4.3-4.5.

Integrating forward our $\mathrm{HVC_I}$ model using a number of estimated parameter sets produces indistinguishable (or nearly so) time evolution in the voltage given a stimulating current protocol. For example, when multiple excitatory ionic currents such as I_{CaT} and I_{Na} with similarly fast activation times are simultaneously present in a model, different combinations of I_{CaT} and I_{Na} can lead to voltage behavior which appears the same. Because information about the flux of ionic currents is not typically available when taking recordings from real neurons, we first examine these potential shortcomings when estimating the accuracy of the inferred parameters in a single neuron model using synthetic data, where all the unrecorded processes are known to the experimenter but withheld from the data assimilation algorithm.

There are a few ways to characterize how estimates which produce accurate predictions might be structured.

1. Plot the max attained amplitude and time averaged magnitude of individual ionic currents for each set of parameters as compared to the true value.

2. Find whether estimations producing accurate predictions have significantly different values of the cost function.

3. Plot the estimated I-V curves for individual ionic currents and compare them to the true I-V curves.

4. Compare the value of all sets of estimated parameters to the true set of values.

The results of twin experiments with model $\mathrm{HVC_I}$ neurons suggest that when data assimilation produces a set of high quality predictions in the voltage traces with different estimated parameter sets, there is clustering in all three of these measures. When

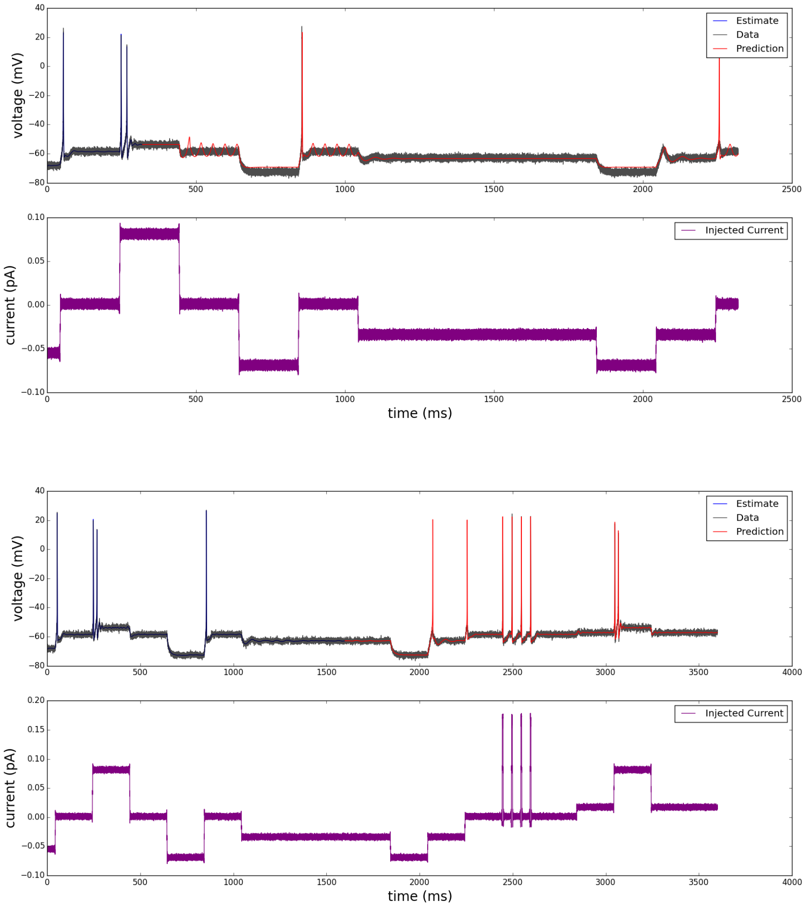

Figure 4.3: Typical data assimilation results using a step current on a synthetic HVC$_I$ neuron. Substantially varying model parameters produce similar quality results. The variation is larger when a step current rather than a complex current is used. Data and a representative estimate and prediction are plotted for the 50 kHz and 10 kHz condition, top and bottom graph respectively. The graphs show membrane voltage (top of each graph) in response to injection of a *step* current waveform (bottom of each graph). The *black traces* show recorded voltage, and the *blue traces* show estimated voltage from the data assimilation procedure for times between 0-320 ms or 0-1,600 ms, during which all state variables and parameters of the model were estimated. The *red traces* show the voltage predicted by integrating the completed model with estimated parameters and state variables forward in time beyond 320 ms or 1,600ms.

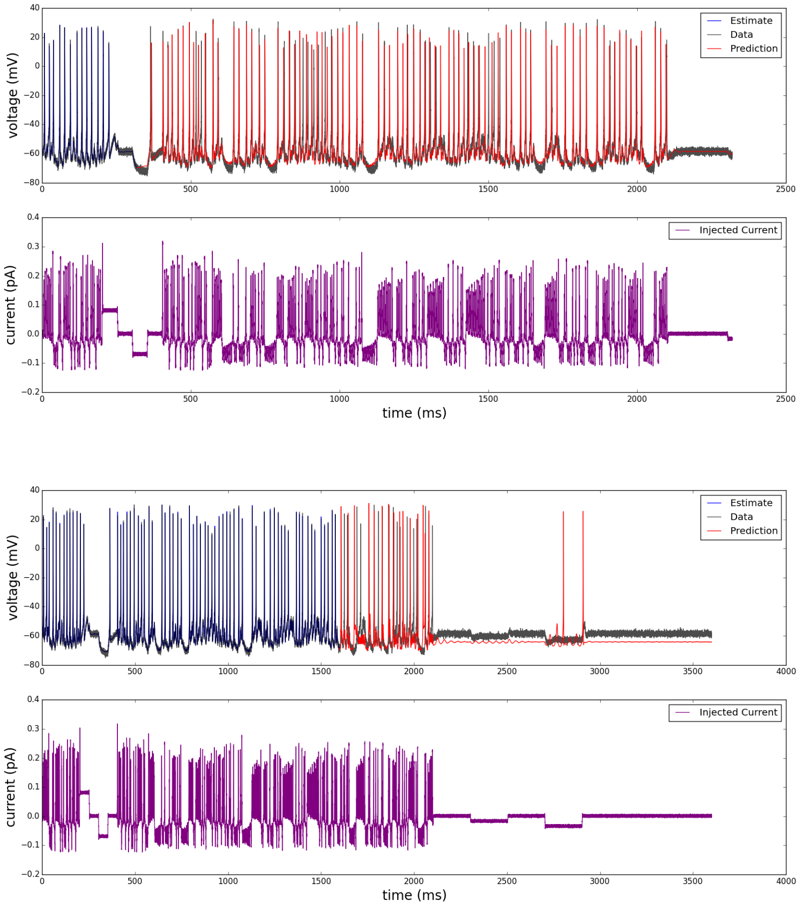

Figure 4.4: Typical data assimilation results using a high frequency complex current on a synthetic HVC$_\mathrm{I}$ neuron. Substantially varying model parameters produce similar quality results. With the high frequency complex current protocol used here, the neuron is not driven far into the subthreshold regime. Passive membrane properties are therefore not estimated well, and the overall prediction quality suffers. A 50 kHz sampling produced results no better than a 10 kHz sampling.

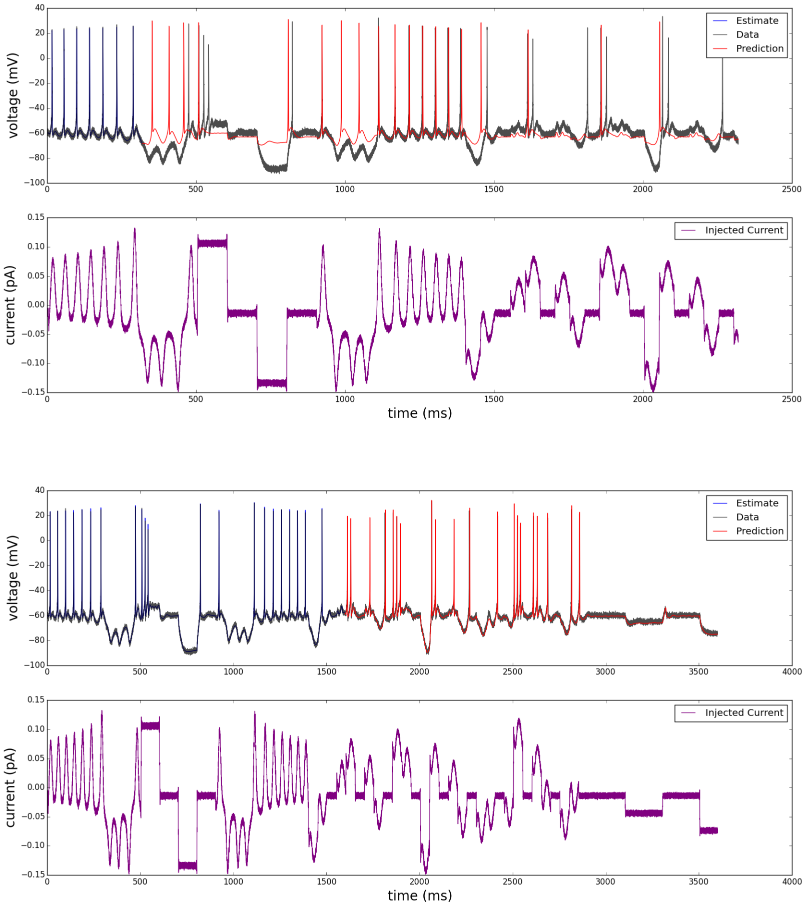

Figure 4.5: Typical data assimilation results using a low frequency complex current on a synthetic $\mathrm{HVC_I}$ neuron. As with other protocols, substantially varying model parameters produce similar quality results. The 50 kHz protocol (top) and the 10 kHz protocol (bottom) use the same number of time points and therefore have the same dimensionality. They both require a similar amount of computation time to finish (around 24 hours on a single compute node), but because the 10 kHz protocol samples additional regimes, including the subthreshold regime, it produces much better predictions.

the accuracy of predictions degrades, the variance of the distributions for each measure increase. Once the predictions have become sufficiently inaccurate, usually when the subthreshold behavior is not matched well and the predicted spike timing is off, the clustering is gone and there are significant deficiencies in the estimations.

Figures 4.6 and 4.9 demonstrate that none of the sets of estimated parameters producing accurate predictions for each stimulating current protocol can be considered to be much more probable than any of the others. This is because the value of the cost function is not very different for each distinct estimated parameter set path, so the integral in the expression for the probability distribution

$$P(x(t_M)|Y(t_M)) = \int dX \exp(-A_0(X,Y))$$

is highly multimodal and does not have any dominant contribution. An expected value for the parameters and estimated initial conditions can still be computed, but because there is no reason to believe that all the minima contributing to the integral have been found or that the expected value of parameters will correspond to the true value, there is not a sound basis for calculating this expected value.

The values of average and max currents and maximally attained values of I-V curves are also systematically observed to be biased upwards of their true values. These data support the hypothesis that large depolarizing currents can be offset by similarly large rectifying currents while maintaining an identical net membrane voltage. Some of the degeneracies in estimated parameter sets were merely due to deficiencies in the training data presented to the algorithm. In such cases, the model produced accurate predictions only when the stimulating current presented after the estimation window was sufficiently similar to the training data, but comparison of the maximally attained, 1-norm, and theoretical I-V curves between different recording epochs showed that some characteristics of the current protocols caused these measures to be closer to their true values despite superficially comparable performance in forward prediction. This demonstrates the importance of validating the model in a number of ways, such as forward integration of the model with a variety of novel stimulating current waveforms in the prediction window.

In the graphs in Figure 4.3 - 4.5, the quality of predictions fell into roughly 3 classes of high, medium, and low quality. High quality predictions matched the voltage waveform of the data exactly. Medium quality predictions showed a close correspondence with subthreshold behavior but missed some spikes. Low quality predictions did not match the subthreshold behavior and missed many spikes.

The 3 stimulation protocols are a step current protocol, a high frequency complex waveform protocol, and a low frequency complex waveform protocol. The *step* current protocol sampled at 10 kHz produced high quality predictions, while the 50 kHz step current protocol produced medium quality predictions. The *high frequency complex* current protocol sampled at 10 kHz produced low quality predictions, while the 50 kHz version produced medium quality predictions. The *low frequency* complex current protocol sampled at 10 kHz produced high quality predictions but produced low quality predictions when sampled at 50 kHz. These patterns can be seen in Figures 4.3 - 4.5.

However, when the plots of the theoretical I-V curves and maximal/1-norm currents were analyzed, the 10kHz-sampled *step* current protocol which produced high quality predictions did less well than the 10kHz-sampled *low frequency complex* current protocol. This shows that complex current protocols better sample the available degrees of freedom in the model than step current protocols. The resulting parameter sets are then likely to generalize better to other stimulating current waveforms.

The *high frequency complex* current waveform may have generated parameter sets less able to generalize well to new driving current waveforms because the subthreshold behavior of the neuron was not adequately explored during the estimation window.

Although the neuron was highly excited during the estimation window, it spent most of its time in similar regions of phase space, along the limit cycle trajectory of a spike. This supports the idea that driving the voltage of a neuron into as many regions of phase space available to the model as possible is crucial when attempting to estimate the parameters and unknown states. The poor performance is likely also due to the fact that the *RC* time constant of the membrane responded only weakly to the rapidly oscillating stimulating current waveform. This information in the voltage waveform was erased by the addition of $\approx$ 1 mV measurement noise, resulting in bad model estimates.

4.2.6 Twin Experiment: Step Current, Sparse

The step current protocol sampled at 10 kHz produced high quality predictions, also doing well on the clustering measures above. Out of 25 paths examined, 21 (84%) produced high quality predictions. In each path examined, all anneal steps producing good predictions were retained. This is likely to be a confounding source of clustering in the measures used because of the similarity of the cost function landscape between annealing steps in the numerical procedure used to locate the lowest minima. A subset of these high quality predictions produced the best predictions, but in order to explore the distribution of the clustering methods above and to simulate conditions in a real experiment where unresolved processes are not known and low error on the training and validation data does not necessarily imply good generalization ability of the model, all accurate predictions were retained. One strategy to evaluate the ability of the algorithm to resolve unknown processes would be to choose only parameter sets producing accurate predictions that correspond to the lowest value of the action in the action level plot of Figure 4.6. This strategy is not taken, however, because in real experiments, parameter sets corresponding to the lowest action values could be overfit to the training and validation data and generalize poorly to different stimulating currents.

Estimated parameters were symmetrically distributed around their true values in many cases with relatively small variance, or were clustered in a region close to their true values. The step current protocol produced the best quality of parameter estimates, in terms of the squared deviation of the estimates from their true values. Yet, the theoretical I-V curves were of lower quality than their low frequency chaotic current protocol counterparts. Among the very best of these predictions, the variance of the clustering measures would likely be smaller. The quality of predictions increased as the number of anneal steps increased, with no degradation in their quality in the range of beta completed in 24 hours.

The primary difference between this set of predictions and the other set of high quality predictions in the next section is that the action level plot plateaued (Figure 4.6) and did not begin to increase again (Figure 4.9).

In the set of estimated parameters, we found that the ionic current amplitudes were estimated to be too large. This is because excessively large estimated depolarizing currents can sometimes be offset by excessively large rectifying currents. Other possible symmetries involve an increase in the maximal conductances and/or larger widths in the steady state activation functions for each ionic current offset by an increased value of the thresholds. In fact, this was often observed to be the case.

Next, the steady state value of the ionic currents as a function of voltage was calculated by substituting the steady state activation and inactivation functions in the expressions for the ionic currents. The pattern here was that the I-V curves for the excitatory currents I_{Na} and I_{CaT} were shifted to the left, as shown in Figure 4.7. The estimated I-V curves tended to match the shape of the true I-V curves less well using the step current protocol than the curves obtained from the complex current protocol, shown in the next section. This is likely due to the fact that the range of parameter sets which assimilate the data from the simpler step current drive is larger than the range of

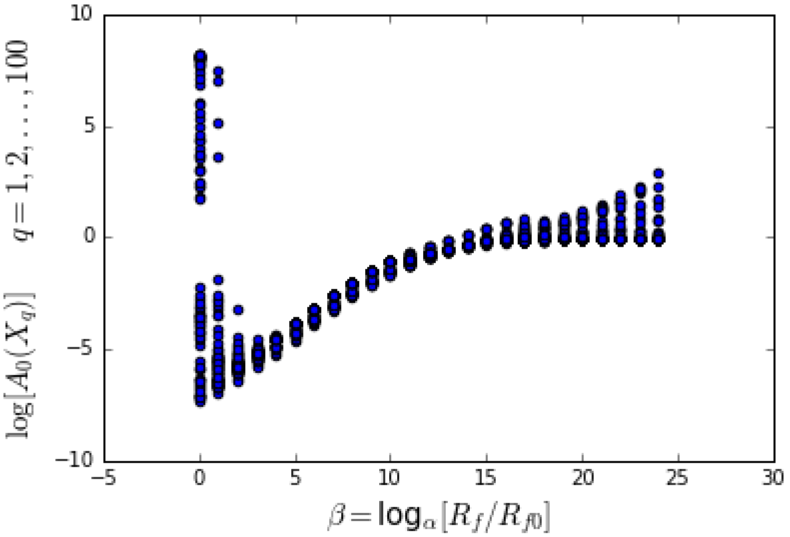

Figure 4.6: Action level (cost function) plot for the sparsely sampled step current protocol. The natural logarithm of the action is plotted against β for $\alpha = 2$. None of the estimated paths produce dominant contributions to the variational integral because the corresponding values of the cost function are not appreciatively different. The action flattens out with increasing β, and in contrast to the situation in Figure 4.9 in the low frequency complex current protocol, does not begin to rise again. Although these models described the data in the estimation and prediction window well, comparison of the estimated theoretical I-V curves with their true value shows that these estimated models are inferior to the models corresponding to the lowest action levels of Figure 4.9. This demonstrates the importance of validating an estimated model in many ways, such as with a variety of novel stimulating currents in the prediction window.

parameter sets assimilating the low frequency complex current protocol, which puts more constraints on the values of the parameters in the assimilation window.

The steady state activation curves for I_K and I_H are shown in Figure 4.8. The I_H waveform was not estimated as well as other current wave forms, probably due to the fact that it is a current that operates on a much slower timescale and could not be sufficiently sampled in the 1600 ms recording epoch. It also operates at hyperpolarized voltages, while the stimulating protocol used did not explore this regime of the membrane voltage sufficiently.

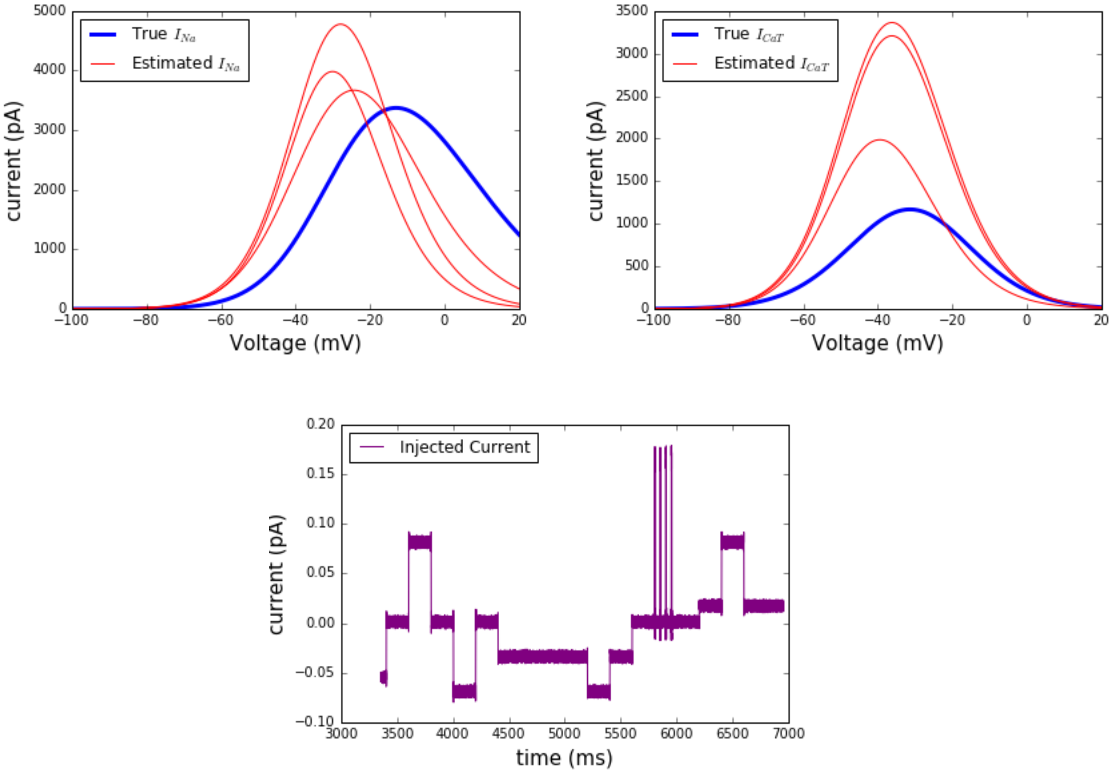

Figure 4.7: Steady state activation functions for I_{Na} (top) and I_{CaT} (bottom) for the sparse step current stimulating protocol. The red curves are sample curves calculated from the estimated parameter sets, and the blue from the parameter sets used to generate the data. The estimates for I_{CaT} appear to cluster around the true value, while for I_{Na} are shifted a bit to the left due to σ_m being substantially overestimated (not shown) with little compensation in other kinetic parameters. Overall, the shape of the estimated theoretical I-V curves using the step current protocol are less like the true theoretical I-V curves than those estimated using the low frequency complex current protocol of Figure 4.10

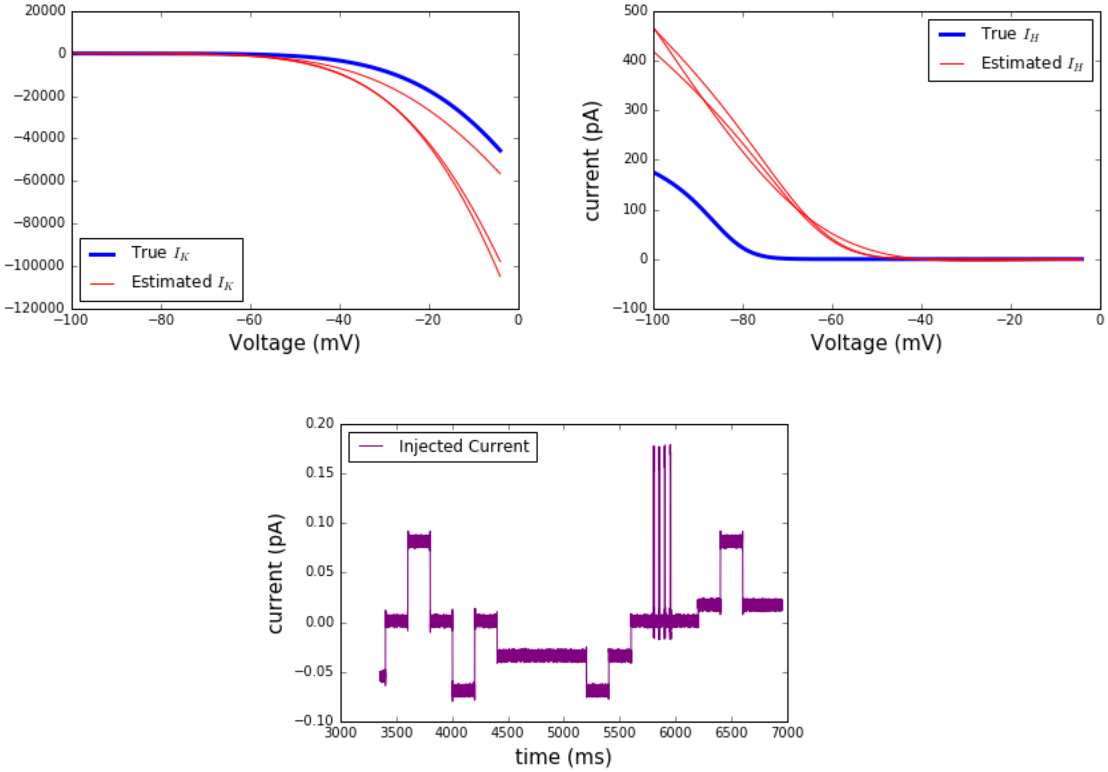

Figure 4.8: Steady state activation functions for I_K (top) and I_H (middle) for the sparse step current stimulating protocol (bottom). The red curves are sample curves calculated from the estimated parameter sets, and the blue from the parameter sets used to generate the data. I_K seems to be well estimated, though here again σ_n was substantially overestimated. To compensate, g_K was often underestimated.

4.2.7 Twin Experiment: Low Frequency Complex Current, Sparse

The low frequency complex current protocol sampled at 10 kHz produced high quality predictions and did well on the clustering measures. Out of 100 paths examined, 11 (11%) produced high quality predictions. A subset of these high quality predictions produced the best predictions in the prediction window, but to best explore the distribution of the clustering of the model, all the levels of quality of predictions were retained.

The graph showing clustering of parameters is once again very large, as there were $\sim$40 parameters estimated. Therefore, this graph is omitted. Estimated parameters were clustered in a region close to their true values, but were systematically overestimated or underestimated. In no case were all of the parameters correctly estimated. Although the quality of predictions monotonically increased with the anneal step (β value) in the range of β completed in 24 hours, the estimated parameter sets were always wrong and therefore the action began to rise again after an initial plateau, as shown in Figure 4.9. Despite this inaccuracy, the obtained predictions are still very good, suggesting that the parameter sets obtained are symmetries in the model given the data. Another possible explanation for the inaccurate parameter estimations is that the Gaussian noise added to the data in the estimation window degrades the quality of the estimations.

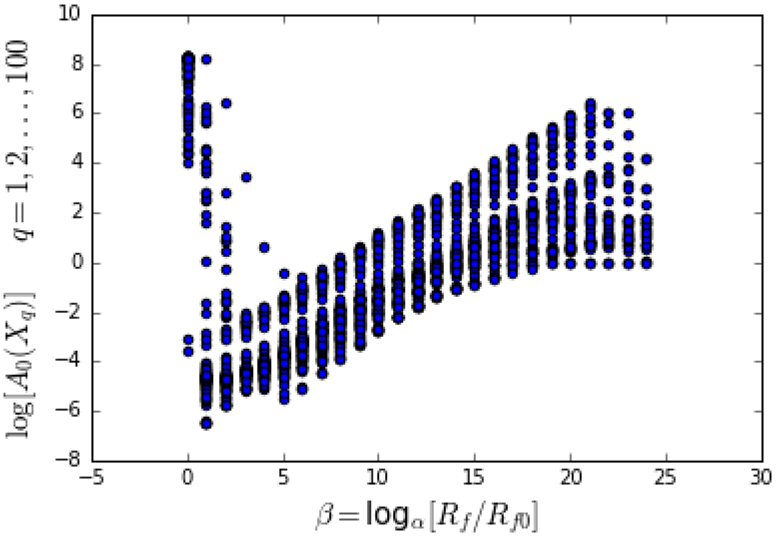

Figure 4.9: Action level (cost function) plot for the sparsely sampled complex current protocol. The natural logarithm of the action is plotted against β for α = 2. None of the estimated paths produce dominant contributions to the variational integral because the corresponding values of the cost function are not appreciatively different. The action flattens out but begins to rise again with increasing β, in contrast to the situation in Figure 4.6. This rising of the action indicates that the model corresponding to the minimum found in the high dimensional landscape is slightly wrong. The paths corresponding to these lowest action levels were still superior to those of Figure 4.6 as measured by forward prediction and by the shape of the estimated theoretical I-V curves compared to their true shapes, a result of the superior stimulation protocol used.

Next, the steady state value of the ionic currents as a function of voltage was calculated. The pattern here was that the I-V curves for the excitatory currents I_{Na} and I_{CaT} had an estimated shape which was mostly correct, but with a too-large amplitude. The curves for the complex stimulating protocol are shown in Figure 4.10. I_K also compares well with its true value, but I_H matches less well, especially at subthreshold voltages, which are not well sampled during the estimation procedure.

The I-V curves for I_K and I_H are shown in Figure 4.11. Again, the I_H waveform was not estimated as well as other current wave forms, probably due to the fact that it is a current that operates on a much slower timescale and could not be sufficiently sampled in the 1600 ms recording epoch. It also operates at hyperpolarized voltages,

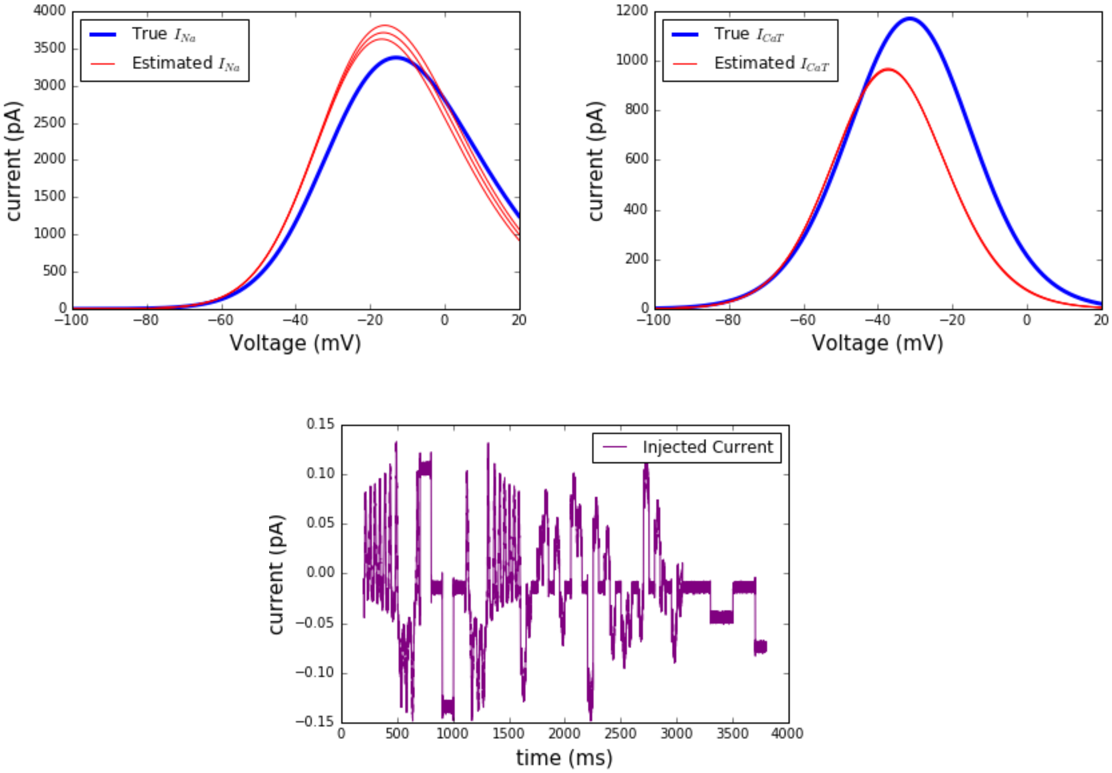

Figure 4.10: Steady state activation functions for I_{Na} (top) and I_{CaT} (middle) for the sparse complex current stimulating protocol (bottom). The red curves are sample curves calculated from the estimated parameter sets, and the blue from the parameter sets used to generate the data. The magnitude of the current was generally overestimated. One respect in which these steady state curves seem to be improved upon the estimations using the step current protocol is that their average shape is much closer to the true shape, if sometimes estimated to be of larger magnitude, and they are much closer to the true value in regions sensitive to turning subthreshold stimulating into an action potential.

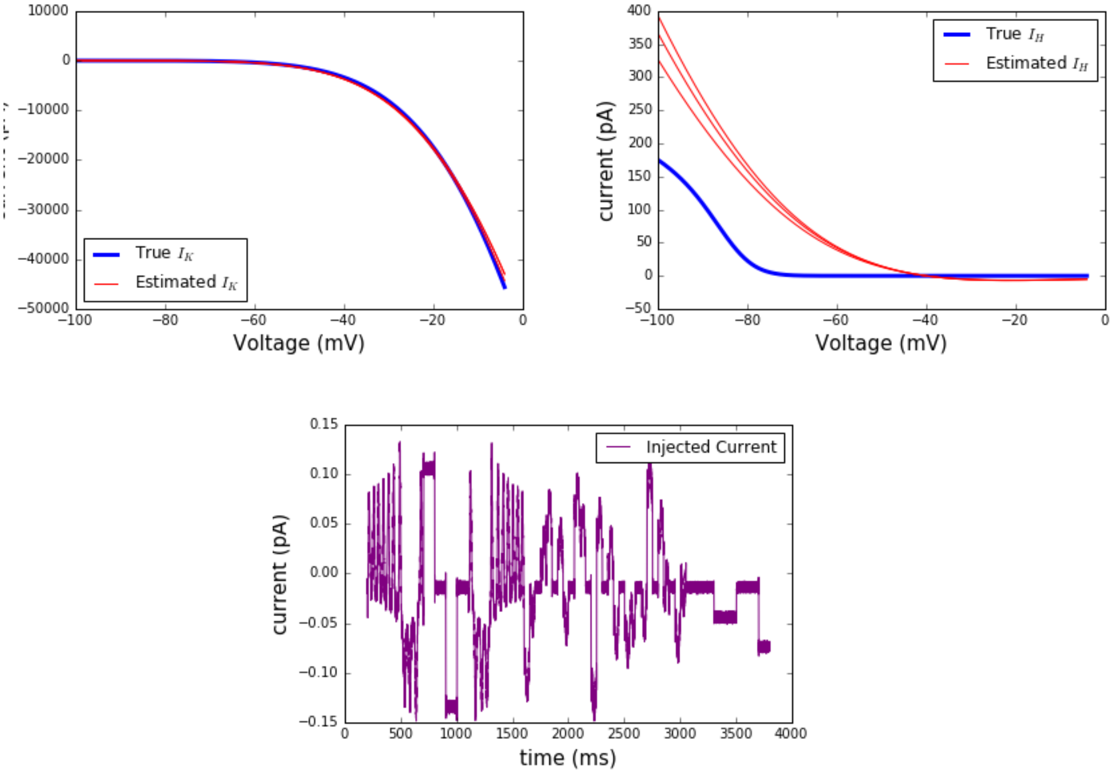

Figure 4.11: Steady state activation functions for I_K (top) and I_H (middle) for the sparse complex current stimulating protocol (bottom). The red curves are calculated from the estimated parameter sets, and the blue from the parameter sets used to generate the data. I_K seems to be well estimated, though here the parameters σ_n, g_K and θ_n were consistently overestimated. I_H, however, did not seem to be as well estimated, probably due to the fact that its overall contribution to net behavior is small.

while the stimulating protocol used did not explore this regime of the membrane voltage sufficiently.

4.3 Discussion

4.3.1 Conclusions to be drawn from model fit to real data

Attaining a model fit producing accurate predictions to novel stimuli on voltage recordings from real HVC$_I$ neurons, including estimating *all* parameters entering the model nonlinearly such as the kinetic parameters, is a significant achievement. To our knowledge, our methods of data assimilation are the only methods to date that are capable of this.

The model fit also reproduced qualitatively observed behavior of HVC_I neurons, including a sag in the voltage in response to hyperpolarizing current injection and rebound spiking upon release. Additionally, the model was able to reproduce experimentally observed effects of blocking I_H or I_{CaT}. This shows that the corresponding degrees of freedom in the model were stimulated during data assimilation with the recorded voltage waveform.

One potential limitation of the current methodology is that all voltage recordings are obtained by injecting a custom current waveform into the neuron which then drives the voltage. Even with a well selected complex current waveform, the regions of phase space of the model neuron are limited to those that can be stimulated at subthreshold membrane voltages and those that are attainable by relatively stereotyped limit cycle trajectories traced out during voltage spikes. A potential methodological improvement could be to control the membrane voltage directly and record the requisite injected current for bringing the membrane voltage to the control values. With cleverly designed control voltage waveforms, the region of phase space traversed and available to the assimilation algorithm could be significantly increased, constraining the number of theoretical I-V and relaxation curves compatible with the data.

4.3.2 Conclusions to be drawn from data analysis of real data

Out of the six different analyses performed using three stimulating current protocols eliciting a voltage measured from HVC_I *in vitro*, the low frequency current protocol at 10kHz (1000.1 ms) was the most successful [7]. It appears that the improvement in prediction and estimation from longer estimation windows due to lower frequency sampling surpasses any decline in prediction and estimation quality due to lower sampling rates, which means lower dimensionality of the optimization problem to be solved. It can also be concluded that a step current is less effective at sampling the I-V curves and relaxation properties of a neuron to clearly define minima in the action and prevent the formation of isosurfaces in the cost function which lead to inaccurate models. A chaotic current, with a region of extended negative current to explore subthreshold membrane voltages activating I_H and I_{CaT}, is required for higher quality predictions.

In general, a chaotic current with slow oscillations relative to the RC time constant of the membrane and an extended hyperpolarizing region within the estimation region are needed for the best predictions. Additionally a longer estimation window is more important than the additional resolution in time of a 50 kHz sampling rate as opposed to a 10 kHz sampling rate.

In the analysis on real data, the estimated magnitude of I_H is less than or equal in magnitude to the leak current. If a simpler model is desired, this suggests that I_H may be removed from the model. Alternatively our model is missing some other crucial component, or the parameterization of I_H in the model is not realistic.

The magnitude of I_H could be as small as it appears, with its main effect in the subthreshold region, being small but contributing significantly to the behavior of the neuron. I_H and I_{CaT} are complementary and crucial mechanisms in the dynamics of some neural circuits [33, 54]. Future work should examine the role of I_H in the dynamics of HVC.

4.3.3 Conclusions to be drawn from data analysis of synthetic data

In our HVC$_I$ neuron model, different combinations of parameters could produce accurate predictions and effectively identical voltage traces. Plotting the theoretical I-V curves from the estimated sets of parameters against their true value shows that max conductances g, activation and inactivation thresholds θ, and widths σ can trade off with each other to produce I-V curves which are similar in shape, and which produce nearly identical responses in the membrane voltage to different values of the injected current.

There also exist model degeneracies in parameters given the assimilated and validation data. We found that data assimilation on step currents could produce a variety of models producing accurate predictions when current waveforms similar to those found in the estimation window were presented to the neuron. However, the resulting estimated theoretical I-V curves were less similar to the true theoretical I-V curves than in the case where a low frequency complex current waveform was presented to the neuron. If the neuron model estimated from the voltage waveform elicited by the step current stimulating protocol is presented with one of the chaotic current stimulating protocols, it is likely to produce inaccurate predictions.

When the voltage waveform produced by driving the neuron with the high frequency stimulating current was presented to the assimilation algorithm, it did not produce accurate estimated neuron models. This could have been caused by the fact that the high frequency stimulating current did not sample the subthreshold degrees of freedom of the neuron model, producing parameter sets able to predict accurately only in stimulating current regimes producing spiking, but doing relatively poorly when tested in subthreshold regimes. Overall, this stimulating current produced estimated theoretical I-V curves far from their true value.

Another factor explored in our twin experiments on HVC_I neurons was the influence of the tradeoffs in higher or lower sampling rates given limited computational resources. We found that for assimilation on the HVC_I neuron model, 10 kHz time resolution did better than 50 kHz. In this case, this means that the loss in information transfer from data to the model by reducing the time resolution was outweighed by the increase in information transfer by stimulating additional degrees of freedom in the model.

Chapter 4 is reproduced from material as it appears in a preprint on ArXiV. “HVC Interneuron Properties from Statistical Data Assimilation. Daniel Breen, Sasha Shirman, Eve Armstrong, Nirag Kadakia, and Henry DI Abarbanel. The dissertation author was the first author on this preprint.

Chapter 5

Use of Data Assimilation for Inference of CA1 Neuron Pathology in 3xTg Mouse Model of Alzheimers

Experiments used to characterize altered electrical and molecular properties of neurons in Alzheimer's disease are laborious and may fail to detect variability and multifaceted differences in and across populations of neurons. We present techniques of data assimilation together with a data mining approach to investigate pathology in the 3xTg mouse model of Alzheimer's disease combining both amyloidopathy and tauopathy. We find differences in individual intrinsic excitability and use our biophysically tuned models to propose mechanisms underlying these differences. The scientific contribution is primarily methodological, as we presently wish only to show a proof of principle of the methods. We set a foundation to extend the methods towards investigating alterations in Ca^{2+} dynamics and medical applications.

While the etiology of Alzheimer's disease (AD) is unknown, the production of beta-amyloid peptides ($A\beta$) [30] and disruption of Ca^{2+} signaling [53, 79] in the brain are involved. Histopathological features of AD, including amyloid and tau pathology, interact with the calcium signaling dysregulations in a vicious spiral [82, 80].

What is now lacking in the understanding of neuropathology are the links between underlying causes at the molecular, single cell, and circuits level of the disease and the behavioral symptoms. Computational modeling offers the potential to link across the different levels of genes and drugs, synapses and neurons, and cognition and behavior to provide a unified, testable basis for formulating therapeutic strategies [23, 18, 6, 58, 62, 99].

Mathematical models have been applied to the study of Aβ related excitability changes [18, 62, 99]. Biophysical neuron models such as conductance based Hodgkin Huxley (HH) type models contain interpretable and biologically meaningful parameters capable of reproducing experimentally observed membrane dynamics. Such models thus provide a framework in which to identify changes in biophysical properties of diseased cells in a manner that elicits targets for potential therapeutic intervention.

Neurophysiological studies often focus on how AD pathology disturbs synaptic function [12, 16, 43, 74, 75, 76, 98]. More recently, alterations in intrinsic neuronal properties have been examined [6, 26, 57, 10, 18, 85, 86, 36, 42, 92, 72]

We wish to identify quantitative changes in the detailed electrical membrane dynamics between wild type and diseased cells where the latter have been genetically modified, in mice, to produce analogs to human Alzheimer's responses. These 3xTg-AD mice express mutant presenilin (PS1), Aβ precursor protein (APP), and tau, and develop both plaques and tangles in an age and region dependent manner [63, 64]. The goal is to further develop HH like models of cortical cells implicated in Alzheimer's disease and to inform the use of well designed experiments which will quantify changes, including intracellular Ca^{2+} dynamics of these cells. Models will incorporate enough biological realism to be useful tools in the neuropharmacological process of drug design, such as by prescreening a number of compounds for bench testing in the laboratory. An example would include using the estimated models as a readout for alterations in Ca^{2+} dynamics and molecular signaling pathways, an area that has been identified as a target mechanism where the initial stages of Alzheimer's can be tracked and where further degeneration of 3xTg cellular behavior can be potentially identified and treated [15]. The present work investigates alterations in neuronal intrinsic excitability, and sets foundations which can be used to attend to alterations of Ca^{2+} dynamics in future work.

We have previously developed methods of statistical data assimilation which have been successful in incorporating information in complex chaotic and neural systems, from incomplete data sets [88, 45, 1, 44, 56, 94, 93, 39, 61].

Other methods of parameter fitting in neuron systems include hand-tuning [78], parameter space exploration [5, 68, 69], gradient descent [5], and evolutionary algorithms [87, 41, 3]. These methods have yielded the estimation of many states and parameters as well, but with one or more of the following limitations [17]:

- parameter identifiability issues are not addressed
- the parameters must be heavily constrained, reflecting unlikely prior knowledge
- only parameters which enter the model linearly can be accurately estimated
- make adequate “best fits” but fail to predict

The first and last are crucial shortcomings - work in single neuron modeling often does not raise the issue about the uniqueness of estimated parameters [17], nor does it often consider the subjectivity of a good fit. Hand tuning is highly subjective, and even if a ‘good’ parameter set can be found, it is never certain if there is a better one that has not been discovered. Methods incorporating a heuristic metric of error, including parameter space exploration and evolutionary algorithms, may miss best parameter sets and have poor predictive power.

The approach advanced here directly addresses the issue of testable model fitting in neural systems. Previously [88, 45, 44, 56, 39, 61], it has demonstrated success in estimating all of the parameters in HH conductance models, including those that enter nonlinearly such as the gating kinetics describing the ion currents. They have also been shown capable of predicting precise waveform information of the voltage of a neuron evolving in time according to a novel stimulating current.

5.1 Methods

5.1.1 CA1 Neuron Model

One of the limitations of minimizing a smooth loss function is the danger of getting stuck in local minima. This tends to yield a number of different parameter sets which fit the data well and have good predictive power. A more complex model including more ionic currents would have the advantage of incorporating some elements of biological realism. However, we are interested in finding alterations in parameters and features which can be used to distinguish populations of neurons in a biophysical context, so we must take into consideration this parameter identification problem.

We subjected different neuron models, some including Ca^{2+} dynamics, to our methods of parameter estimation. Although we wish to identify alterations in Ca^{2+} dynamics in 3xTg cells, these models encountered substantial parameter identifiability obstacles that remain to be overcome. Instead, we settled with a basic HH model, as it is a biophysically realistic model which provides a simplified arena in which to cope with obstacles of parameter identifiability. The basic HH model did better at reproducing many features of the data without the parameter identification issues of more complicated models. This will set a foundation for future refinements including calcium and calcium activated potassium currents in the model, which are able to address investigating alterations in calcium dynamics.

In our experiments collecting voltage data, we use a current clamp setup, where voltage traces are measured with a known injected current. Intuitively, one example of parameter unidentifiability is that different depolarizing and repolarizing ion currents can contribute to the opening and closing of ion channels underlying voltage response to injected current. Using only measured voltage, it is difficult to determine what the relative contributions of different ionic currents are to characteristics of features of the voltage waveform.

Our CA1 neuron model is thus a basic Hodgkin-Huxley (HH) neuron with only sodium (I_{Na}), potassium (I_K), and leak (I_L) currents:

$$\begin{aligned}
C_m \frac{dV(t)}{dt} &= g_{Na} m(t)^3 h(t)(E_{Na} - V(t)) + g_K n(t)^4 (E_K - V(t)) \qquad (5.1)\\
&\quad + g_L(E_L - V(t)) + I_{injected}(t)\\
\frac{dx(t)}{dt} &= \frac{x_\infty(V) - x(t)}{\tau_x(V)}\\
x_\infty(V) &= 0.5(1 + \tanh(\frac{V - \theta_x}{\sigma_x}))\\
\tau_x(V) &= t_{x0} + t_{x1}(1 - \tanh^2(\frac{V - \theta_x}{\sigma_{xt}})) \qquad (5.2)
\end{aligned}$$

Here $V(t)$ is the membrane voltage, C_m is the membrane capacitance, and $I_{injected}(t)$ is a known stimulating current injected into the neuron in a current clamp setup. g_i and E_i denote the maximum conductance and reversal potential for current i, respectively.

The equations for x are a shorthand for the kinetics of the gating variables m, h, and n. $x_\infty(V)$ is the voltage dependent steady state activation which depends on θ_x, the voltage at half activation, and σ_x, the width of this activation. $\tau_x(V)$ is the voltage dependent relaxation time, describing the rate that the gating variables change to their steady state values.

Errors in the descriptions of the dynamics of these channels in HH models and noisy measurements contribute to the parameter identifiability problem. We show, using our data assimilation algorithm described in section 5.1.2, that we can find sets of parameters in a basic HH model capable of reproducing nearly exactly the shape of the waveform of our recorded data. Correctly predicted features include the shape of the action potential waveform, afterhyperpolarization effects, refractory periods, spike timing, and subthreshold variations. An example of a prediction is displayed in Figure 5.2.

Our strategy here is to estimate ensembles of models and look for similarities and differences in patterns in these ensembles across distinct populations of neurons. Since there are many sets of model parameters which are compatible with the observed data, we do not attempt to pinpoint a unique ‘correct set’ of model parameters. Each set of model parameters possessing good predictive power is assumed to be as good as any other set. We do retain some element of subjectivity in that estimated models are said to be compatible with the data when predictions, generated by integrating the model beyond the estimation, are judged to match the subthreshold voltage variations and spike timing. This subjective evaluation was generally easy as matches to the data were either very good or very poor.

The criteria used to make this evaluation included checking the value of a number of features of the estimated models. For example, predictions were checked for action potentials with a duration within a biologically realistic range (no more than several ms), whether subthreshold voltages and spike amplitudes were within a few millivolts of the correct value. Another useful measure involved checking whether the value of the objective function, defined in section 5.1.2, was smaller than a cutoff value. The cutoff value was easily determined as the objective function tended to take either large or small values which corresponded to very poor or very good fits, respectively.

5.1.2 Methods of Data Assimilation

Data assimilation refers to analytical and numerical procedures in which measurements and models are combined to infer knowledge about a system which is not available in the measurements alone. Information in measurements is transferred to model dynamical equations selected to describe the processes thought to produce the data. The problem is typically formulated as follows. We formulate a model describing the system with state variables $\mathbf{x}$ defined at times $t_0, t_1, ..., t_T$, and we seek to infer the model state variables at the end of the estimation window $\mathbf{x}(t_T)$ and the unknown model parameters $\mathbf{p}$.

One of the difficulties commonly encountered in data assimilation arises from the fact that systems in the real world almost always contain processes that the modeler is ignorant about or cannot represent. Another difficulty comes from the fact that measurements are noisy, which limits the ability to infer properties of systems even in the presence of perfect models. Estimating properties of the dynamics of systems which are nonlinear, such as neurons, only compounds these problems. Previously, we have developed methods of data assimilation capable of dealing with gaussian measurement noise when the neural system under study is perfectly described by a basic HH model [88, 45, 44, 39]. As our focus here is on data from real neurons, we cannot hope to have a perfect model, but past studies provide a firm basis on which to build our understanding here.

We have applied a variational approach which has been used successfully for estimating models of neurons, using data from neurons in region hvc of the avian song system, in [56, 61] and remains a reliable choice in complex or simpler models, including our basic HH model of CA1 neurons here. A critical concern in variational approaches to nonlinear dynamical systems is that the familiar least-squares objective function may give an irregular search surface with many local minima. We addressed this problem here by including a balanced synchronization term, $u(t)(V_{data}(t) - V(t))$, in the model dynamics, which regularizes and minimizes the influence of local minima by ensuring that the solution set defines a model that is capable of synchronizing with the data. The objective function includes a penalty for this regularization and is taken as

$$\sum_{t=1}^{T} (V_{data}(t) - V(t))^2 + u(t)^2 \tag{5.3}$$

where $V_{data}(t)$ is the measured voltage, $V(t)$ is the voltage output from the model, and T is the number of discretely sampled time points in the estimation window. $u(t)$ describes the magnitude of the synchronization term.

The neuron model that we investigated is not known to be chaotic for the biophysical range of our parameters, but it is nonlinear, and in the high-dimensional search space, there may well exist chaotic regions that must be explored by the optimization routine, and these will benefit from this regularization. If a solution to the model can be found that is consistent with the data, the value for the control parameter $u(t)$ should become small relative to the model dynamics at the end of the optimization. The quality of the model was tested by setting $u(t)$ to zero and integrating the model forward from the end of the assimilation period using the estimated parameter and state values.

The implementation of the data assimilation algorithm was accomplished through the use of the open source software package IPOPT (Interior Point OPTimizer) with the linear solver ma57 [90].

5.1.3 Detecting Altered Features and Mechanisms

We are interested in identifying features of membrane electrical behavior in our neuron models which are different between two strains, 3xTg and nonTg. Using the estimated models, the number of features was expanded from the set of 20 estimated parameters in the models to include approximately 50 total features.

The predictions, obtained by integrating the model forward past the end of the estimation window, is the blue trace such as those shown in Figure 5.2.

We trained random and regression forest models using the features and model parameters to predict labels of interest, including the known type of each neuron (3xTg or nonTg), threshold voltage, spike half-width, and other features. We used the feature importance attribute of the trained models to determine the relative importance of each feature in predicting the value of each label.

We also calculated kernel density functions over the data points in three dimensional subspaces of the 50 dimensional feature space. We used these density functions to estimate the overlap of three dimensional probability distributions describing the 3xTg and nonTg neurons. Feature importances significantly reduce the number of feature subspaces to examine when searching for differences between the two populations.

We have used regression forests to detect predictors of altered features between 3xTg and nonTg strains. These can include model parameters and other features. An example includes discovering other features and model properties predicting intrinsic excitability. Threshold voltage and the rheobase are examples of features quantifying intrinsic excitability. By finding predictors of quantities such as these, we are able to propose mechanisms causing the observed differences in the two populations of neurons. We also computed correlations among these model parameters and features, though in some cases relationships are nonlinear.

It may be that features detected to be altered differ only between subsets of 3xTg and nonTg. In this case, if two features A and B are found to be altered, it may be that A is altered in some neurons while B is unaltered in a other neurons. Then, in some neurons, feature B may be altered, while A is unaltered in others. A decision tree may reveal this kind of structure, while a table of variable importances will not. For clarity and simplicity of the presentation, we do not include decision trees, but mention in passing that splits along different features descending down the decision tree supported this notion.

In our set of recordings, an epoch is defined as each instance in which recordings are obtained from a neuron by eliciting a response with a stimulating current. Each epoch usually involves the neuron being stimulated with a new injected current.

Using our methods of data assimilation, we estimated model parameters for 56 epochs from 5 3xTg and 6 nonTg neurons. Recordings were collected from two preparations. Each preparation contained one 3xTg and one NonTg mouse, both male. One preparation used 3 month old mice, while the other preparation used 1 month old mice. For each data set, we initialized our data assimilation search procedure with 500 distinct sets of parameters sampling each parameter from a uniform distribution defined between the parameter bounds. At the end of search procedure, many of these parameter value initializations did not result in an estimated model which passed the validation tests described in section 5.1.1. These were not included in the processed data set used to draw inferences about differences between 3xTg and nonTg neuron populations.

In order to balance out the contribution from each neuron and class in the data, we sampled with replacement from the pool of estimated parameter sets so that each neuron contributed an equal number of samples to its respective class. Finally, we ensured that the 3xTg and nonTg pools had equal size.

5.2 Results

5.2.1 Estimating Ensembles of Model Parameters

Much of the power of the analysis presented here derives from the data assimilation procedure described in section 5.1.2 and elsewhere. When integrated forward, the estimated model parameter values accurately predict spike timing, subthreshold variation, afterhyperpolarization, spike amplitude and other features within the range of biological trial by trial variability. This even holds true with a basic HH model with only the usual I_{Na} and I_K currents, despite missing many other currents which are present. Figures illustrating the data assimilation protocol on the data set are displayed in Figures 5.1 and 5.2. The model fit and prediction of Figure 5.2 is among the best produced, though not by much. The technical difficulty of estimating the many parameters of even a basic HH model is challenging because many of the parameters enter the equations of motion nonlinearily. Although the model fits and predictions are often judged subjectively to be quite good, they are not perfect and make mistakes in the prediction. This is illustrated in Figure 5.3, where some of the action potentials and features produced by other currents, including I_{SK}, are not accurately reproduced.

We do not find that all 3xTg models can be cleanly separated from nonTg. The differences between estimated parameter sets in the two classes may sometimes be subtle or nonexistent. This could reflect the fact that 3xTg and NonTg do not always have strikingly different electrophysiological behavior, or it may be due to present limitations of the model and the fitting procedure. Because only 11 neurons were analyzed, some of the differences found may be spurious due to the low sample size. These include some of the passive membrane properties, as discussed in section 5.2.2. Still, in some aspects, clear differences in the neural dynamics do emerge.

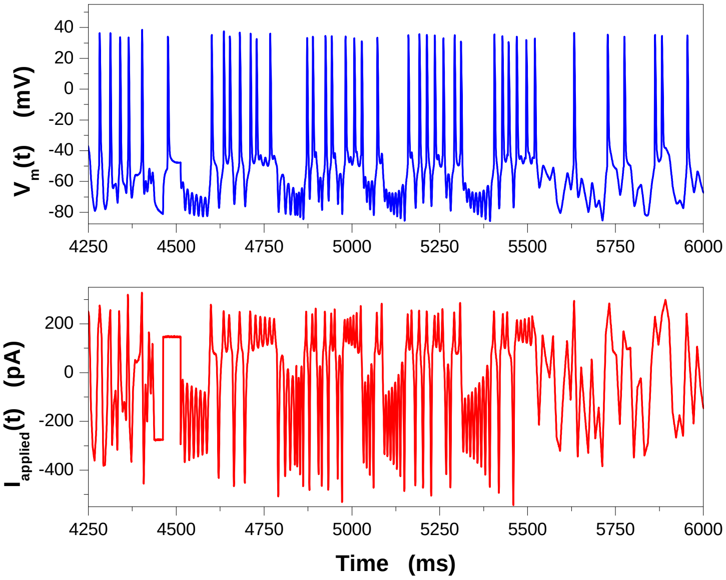

Figure 5.1: Recorded voltage data and injected current. The stimulating current waveform used in the data assimilation procedures is a combination of a pseudo-noisy current and the output of a chaotic model. The pseudo-noisy current was created by uniformly sampling current values about every 20 ms and linearly interpolating in between. The chaotic current waveform is the output from the Lorenz63 model used as a stimulating current.

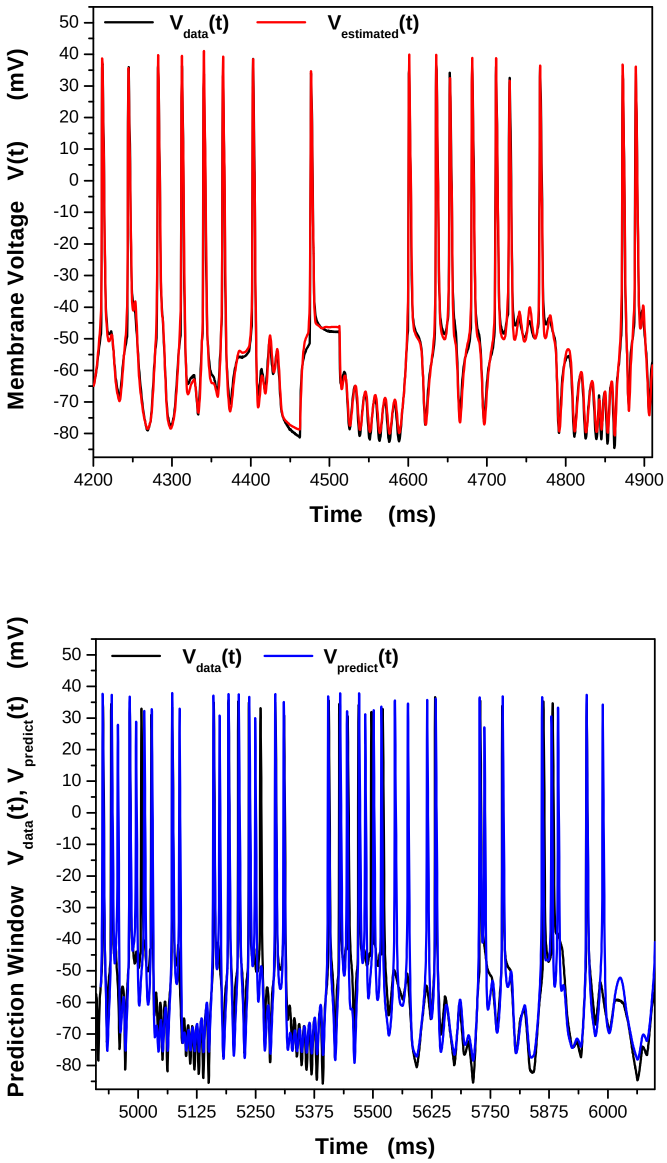

Figure 5.2: Recorded voltage data (black), model estimate (red), and prediction (blue) obtained after the minimization procedure. The quality of the model estimate alone does not give any information about whether the estimated parameters are good. To evaluate the goodness of fit of the parameters to the data, a prediction must be generated by integrating the equations of motion, using the estimated parameters, starting from the end of the estimation window.

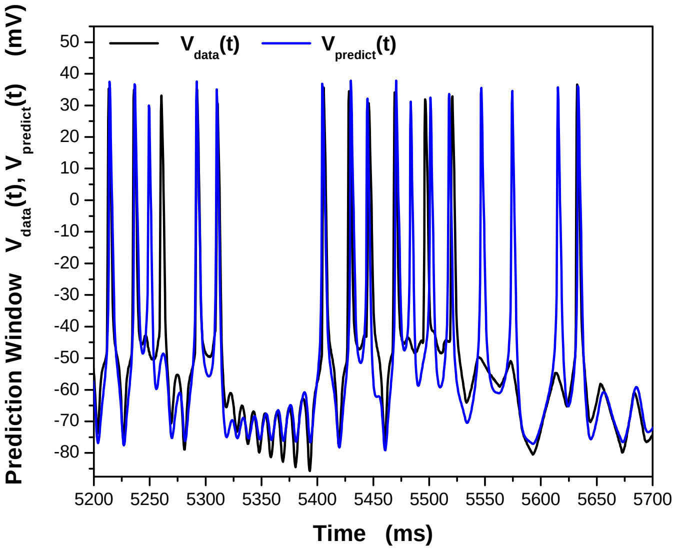

Figure 5.3: Estimated voltage (black) and prediction (blue), obtained by integrating the estimated model forward in time. The prediction is displayed in higher resolution to clearly that while the prediction is very good, it also does not always perfectly reproduce all spikes and features like the AHP likely produced by the calcium dependent potassium current I_{SK}. A future area of improvement would be to reproduce the AHP following durations where the neuron is active. Refinements to the model and methodology overcoming model identifiability issues not adequately addressed by other work will be necessary.

5.2.2 Identifying Differences Between 3xTg and nonTg Neurons

Random Forest Analysis

Figure 5.5 shows the separation achieved in a subspace including the membrane time constant $\tau_{mem} = C_m/g_L$, the leak conductance g_L, and the membrane capacitance C_m. Among the full set of features tested for differentiating classes of neurons, this and other subspaces including these features are excellent at distinguishing between the two populations as measured by variable importance calculated from the trained random forest classifier. However, the boundary that one can draw between the NonTg and 3xTg strains in subspaces including these features is very complicated. A complicated separating subspace between the two strains may simply be the result of the fact that the number of neurons analyzed is not very large. Therefore, we hesitate to draw any conclusions about whether this separation is biologically significant.

The fact that membrane capacitance C_m, leak conductance g_L, and membrane time constant τ_{mem} are excellent predictors of neuron strain because the number of neurons is low caused us to exclude them from the calculation of variable importances predicting neuron type in Table 5.1. Without the influence of these parameters, rheobase, threshold voltage, $I_{DC} = g_L E_L$, and features measuring excitability and properties of I_{Na} and I_K during action potentials are the most important distinguishing characteristics. Rheobase is defined as the lowest value at which a sustained step current first causes an action potential. The voltage at which this occurs is the threshold voltage. Rheobase and threshold voltage are depicted in Figure 5.4, where simulated and measured traces are compared. g_L is the leak conductance, while E_L is the voltage at which the leak current I_L reverses.

The time derivative of the ionic currents were calculated at the threshold voltage during an action potential, denoted $\dot{I}_{Na}(V_{thres})$. There are two times at which the

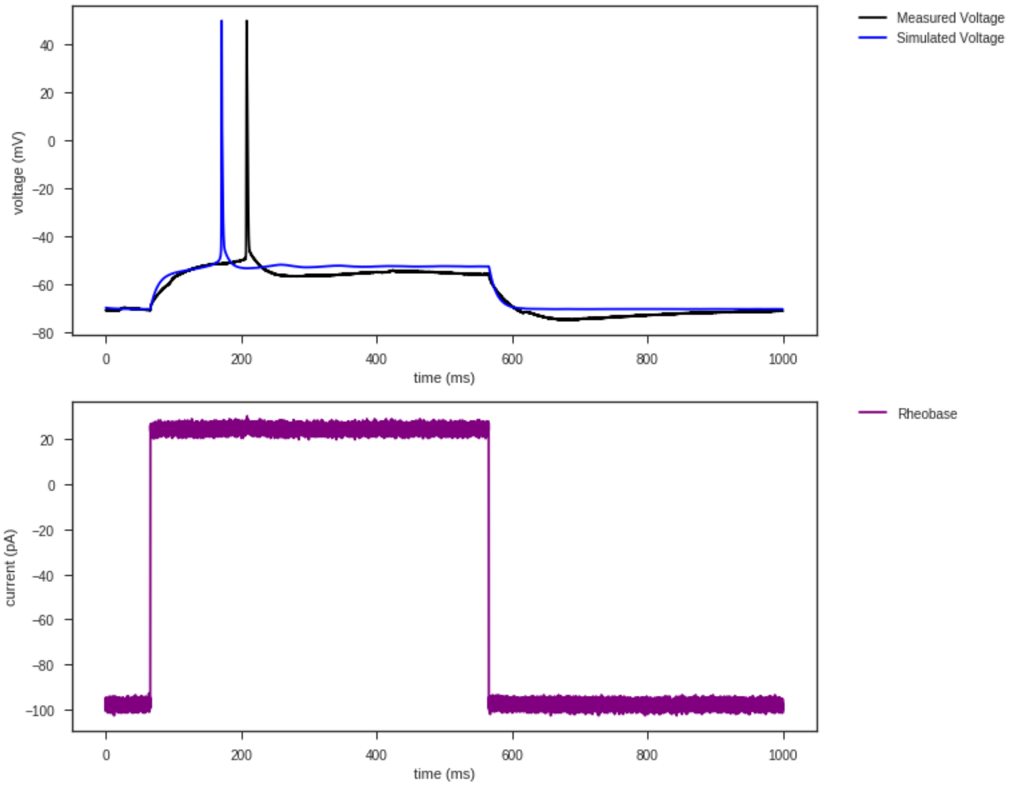

Figure 5.4: Injected current (purple) eliciting measured (black) and simulated (blue) voltage responses. The simulated trace (blue) is obtained by integrating forward the model equations with the estimated parameters. In the experiments, the voltage is held at -70 mV. Rheobase is then defined as the smallest magnitude of the step that produces an action potential, here approximately 120 pA. The threshold voltage is the voltage at which the action potential occurs, here at approximately -50 mV. Although the model reproduces some features of the voltage trace well, some ionic current currents including I_{SK} are likely missing.

Table 5.1: Features predicting neuron strain (3xTg or NonTg) are calculated using the variable importance from a random forest classifier. Rheobase is the lowest sustained step current required to elicit an action potential, while threshold voltage is the voltage at which the action potential is first elicited.

Features Predicting Strain	Var Imp	Centroid Difference
Rheobase (pA/pF)	0.122	-1.648
Threshold Voltage (mV)	0.070	-6.301

voltage passes through the threshold value, one during the upstroke, and one during the downstroke. There was a pattern detected of altered values between the two strains in the time derivaties of I_{Na} and I_K at threshold voltage. However, the magnitudes of I_{Na}, I_K, and their time derivatives were not found to be altered. It is not clear why $\dot{I}_{Na}(V_{thres})$ and $\dot{I}_K(V_{thres})$ are effective predictors of neuron strain, but altered threshold voltage may partly account for this.

I_{DC} is associated with a measure of excitability distinct from, though sometimes related to, reduced threshold voltage or rheobase. In the model, increasing E_L increases I_{DC}, which can counteract repolarizing currents such as I_K during trains of action potentials and reduce depolarizing currents needed to generate an action potential. Larger I_{DC} does not consistently predict decreased threshold voltage and current in the set of estimated models. Instead, I_{DC} may be related to these values in more complicated ways. In some cases, I_{DC} is increased as a compensatory mechanism to retain a similar threshold voltage when properties of I_K and I_{Na} are altered. Still, in Table 5.4, increased I_{DC} does have an overall weak negative correlation with rheobase. A similar explanatory principle may account for why other features effectively predicting the value of a given feature in a regression tree structure show only weak correlations in Tables 5.3 and 5.4.

Alterations in the kinetic properties of I_{Na} and I_K in the two strains of neurons were also detected. Altered parameters are listed in Table 5.2. The regression forest models of Tables 5.3 and 5.4 show that θ_m, θ_h, and σ_h are important predictors of threshold voltage and rheobase. These are features which are different between neuron

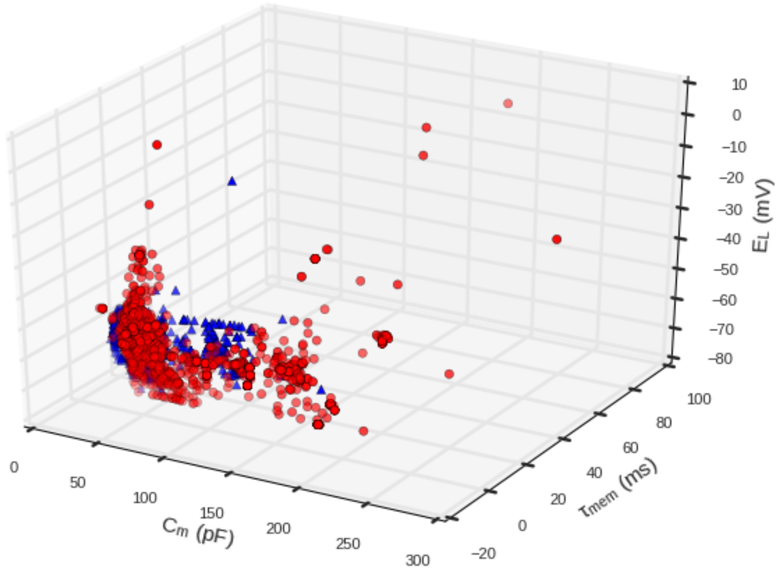

Figure 5.5: The (C_m, $\tau_{mem} = C_m/g_L$, E_L) random variable subspace gives one of the best separations between the two strains consistent with the calculated variable importances. There are more subspaces in which the two strains separate, but we focus on this subspace as these features are biologically interpretable while showing that the methods presented in the paper detect a separation between the two strains. In this subspace, the separation is complex and the data is sparse, so strong conclusions should not be drawn from it. 3xTg estimated values are plotted in red, nonTg cells in blue. Other subspaces, such as those of Figures 5.6 and 5.7, give separations of 3xTg and nonTg values which are simpler with attendant mechanistic interpretations we describe.

Table 5.2: Model parameters predicting neuron strain. $I_{DC} = g_L E_L$ is equivalent to an applied DC current, E_L is the potential at which the direction of the leak current I_L reverses. σ_h is the width of the voltages at which the gating variable h describes the inactivation of I_{Na}, $theta_m$ is the half-activation voltage of m describing the activation of I_{Na}, and σ_{ht} is the width of voltages where the relaxation time are large.

Params Predicting Strain	Var Imp	Centroid Difference
I_{DC} (pA/pF)	0.134	0.954
E_L (mV)	0.117	-1.322
σ_h (mV)	0.096	10.672
θ_m (mV)	0.087	-3.701
σ_{ht} (mV)	0.058	4.535

Table 5.3: Feature importances for threshold voltage using only the model parameters as predictors. $theta_m$ is the half-activation voltage of m describing the activation of I_{Na}. σ_h is the width of voltages where h describes the inactivation of I_{Na}. The appearance of θ_m and σ_h shows that in the dataset, variations in properties of I_{Na}, rather than I_K, account for variations in the threshold voltage. Interestingly, a longer membrane time constant τ_{mem} is perhaps associated with lower threshold voltage.

All Params Predicting Threshold Voltage	Var Imp	Correlation
θ_m (mV)	0.208	0.392
τ_{mem} (ms)	0.127	-0.320
σ_h (mV)	0.081	-0.363

strains in the random forest model of Table 5.1. The two random forest models predicting strain then complement each other, with the regression forest models predicting rheobase and threshold voltage helping to provide a coherent picture of the altered mechanisms underlying these observed differences.

In the (θ_m, σ_m) subspace, shown in Figure 5.6, a separation between 3xTg and nonTg cells is apparent and provides evidence that altered kinetics in I_{Na} results in a lower threshold voltage in 3xTg. Inspecting Figure 5.6 suggests that there is a weak linear relationship between θ_m and σ_m. One can test θ_m and σ_m for collinearity from an analysis of eigenvalues and eigenvectors of the feature correlation matrix. The correlation matrix **C** is has elements C_{ij} which are the correlation between random variables i and j. Here, random variables i and j are features and parameters of the estimated neuron models. When an eigenvalue of this matrix is close to zero, the nonzero indices of the corresponding eigenvector shows which variables are collinear. Therapeutic remedies could be screened in one way by checking for whether they can shift the value of θ_m and/or σ_m in 3xTg cells into the nonTg regime of Figure 5.6.

Strain Probability Distribution Overlap

We calculated Gaussian kernel density estimates (KDEs) of the probability distributions over three and four dimensional feature subspaces, grouped by strain. Subspaces

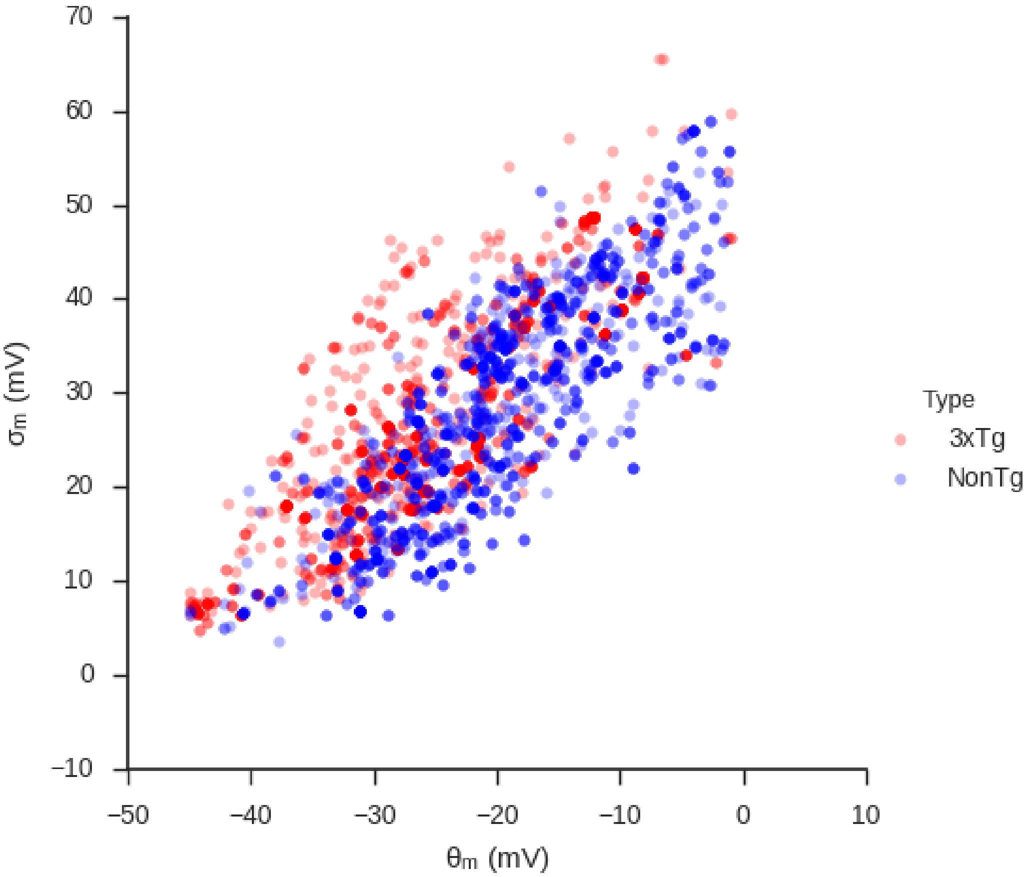

Figure 5.6: We can see by inspection of the figure that θ_m and σ_m are modestly collinear. This can also be seen by calculating the eigenvalues and eigenvectors of the correlation matrix between all features. In two dimensions the separation between the two classes of neurons becomes stronger than in the case of either of θ_m or σ_m individually.

Table 5.4: Feature importances for rheobase using only model parameters. I_{DC} is itself a current and τ_{mem} describes the extent to which the membrane smoothes excitatory currents. Alterations in I_{Na}, rather than I_K, are present again as well. For example, larger g_{Na} predicts a lesser rheobase. The values of θ_m and θ_h, parameters describing voltages where I_{Na} become activated and inactivated, respectively, are often more positive when more current is required to elicit an action potential.

All Params Predicting Rheobase	Var Imp	Correlation
θ_h (mV)	0.154	0.183
τ_{mem} (ms)	0.111	-0.283
I_{DC} (pA/pF)	0.098	-0.146
g_{Na} (nS)	0.077	-0.295
θ_m (mV)	0.063	0.224

Table 5.5: Overlap was calculated between probability density functions for the two strains in three and four dimensional subspaces representing a particular strain. Subspaces achieving the best separation are in accordance with the top ranked random forest variable importance features. The probability distributions are kernel-density estimates using Gaussian kernels. Details about the specific implementation can be found in the scipy documentation. We repeated the calculation for all four dimensional feature subspaces. In four dimensions the subspace (Threshold Voltage (mV), $\dot{I}_{Na}(V_{thres})$ Downstroke (pA/pF/ms), $\dot{I}_{Na}(V_{thres})$ Upstroke (pA/pF/ms), $\dot{I}_{K}(V_{thres})$ Downstroke (pA/pF/ms)) has the minimal overlap of 3.6%.

3D Subspace	Overlap Percentage
(C_m (pF), E_L (mV), τ_{mem} (ms))	0.213
(Rheobase (pA/pF), Threshold Voltage (mV), I_{DC} (pA))	0.368

achieving the best separation were consistent with highly predictive features of the random forest models (Table 5.5). As the overlap percentages alone do not give sufficient information to determine the nature of the separations, visualizing the feature separations is mandatory. Passive membrane properties tended to give complex separation surfaces (Figure 5.5), while separations using kinetic properties of ionic current tended to be simpler (Figure 5.7).

As the dimension of the feature subspaces increases, their number increases factorially, while computational cost of calculating the KDEs grows exponentially. Therefore, it is not feasible to scale this approach up to higher dimensions, and evaluating feature importances using random forest models is preferred.

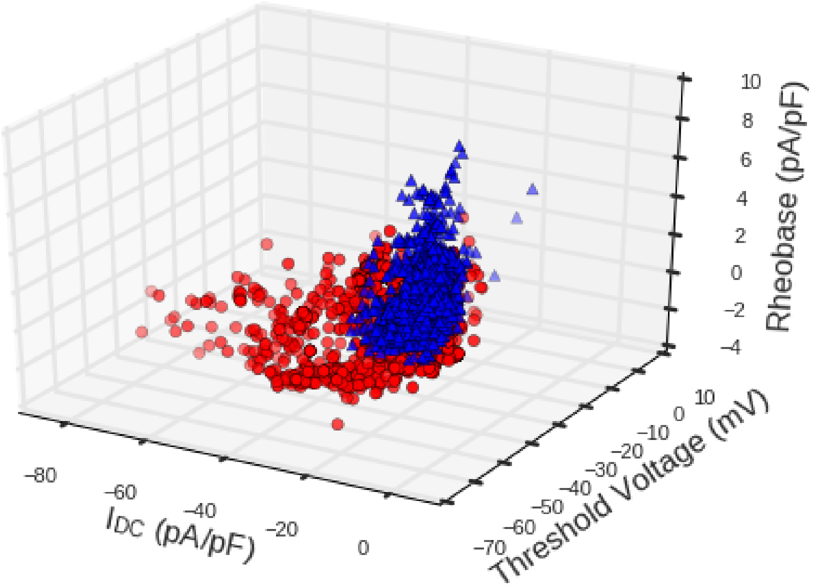

Figure 5.7: One subspace demonstrating a simple separation with minimal overlap between the neuron strains. The separating features include those associated with neuron excitability and are consistent with features selected by the random forest model.

5.3 Discussion

Electrophysiological experiments typically carried out to characterize membrane properties are tedious and laborious. The assimilation filter and data mining procedure presented here together constitute a less time consuming and expensive approach. With this approach, it is possible to infer the electrical properties, in detail, across a sizeable population of neurons. This makes it possible to characterize variability and to discover multifaceted differences which might otherwise be lost due to low sample size and averaging. This is a step towards further understanding neuronal dysfunction in AD, with possible application to other neuronal disorders, potentially opening a path towards evaluating the efficacy of novel potential therapeutics or, conversely, detecting adverse of effects of compounds targeting other mechanisms.

While the Aβ hypothesis continues remarkably unabated, as yet no Aβ-directed therapeutic has been successful in human clinical trials. Many of the existing therapeutics alter properties of ion channels. Principal therapeutics remain cholinesterase inhibitors, meant to boost acetycholine (ACh) levels. ACh acts through nicotinic receptors (ligand gated ion channels) and muscarinic receptors. These are G-protein coupled receptors that could influence Ca^{2+} signaling through IP_3R, neuron excitability through M-channels and G protein regulated channels. Another therapeutic is memantine, an NMDA-sensitive receptor antagonist. NMDA-Rs are ligand gated ion channels known for their Ca^{2+} flux. These drugs may temporarily treat some of the symptoms of AD, but they do not halt the progression. It is clear that we need a more comprehensive understanding of the impact of drugs on neuronal ion channel functionality. A means to characterize properties of ion channels through a simple, rapid measurement, rather than picking each apart one channel at a time, would be of considerable benefit.

A present obstacle is the model identifiability problem arising from the numerical difficulty of assimilating data to biophysical models which are too complex. Here we have limited our analysis to detecting alterations in intrinsic excitability of neurons. Consistent with other work bath applying extracellular Aβ to 1 month old WT mice, we have found a hyperpolarized threshold voltage [85, 86] as well as a reduced rheobase in 3xTG mice. However, our study uses relatively young mice, with little or no Aβ accumulation and no neurofibrillary tangles. Our combined modeling and data mining approach therefore may have identified detectable electrophysiological functional differences between young AD mouse model neurons and nontransgenic controls. Our study additionally provides possible mechanisms underlying these altered features. We mainly find alterations in the I_{DC} component of the leak current and I_{Na}. Others have found a reduction in I_{Na} [10] or reductions in I_K [72]. In contrast, we primarily see alterations in the gating kinetics of I_{Na}.

We do not find an alteration in spike width. Reductions [85, 86] and broadening [72] in spike width have been found elsewhere. The discrepencies are possibly due to differing mouse models used across studies.

An increased AHP through Ca^{2+}-sensitive K^+ conductances has been noted elsewhere [36, 81]. Our model with only I_{Na} and I_K cannot detect this AHP, as it requires a slower current such as I_{SK} which is sensitive to intracellular Ca^{2+}. Incorporating intracellular Ca^{2+} dynamics will be the goal of future work.

Chapter 5 is reproduced from material as it appears in a preprint being prepared for publication. Use of Data Assimilation for Inference of CA1 Neuron Pathology in 3xTg Mouse Model of Alzheimers. Daniel Breen, Clark Briggs, Grace Stutzmann, and Henry DI Abarbanel. The dissertation author was the first author on this preprint.

Chapter 6

Future Opportunities

There are many things that I didn't have the time or the expertise to try. Some of these things, such as using data from voltage clamp and incorporating I_{SK} into the CA1 neuron model introduced in Chapter 5, seem straightforward and promising. Others, such as implementing new optimization algorithms, would involve expertise in subjects such as nonlinear optimization and high-performance computing. In this chapter, I discuss some possible directions for future work.

6.1 Data Assimilation using Voltage Clamp

In patch clamp experiments, the usual methodology has been to design and inject a current waveform into the neuron, eliciting a voltage response which is measured and recorded. This is called current clamp. Another experiment that is possible in patch clamp experiments is to instead control the membrane voltage directly, measuring the injected current that is required to pass through the membrane to cause the desired change in voltage. The reason that this experiment, called voltage clamp, might be preferable is that in current clamp mode, it is not always possible to distinguish the currents which cause depolarization or repolarization of the cell. The intuition is that there are so many ways to get a neuron to spike that fitting the model to a spike waveform might not constrain the parameters very much. For example, the transient sodium current I_{Na} and the L-type calcium channel I_{CaL} might both be responsible for depolarizing the cell membrane. In current clamp mode, this can cause model identifiability issues of the flavor discussed in previous chapters. Models with equal predictive power will be estimated using the data, but with quite different relative strengths of I_{CaL} and I_{Na}.

Voltage clamp mode can overcome these issues, as it is possible to control the voltage within a narrow range so that just one of the currents, say I_{Na}, activates, deactivates, or inactivates. In that case, the strength and gating kinetics of I_{Na} could then be characterized independently of I_{CaL}. After exploring the range of voltages at which I_{Na}, the voltage could be perturbed to find the regime at which I_{CaL} becomes active. Then I_{CaL} could be characterized, and the process could be repeated for other ionic currents. This is similar to experiments involving chemical agents which isolate the effects of a specific current by blocking the others. However, this ‘electrical pharmacology’ has an advantage over methodology involving chemical blockers. Chemical blockers can be only partially effective or alter other properties of the cell when they achieve their effect. Additionally, the effects of chemical blockers can be irreversible, which is highly undesirable in slice recordings as it can render the rest of the slice as useless. It is therefore desirable to have a technique which has the capability of isolating the effects of specific currents while simultaneously leaving the cell and the rest of the preparation unaltered.

An example of a voltage clamp protocol, with an explanation of reasoning underlying its design and potential use, is given in Figure 6.1.

Showing that the data assimilation procedure can work using data recorded in voltage clamp mode would likely be straightforward. Thus, this represents a good opportunity for future work. One potential limitation of voltage clamp is that, as a result of the fact that perturbations in the membrane voltage at the site of the patch clamp will

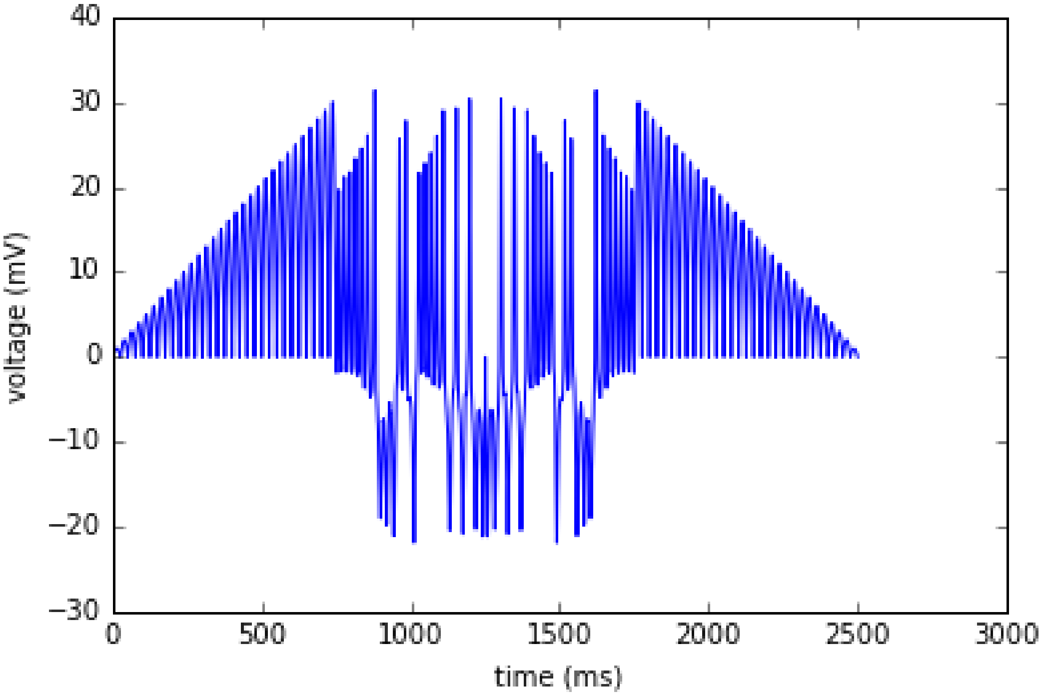

Figure 6.1: An example of a candidate voltage clamp protocol. The idea behind this waveform is to gradually ramp up the voltage until the threshold for activating the current is found. There are 30 steps, each incrementing by one volt after one oscillation. The steps return to baseline in order to explore the relaxation properties of whatever current is under investigation. A consideration in the design of voltage clamp protocols is that the voltage must not change too quickly. A sudden change in voltage will give erroneous values of the injected current. This artifact manifests because of charging of the pipette used to patch onto the cell. In order to use this waveform, it would be voltage shifted and temporally stretched to fit a particular problem. The voltage might be shifted by -100 mV, very hyperpolarized potentials, to explore the hyperpolarization activated cation current I_H and the passive properties of the membrane, including leak conductance and the RC time constant at subthreshold potentials. Similarly, the protocol would be shifted -70 mV to explore I_{CaT} or I_{NaP}, -55 mV to explore I_{Na}, and somewhere from -25 mV to -5 mV to explore high threshold calcium currents such as I_{CaL}.

decay spatially with increasing distance from the pipette, only local ion channels will be affected by voltage perturbations. Therefore, an experiment using voltage clamp might miss types of ion channels that are not near the site of the clamp, which is typically at or near the soma. Therefore, it may be desirable to check the results of data assimilation using a voltage clamp protocol against estimates obtained using a current clamp protocol.

6.2 Biophysical Neuron Modeling

6.2.1 Modeling Cellular Mechanisms of Alzheimer's Disease

An aspirational goal is to validate quantitative models which incorporate proposed disease mechanisms. We investigate potential mechanisms underlying one specific type of Alzheimer's which is developed much earlier in life and is genetically inherited. This type of Alzheimer's is called familial Alzheimer's disease. In mice, a condition resembling hereditary Alzheimer's which is used as a model for the disease can be generated in mice with expressing three human transgenes, PSEN1, APP, and MAPT. This is the triple transgenic mouse model, denoted 3xTg.

PSEN1, or Presenilin-1, is a gene encoding information about transmembrane proteins called presenilins which are involved in cleaving APP, the amyloid precursor protein. APP is also a transmembrane protein, which when cleaved gives rise to Aβ peptides. The presence of amyloid beta (Aβ) species are highly correlated with Alzheimer's disease. Aβ gathers into clumps which are the characteristic Aβ plaques found in the brains of people with Alzheimer's disease. Mutations in PSEN1 cause an increase in the Aβ42/Aβ40 ratio. Aβ42 is the peptide form which is the far more pathogenic Aβ oligomer and its presence is highly correlated with Alzheimer's disease.

The majority of pathogenic APP mutations cluster near the cleavage sites of the proteases β-secretase and γ-secretase, and generally increase total Aβ levels and/or the

Aβ42/Aβ40 ratio.

The physiological function of the tau protein is to bind tubulin and stabilize microtubules, which form the skeleton of the cell. In this way it supports cell differentiation, polarization, and other processes involving the cytoskeleton. Although it is not clear what the pathologic mechanism of the mutation of the MAPT gene is, it is thought that an aggregate propagates across synapses, leading to pathological aggregation and eventually to microtubule breakdown and neurodegeneration, causing the characteristic neurofibrillary tangles of Alzheimer's disease.

We investigate one proposed disease mechanism, altered dynamics of calcium concentration inside the cell caused by mutations in PSEN1 and APP. Increased baseline calcium levels could be mediated by calcium permeable pores formed by Aβ pores through the plasma membrane.

Presenilin can function as an endoplasmic reticulum (ER) membrane calcium leak channel, helping to maintain calcium homeostasis in the ER. The ER is a large organelle in cells which serves many functions. The one we are concerned with is that the calcium concentration in the ER is hundreds of times larger than in the surrounding cytosol. We would then expect that alterations in the functions of the ER could lead to pathology associated with alterations of calcium dynamics of the cell.

Mutations in presenilin may help cause Familial AD by rendering presenilin unable to pass calcium ions, though this is a controversial mechanism. Alterations in the calcium leak current in the ER and calcium balance between the ER and cytoplasm in the model may reflect this disease mechanism.

Presenilin also interacts with ryanodine receptors (RyRs), causing the ryanodine receptor to be more active at high calcium concentrations when the channel normally would be inhibited. RyRs are a class of large proteins cause release of calcium from the ER when calcium levels increase inside the cell. This is called calcium induced calcium release (CICR). Increased sensitivity of RyRs to CICR from the ER is observed in Alzheimer's disease. In a model, this could be reflected in alterations to the parameters governing the kinetic properties of the RyRs. An increased number of neuronal RyRs is another observed disease mechanism.

Inositol triphosphate (IP_3) is a messenger associated with calcium dynamics. Presenilin alters IP_3 receptor (IP_3R) activity, manifesting as an increase in the open probability of IP_3Rs at threshold IP_3 concentrations and a leakage of calcium from the ER "at rest" conditions. This may manifest in the model as a change in kinetic parameters governing the kinetics of IP_3Rs.

Finding alterations in these processes using measurements of calcium concentration and a detailed model of cellular calcium dynamics, such as that discussed in chapter 2 of this dissertation, is a goal of future work. It is possible, but technically challenging, to obtain measurements of the calcium concentration as a function of time. Additionally, at the time of writing, these measurements can be done at a frequency of order 10-10^2 Hz, far below the resolution of patch clamp recordings, 10^4 - 10^5 Hz. This limits the present scope of experiments that can be done. When better data becomes available, it could be desirable to drive the neuron in voltage clamp mode, rather than current clamp, to drive the calcium dynamics into a greater range of phase space. This will help to constrain the range of estimated model parameter values.

6.2.2 Refinements to Neuron Modeling

CA1 neurons are a class of large pyramidal neurons located in the hippocampus of the mammalian brain. The hippocampus is involved in memory and is a vulnerable region in Alzheimer's disease. It is thought that the L-type calcium current and potassium current increases with age in mouse models of Alzheimer's disease. Additionally, due to increased influx of calcium through the L-type channels, one set of calcium-dependent potassium channels is spatially close to L-type channels and therefore coupled to them. They constitute the calcium dependent potassium current I_{SK}, which is thought to be increased as well. Therefore, including these currents and a basic description of calcium dynamics in our neuron model is desirable.

Our baseline H-H model includes, in addition to I_{SK} and I_{Ca_L}, I_{Na} and I_K. I_{Na} is a fast activating and inactivating sodium current, while I_K is a current summarizing the effects of the delayed rectifier and other fast potassium channels. Although many other subthreshold ionic currents exist, trial and error suggests that some currents are sufficiently similar with another that a single measurement of the voltage is not sufficient to correctly distinguish the contributions from these currents. These currents include the potassium current I_A, I_K, and I_{BK}. I_H is difficult to distinguish from slow-afterhyperpolarization current I_{sAHP} and the muscarinic current I_M.

The proliferation of sets of model parameters within the H-H framework giving equal predictive power suggests that conductance based models may be too simplistic. If this is an underlying reason, this suggests that model parameters should not be taken literally and that simply adding more such currents does not go far in bridging the complexity gap between model and system which would remediate the issue of model error sufficiently. Many other components, including correctly modeling interactions among subunits of gating proteins, adding in many more compartments to capture the spatial information of neurons, adding flexibility to the shape of gating variables and relaxation times, including cell signaling processes modulating the kinetics of channels, and many others, would have to be included for parameters of a neuron model to be sufficiently interpretable to justify the hypothesis that adding complexity could aid the inference problem.

Other less fundamental, but still challening, obstacles include the fact that even small amounts of error in the voltage could derail estimating the parameters of a ‘sufficiently accurate', but very complex model. The numerical difficulty of the optimization procedure is another large hurdle to such an enterprise. In order to keep the numerical difficulty of the problem controlled and for the reasons discussed above, we opt out of including additional complexity in our model. Then the equations describing our H-H model are the following:

$$C\frac{dV}{dt} = (I_{Na} + I_K + I_{SK} + I_{CaL} + I_L + I_{inj})$$

This equation simply describes current conservation across the cell membrane. Given a measurement of only the voltage V as a function of time, we can determine the values of these ionic currents as a function of time. These ionic currents have the following form:

$$\begin{aligned} I_{Na} &= g_{Na} m^2 h (E_{Na} - V) \\ I_K &= = g_K n (E_K - V) \\ I_{SK} &= g_{SK} q (E_K - V) \\ I_{CaL} &= g_{CaL} s^2 \Phi_{GHK}(\mathrm{Ca}^{2+}, V) \end{aligned}$$

The various g_X are the maximal conductances for each ionic current. E_X are the reversal potentials. m, h, n, q, and s are the activation and inactivation variables for these currents. All but q have voltage dependent sigmoidal activation functions which relax to equilibrium with a voltage dependent time constant. q_∞ is a calcium dependent hill function, $q_\infty = \frac{Ca_{SK}^2}{Ca_{SK}^2 + K_{SK}^2}$, attempting only to model the calcium pool Ca_{SK}, considering the background calcium concentration of the bulk cytosol to be constant. q relaxes to its equilibrium value with a calcium dependent time constant. $\Phi_{GHK}(\mathrm{Ca}^{2+}, V)$ is the Goldman-Hodgkin-Katz flux formula, a more accurate form for describing calcium currents.

We could also consider a model lacking a term I_{CaL}, as its contribution to changes in membrane voltage might be small, and therefore negligible. Then the influx of calcium might be modeled as being proportional to I_{Na} instead of I_{CaL}, or some other measure. This would have the advantage of eliminating a potential source of underdeterminacy in the data assimilation problem and reducing the number of parameters to be estimated in the model. Then we would merely have for the voltage equation:

$$C\frac{dV}{dt} = (I_{Na} + I_K + I_{SK} + I_L + I_{inj})$$

As discussed in Chapter 5, biologically realistic computational models can be used in basic and clinical research, such as in the neuropharmacological process of drug design by prescreening a number of compounds for bench testing in the laboratory. By understanding the relationship between detailed biological mechanisms and the parameters in the computational models, the response of the neuron to different treatments can be predicted by simulation, guiding intuition and leading to better understanding of disease mechanisms.

Because of the numerical difficulty of data assimilation, computational resources limit the complexity of a model. Additionally, a model with many terms representing cell signaling and ion currents may be too underdetermined a problem, so that results of data assimilation cannot be trusted.

Model fits have made predictions about differences in the electrophysiological properties of recordings made from Alzheimers mice and wildtype controls, including the enhancement of calcium and potassium ion currents that have been experimentally observed previously in aged mice and in mouse models of Alzheimers disease. These results are displayed in Figures 6.2 and 6.3.

Despite being an underdetermined problem, we can examine the data for overall patterns in engineered features or the estimated parameters. Here, we might guess that the currents I_{Ca_L}, I_{SK}, and I_K are all increased in the 3xTG animals compared to controls. As we saw in chapter 5, data assimilation may be capable of answering broad questions about unobserved, altered processes within a system, so long as the model is approximately correct and not too complex for the number of measurements that are available. Here, however, the distribution over the possible range of the magnitudes of ionic currents is very large. There is an unaccecptably high variation in the contribution of different depolarizing and repolarizing currents. Sometimes, model estimates show I_{Ca_L} as primarily responsible for depolarizing the cell. Other sets of model parameters suggest this role actually belongs to I_{Na}, and still others give intermediate roles for I_{Na} and I_{Ca_L}.

For the purpose of making inferences about the properties of ionic currents for biophysical and medical applications, additional refinements to the model, improving the inference algorithm, and/or using additional measurements obtained from physiological experiments are necessary and a goal of future work. It may also be necessary to find appropriate statistical measures for hypothesis testing and uncertainty about observed differences between populations.

6.3 Probing Model Identifiability by Exploring the Surface of the Action

Recent empirical and theoretical work argues that in the case of training large artificial neural networks, local minima are not a significant problem [20, 28]. In the past, training neural networks, which involves solving large-scale non-convex optimization

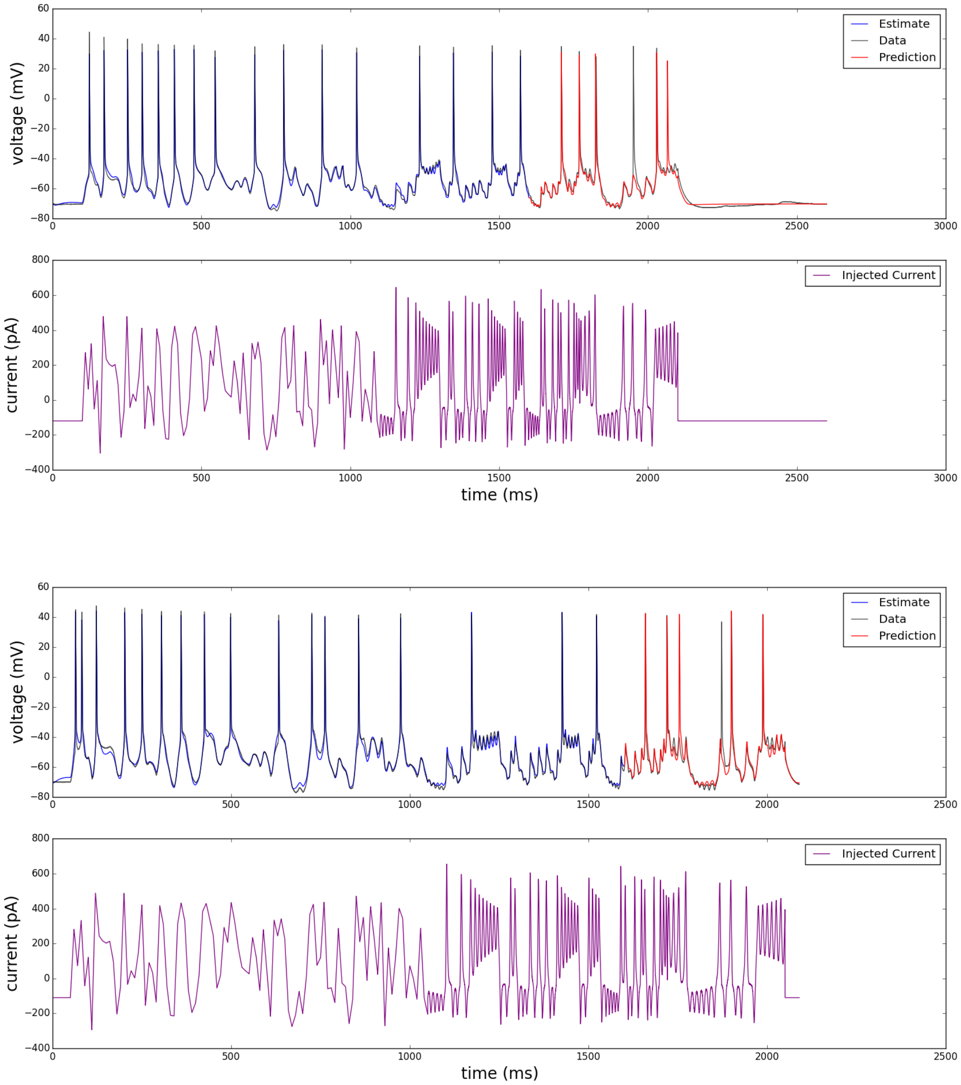

Figure 6.2: Data assimilation window and prediction window for a recording from 3xTG (top)and wildtype (bottom) mouse hippocampal CA1 neurons. Membrane voltage in response to injection of a complex current waveform (purple). The *black trace* shows recorded voltage, and the *blue trace* shows estimated voltage from the data assimilation procedure for times between 0-1,600 ms, during which all state variables and parameters of the model were estimated. The *red trace* shows the voltage predicted by integrating the completed model with estimated parameters and state variables forward in time beyond 1,600ms.

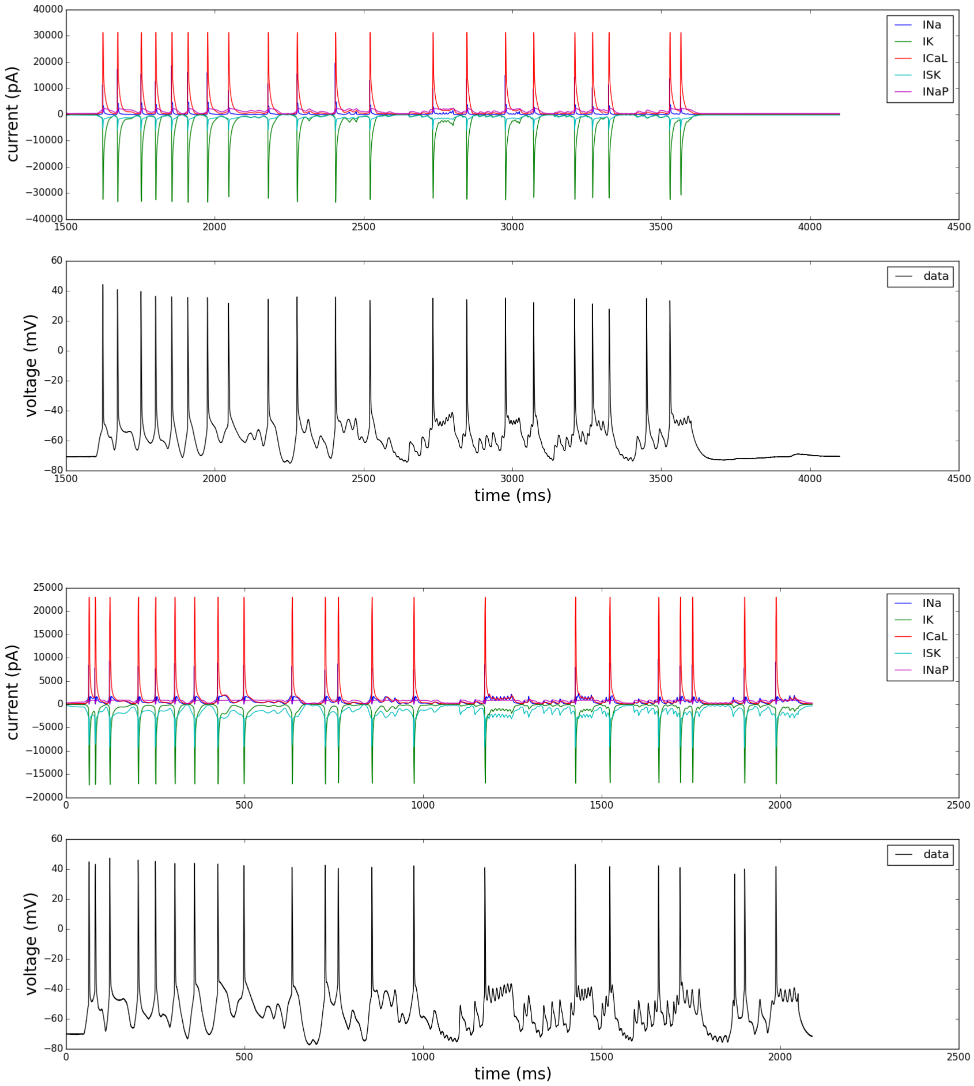

Figure 6.3: Ionic Currents for 3xTG (top) and wildtype neurons (bottom) inferred from their forms given in the model. Data is plotted in black. The magnitude of the calcium currents and potassium currents are much larger in the 3xTG than in the wildtype mouse, according to these particular model parameters. However, the reverse conclusion could also be drawn be examining other sets of model parameters producing accurate predictions. The estimated and predicted traces from these particular model fits are given in Figure 6.2.

problems like those considered in this dissertation, was long thought to be extremely difficult because of a proliferation of local minima and other obstacles. Recently, however, it has been found that on a straight path from initialization to solution, even state of the art neural networks never encounter any significant obstacles [28]. Finding this direction, of course, is not easy. Yet, although finding a good direction to search in a high dimensional space is a difficult problem, it is much easier than navigating an error surface that has complicated obstacles such as local minima or saddle points. This suggests that more mundane problems including poor conditioning are the primary difficulties in training neural networks.

Optimizing neural networks generally results in a proliferation of different solutions achieving very good performance. As some metric, such as predictive accuracy on some task, is generally important, such a proliferation of many good configuratons of the neural network is not an issue. According to recent work studying neural networks [20], a large number of saddle points, rather than local minima, trap the search procedure. These saddle points are thought to proliferate and become progressively more difficult to escape, with fewer search directions leading to lesser values of the cost function, as the cost function becomes smaller. Therefore, the authors of [20] develop a ‘saddle free’ Newton method which can rapidly escape high dimensional saddle points, unlike gradient descent and quasi-Newton methods.

Carrying out a similar analysis for the cost landscape in the parameter estimation problems considered in this disseration might yield some insight into what goes wrong during the search procedure. In addition to checking for a complicated structure of the landscape between random initializations and a solution, it would also be of interest to interpolate between any two solutions with different model parameters to check whether there are barriers of high cost in between them. It may be that many solutions yielding similar values of the cost function and having similar predictive power lie on a connected manifold. Applying other search procedures, such as the ‘saddle-free’ Netwon method of [20], may yield improvements over the present implementation of annealing using the open source software IPOPT. IPOPT deploys an interior point algorithm, including the use of line search and Newton’s method.

6.4 Parallelization

Another potential direction is to parallelize the computations to speed up the process of convergence to local optima. One possible method would be to parallelize the linear algebra computations while minimizing the action using a GPU. This would involve writing a custom algorithm, such as one using an interior point method. Another potential direction would be to develop highly parallelizable Monte Carlo algorithms to explore the action surface.

6.5 Characterizing a Network of Silicon Neurons

In Chapter 3, we characterized an individual silicon neuron on a neuromorphic chip and simulated data from an $\mathrm{HVC_I}$ neuron on it, all using data assimilation. The chip, NeuroDyn, consists of a small network of 4 such neurons which are synaptically connected. One natural extension to the work presented in this dissertation is to assimilate data from other kinds of neurons and emulate those neurons on the chip. Another is to characterize the synapses on the chip and emulate biological synapses in the silicon substrate. An additional direction might be emulating several synaptically connected biological neurons on the chip simultaneously and demonstrating that the methods of data assimilation can be extended to characterizing small real world circuits.

Bibliography

[1] Henry Abarbanel. *Predicting the future: completing models of observed complex systems*. Springer, 2013.

[2] Henry DI Abarbanel, P Bryant, Philip E Gill, Mark Kostuk, Justin Rofeh, Zakary Singer, Bryan Toth, Elizabeth Wong, and M Ding. Dynamical parameter and state estimation in neuron models. *The dynamic brain: an exploration of neuronal variability and its functional significance, Ch. 8*, 2011.

[3] Pablo Achard and Erik De Schutter. Complex parameter landscape for a complex neuron model. *PLoS Comput Biol*, 2(7):e94, 2006.

[4] Bruce P Bean. The action potential in mammalian central neurons. *Nature Reviews Neuroscience*, 8(6):451–465, 2007.

[5] UPINDER S Bhalla and JAMES M Bower. Exploring parameter space in detailed single neuron models: simulations of the mitral and granule cells of the olfactory bulb. *Journal of Neurophysiology*, 69(6):1948–1965, 1993.

[6] Daniela Bianchi, Addolorata Marasco, Alessandro Limongiello, Cristina Marchetti, Helene Marie, Brunello Tirozzi, and Michele Migliore. On the mechanisms underlying the depolarization block in the spiking dynamics of ca1 pyramidal neurons. *Journal of computational neuroscience*, 33(2):207–225, 2012.

[7] Daniel Breen, Sasha Shirman, Eve Armstrong, Nirag Kadakia, and Henry Abarbanel. Hvc interneuron properties from statistical data assimilation. *arXiv preprint arXiv:1608.04433*, 2016.

[8] Jochen Bröcker. On variational data assimilation in continuous time. *Quarterly Journal of the Royal Meteorological Society*, 136(652):1906–1919, 2010.

[9] Jochen Bröcker and Ivan G Szendro. Sensitivity and out-of-sample error in continuous time data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 138(664):785–801, 2012.

[10] Jon T Brown, Jeannie Chin, Steven C Leiser, Menelas N Pangalos, and Andrew D Randall. Altered intrinsic neuronal excitability and reduced na+ currents in a mouse model of alzheimer's disease. *Neurobiology of aging*, 32(11):2109–e1, 2011.

[11] Laure Buhry, Michele Pace, and Sylvain Saïghi. Global parameter estimation of an hodgkin–huxley formalism using membrane voltage recordings: Application to neuro-mimetic analog integrated circuits. *Neurocomputing*, 81:75–85, 2012.

[12] Marc Aurel Busche, Gerhard Eichhoff, Helmuth Adelsberger, Dorothee Abramowski, Karl-Heinz Wiederhold, Christian Haass, Matthias Staufenbiel, Arthur Konnerth, and Olga Garaschuk. Clusters of hyperactive neurons near amyloid plaques in a mouse model of alzheimer's disease. *Science*, 321(5896):1686–1689, 2008.

[13] Gert Cauwenberghs. Reverse engineering the cognitive brain. *Proceedings of the National Academy of Science (PNAS)*, 110:15512–15513, 2013.

[14] Oscar Cerda and James S Trimmer. Analysis and functional implications of phosphorylation of neuronal voltage-gated potassium channels. *Neuroscience letters*, 486(2):60–67, 2010.

[15] Shreaya Chakroborty and Grace E Stutzmann. Calcium channelopathies and alzheimer's disease: Insight into therapeutic success and failures. *European journal of pharmacology*, 739:83–95, 2014.

[16] Johanna L Crimins, Amy Pooler, Manuela Polydoro, Jennifer I Luebke, and Tara L Spires-Jones. The intersection of amyloid beta and tau in glutamatergic synaptic dysfunction and collapse in alzheimer's disease. *Ageing research reviews*, 12(3):757–763, 2013.

[17] Dávid Csercsik, Katalin M Hangos, and Gábor Szederkényi. Identifiability analysis and parameter estimation of a single hodgkin–huxley type voltage dependent ion channel under voltage step measurement conditions. *Neurocomputing*, 77(1):178–188, 2012.

[18] Viviana Culmone and Michele Migliore. Progressive effect of beta amyloid peptides accumulation on ca1 pyramidal neurons: a model study suggesting possible treatments. *Frontiers in computational neuroscience*, 6:52, 2012.

[19] Arij Daou, Matthew T Ross, Frank Johnson, Richard L Hyson, and Richard Bertram. Electrophysiological characterization and computational models of hvc neurons in the zebra finch. *Journal of neurophysiology*, 110(5):1227–1245, 2013.

[20] Yann N Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in neural information processing systems*, pages 2933–2941, 2014.

[21] ERIK De Schutter and JAMES M Bower. An active membrane model of the cerebellar Purkinje cell. I. simulation of current clamps in slice. *Journal of neurophysiology*, 71(1):375–400, 1994.

[22] Shaul Druckmann, Yoav Banitt, Albert A Gidon, Felix Schürmann, Henry Markram, and Idan Segev. A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data. *Frontiers in neuroscience*, 1:1, 2007.

[23] Michele Ferrante, Kim T Blackwell, Michele Migliore, and Giorgio A Ascoli. Computational models of neuronal biophysics and the characterization of potential neuropharmacological targets. *Current medicinal chemistry*, 15(24):2456–2471, 2008.

[24] Mike Fisher, Yannick Tremolet, Harri Auvinen, David Tan, and Paul Poli. Weak-constraint and long-window 4D-Var. *Techn. Rep*, 655, 2011.

[25] WR Foster, LH Ungar, and JS Schwaber. Significance of conductances in Hodgkin-Huxley models. *Journal of Neurophysiology*, 70(6):2502–2518, 1993.

[26] V Frazzini, S Guarnieri, M Bomba, R Navarra, C Morabito, MA Mariggiò, and SL Sensi. Altered kv2. 1 functioning promotes increased excitability in hippocampal neurons of an alzheimers disease mouse model. *Cell death & disease*, 7(2):e2100, 2016.

[27] Izrail M Gelfand and Sergej V Fomin. *Calculus of variations*. Courier Corporation, 2000.

[28] Ian J Goodfellow, Oriol Vinyals, and Andrew M Saxe. Qualitatively characterizing neural network optimization problems. *arXiv preprint arXiv:1412.6544*, 2014.

[29] Richard HR Hahnloser, Alexay A Kozhevnikov, and Michale S Fee. An ultra-sparse code underliesthe generation of neural sequences in a songbird. *Nature*, 419(6902):65–70, 2002.

[30] John Hardy. The amyloid hypothesis for alzheimers disease: a critical reappraisal. *Journal of neurochemistry*, 110(4):1129–1134, 2009.

[31] Kevin H Hobbs and Scott L Hooper. Using complicated, wide dynamic range driving to develop models of single neurons in single recording sessions. *Journal of neurophysiology*, 99(4):1871–1883, 2008.

[32] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952. 12991237[pmid] J Physiol.

[33] JR Huguenard. Low-threshold calcium currents in central nervous system neurons. *Annual review of physiology*, 58(1):329–348, 1996.

[34] Dezhe Z Jin, Fethi M Ramazanoğlu, and H Sebastian Seung. Intrinsic bursting enhances the robustness of a neural network model of sequence generation by avian brain area hvc. *Journal of computational neuroscience*, 23(3):283–299, 2007.

[35] Ranu Jung, Elizabeth J Brauer, and James J Abbas. Real-time interaction between a neuromorphic electronic circuit and the spinal cord. *IEEE Transactions on neural systems and rehabilitation engineering*, 9(3):319–326, 2001.

[36] CC Kaczorowski, E Sametsky, S Shah, R Vassar, and JF Disterhoft. Mechanisms underlying basal and learning-related intrinsic excitability in a mouse model of alzheimer's disease. *Neurobiology of aging*, 32(8):1452–1465, 2011.

[37] Armstrong Eve Breen Daniel Morone Uriel Daou Arij Margoliash Daniel Abarbanel Henry D.I. Kadakia, Nirag. Nonlinear statistical data assimilation for hvc-ra neurons in the avian song system. *Biological cybernetics*, 2016.

[38] N Kadakia, D Rey, J Ye, and HDI Abarbanel. Symplectic structure of statistical variational data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 2016.

[39] Nirag Kadakia, Eve Armstrong, Daniel Breen, Uriel Morone, Arij Daou, Daniel Margoliash, and Henry DI Abarbanel. Nonlinear statistical data assimilation for hvc_ {\ mathrm {RA}} neurons in the avian song system. *Biological Cybernetics*, pages 1–18.

[40] Nirag Kadakia, Eve Armstrong, Daniel Breen, Uriel Morone, Arij Daou, Daniel Margoliash, and Henry DI Abarbanel. Nonlinear statistical data assimilation for hvc_{ra} neurons in the avian song system. *Biological cybernetics*, 110(6):417–434, 2016.

[41] Naomi Keren, Noam Peled, and Alon Korngreen. Constraining compartmental models using multiple voltage recordings and genetic algorithms. *Journal of neurophysiology*, 94(6):3730–3742, 2005.

[42] TL Kerrigan, JT Brown, and AD Randall. Characterization of altered intrinsic excitability in hippocampal ca1 pyramidal cells of the aβ-overproducing pdapp mouse. *Neuropharmacology*, 79:515–524, 2014.

[43] Igor Klyubin, William K Cullen, Neng-Wei Hu, and Michael J Rowan. Alzheimers disease aβ assemblies mediating rapid disruption of synaptic plasticity and memory. *Molecular brain*, 5(1):25, 2012.

[44] Chris Knowlton, C Daniel Meliza, Daniel Margoliash, and Henry DI Abarbanel. Dynamical estimation of neuron and network properties iii: network analysis using neuron spike times. *Biological cybernetics*, 108(3):261–273, 2014.

[45] Mark Kostuk, Bryan A Toth, C Daniel Meliza, Daniel Margoliash, and Henry DI Abarbanel. Dynamical estimation of neuron and network properties ii: path integral monte carlo methods. *Biological cybernetics*, 106(3):155–167, 2012.

[46] Mark Kot. *A first course in the calculus of variations*, volume 72. American Mathematical Society, 2014.

[47] Michinori Kubota and Ikuo Taniguchi. Electrophysiological characteristics of classes of neuron in the hvc of the zebra finch. *Journal of Neurophysiology*, 80(2):914–923, 1998.

[48] Gwendal Le Masson, Sylvie Renaud-LeMasson, Damien Debay, and Thierry Bal. Feedback inhibition controls spike transfer in hybrid thalamic circuits. *Nature*, 417:854–858, 2002.

[49] Sylvie Le Masson, A Laflaquiére, Thierry Bal, and Gwendal LeMasson. Analog circuits for modeling biological neural networks: design and applications. *IEEE Transactions on Biomedical Engineering*, 46:638–645, 1999.

[50] Daniel Liberzon. *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2012.

[51] Michael A Long, Dezhe Z Jin, and Michale S Fee. Support for a synaptic chain model of neuronal sequence generation. *Nature*, 468(7322):394–399, 2010.

[52] Eve Marder and Dirk Bucher. Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annu. Rev. Physiol.*, 69:291–316, 2007.

[53] Mark P Mattson, Frank M LaFerla, Sic L Chan, Malcolm A Leissring, P Nickolas Shepel, and Jonathan D Geiger. Calcium signaling in the er: its role in neuronal plasticity and neurodegenerative disorders. *Trends in neurosciences*, 23(5):222–229, 2000.

[54] DAVID A McCORMICK and Hans-Christian Pape. Properties of a hyperpolarization-activated cation current and its role in rhythmic oscillation in thalamic relay neurones. *The Journal of physiology*, 431:291, 1990.

[55] Carver Mead. *Analog VLSI and Neural Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1989.

[56] C Daniel Meliza, Mark Kostuk, Hao Huang, Alain Nogaret, Daniel Margoliash, and Henry DI Abarbanel. Estimating parameters and predicting membrane voltages with conductance-based neuron models. *Biological cybernetics*, 108(4):495–516, 2014.

[57] Durga P Mohapatra, Hiroaki Misonou, Pan Sheng-Jun, Joshua E Held, D James Surmeier, and James S Trimmer. Regulation of intrinsic excitability in hippocampal neurons by activity-dependent modulation of the kv2. 1 potassium channel. *Channels*, 3(1):46–56, 2009.

[58] Tom Morse, Nicholas T Carnevale, Pradeep Mutalik, Michele Migliore, and Gordon M Shepherd. Abnormal excitability of oblique dendrites implicated in early alzheimer's: a computational study. *Frontiers in neural circuits*, 4:16, 2010.

[59] Farzan Nadim, Øystein H Olsen, Erik De Schutter, and Ronald L Calabrese. Modeling the leech heartbeat elemental oscillator I. interactions of intrinsic and synaptic currents. *Journal of computational neuroscience*, 2(3):215–235, 1995.

[60] Emre Ozgur Neftci, Bryan Toth, Giacomo Indiveri, and Henry DI Abarbanel. Dynamic state and parameter estimation applied to neuromorphic systems. *Neural Computation*, 24(6):1, 2011.

[61] Alain Nogaret, C Daniel Meliza, Daniel Margoliash, and Henry DI Abarbanel. Automatic construction of predictive neuron models through large scale assimilation of electrophysiological data. *Scientific Reports*, 6, 2016.

[62] Jakub Nowacki, Hinke M Osinga, Jon T Brown, Andrew D Randall, and Krasimira Tsaneva-Atanasova. A unified model of ca1/3 pyramidal cells: An investigation into excitability. *Progress in biophysics and molecular biology*, 105(1):34–48, 2011.

[63] Salvatore Oddo, Antonella Caccamo, Masashi Kitazawa, Bertrand P Tseng, and Frank M LaFerla. Amyloid deposition precedes tangle formation in a triple transgenic model of alzheimers disease. *Neurobiology of aging*, 24(8):1063–1070, 2003.

[64] Salvatore Oddo, Antonella Caccamo, Jason D Shepherd, M Paul Murphy, Todd E Golde, Rakez Kayed, Raju Metherate, Mark P Mattson, Yama Akbari, and Frank M LaFerla. Triple-transgenic model of alzheimer's disease with plaques and tangles: intracellular aβ and synaptic dysfunction. *Neuron*, 39(3):409–421, 2003.

[65] J. Park, S. Ha, T. Yu, E. Neftci, and G. Cauwenberghs. A 65k-neuron 73-mevents/s 22-pj/event asynchronous micro-pipelined integrate-and-fire array transceiver. In *2014 IEEE Biomedical Circuits and Systems Conference (BioCAS) Proceedings*, pages 675–678, Oct 2014.

[66] L Pierre-Simon. Memoir on the probability of the causes of events. *Stat. Sci*, 1(3):364–378, 1774.

[67] RD Pinto, RC Elson, A Szucs, MI Rabinovich, AI Selverston, and HDI Abarbanel. Extended dynamic clamp: controlling up to four neurons using a single desktop computer and interface. *Journal of Neuroscience Methods*, 108:39–48, 2001.

[68] Astrid A Prinz, Cyrus P Billimoria, and Eve Marder. Alternative to hand-tuning conductance-based models: construction and analysis of databases of model neurons. *Journal of neurophysiology*, 90(6):3998–4015, 2003.

[69] Astrid A Prinz, Dirk Bucher, and Eve Marder. Similar network activity from disparate circuit parameters. *Nature neuroscience*, 7(12):1345–1352, 2004.

[70] Michael S Reid, Edgar A Brown, and Stephen P DeWeerth. A parameter-space search algorithm tested on a Hodgkin–Huxley model. *Biological cybernetics*, 96(6):625–634, 2007.

[71] Sylvain Saïghi, Yannick Bornat, Jean Tomas, and Sylvie Renaud. Neuromimetic ics and system for parameters extraction in biological neuron models. In *Proc. 2006 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 4207–4210, 2016.

[72] Federico Scala, Salvatore Fusco, Cristian Ripoli, Roberto Piacentini, Domenica Donatella Li Puma, Matteo Spinelli, Fernanda Laezza, Claudio Grassi, and Marcello D'Ascenzo. Intraneuronal aβ accumulation induces hippocampal neuron hyperexcitability through a-type k+ current inhibition mediated by activation of caspases and gsk-3. *Neurobiology of aging*, 36(2):886–900, 2015.

[73] David J Schulz, Jean-Marc Goaillard, and Eve Marder. Variable channel expression in identified single and electrically coupled neurons in different animals. *Nature neuroscience*, 9(3):356–362, 2006.

[74] Dennis J Selkoe. Alzheimer's disease is a synaptic failure. *Science*, 298(5594):789–791, 2002.

[75] Ganesh M Shankar, Shaomin Li, Tapan H Mehta, Amaya Garcia-Munoz, Nina E Shepardson, Imelda Smith, Francesca M Brett, Michael A Farrell, Michael J Rowan, Cynthia A Lemere, et al. Amyloid-β protein dimers isolated directly from alzheimer's brains impair synaptic plasticity and memory. *Nature medicine*, 14(8):837–842, 2008.

[76] Baiyang Sheng, Xinglong Wang, Bo Su, Hyoung-gon Lee, Gemma Casadesus, George Perry, and Xiongwei Zhu. Impaired mitochondrial biogenesis contributes to mitochondrial dysfunction in alzheimers disease. *Journal of neurochemistry*, 120(3):419–429, 2012.

[77] Michael Sorensen, Stephen DeWeerth, Gennady Cymbalyuk, and Ronald L Calabrese. Using a hybrid neural system to reveal regulation of neuronal network activity by an intrinsic current. *Journal of Neuroscience*, 24(23):5427–5438, 2004.

[78] Cristina Soto-Trevino, Pascale Rabbah, Eve Marder, and Farzan Nadim. Computational model of electrically coupled, intrinsically distinct pacemaker neurons. *Journal of neurophysiology*, 94(1):590–604, 2005.

[79] Grace E Stutzmann. Calcium dysregulation, ip3 signaling, and alzheimers disease. *The Neuroscientist*, 11(2):110–115, 2005.

[80] Grace E Stutzmann. The pathogenesis of alzheimers diseaseis it a lifelong calciumopathy? *The Neuroscientist*, 13(5):546–559, 2007.

[81] Grace E Stutzmann, Antonella Caccamo, Frank M LaFerla, and Ian Parker. Dysregulated ip3 signaling in cortical neurons of knock-in mice expressing an alzheimer's-linked mutation in presenilin1 results in exaggerated ca2+ signals and altered membrane excitability. *Journal of Neuroscience*, 24(2):508–513, 2004.

[82] Grace E Stutzmann, Ian Smith, Antonella Caccamo, Salvatore Oddo, Ian Parker, and Frank Laferla. Enhanced ryanodine-mediated calcium release in mutant ps1-expressing alzheimer's mouse models. *Annals of the New York Academy of Sciences*, 1097(1):265–277, 2007.

[83] Attila Szücs, Pablo Varona, Alexander R Volkovskii, Henry DI Abarbanel, Mikhail I Rabinovich, and Allen I Selverston. Interacting biological and electronic neurons generate realistic oscillatory rhythms. *Neuroreport*, 11(3):563–569, 2000.

[84] O Talagrand. 4D-VAR: four-dimensional variational assimilation. *Advanced Data Assimilation for Geosciences: Lecture Notes of the Les Houches School of Physics: Special Issue, June 2012*, page 1, 2014.

[85] Francesco Tamagnini, Janet Novelia, Talitha L Kerrigan, Jon T Brown, Krasimira Tsaneva-Atanasova, and Andrew D Randall. Altered intrinsic excitability of hippocampal ca1 pyramidal neurons in aged pdapp mice. *Frontiers in cellular neuroscience*, 9, 2015.

[86] Francesco Tamagnini, Sarah Scullion, Jon T Brown, and Andrew D Randall. Intrinsic excitability changes induced by acute treatment of hippocampal ca1 pyramidal neurons with exogenous amyloid β peptide. *Hippocampus*, 25(7):786–797, 2015.

[87] Anna M Taylor and Roger M Enoka. Optimization of input patterns and neuronal properties to evoke motor neuron synchronization. *Journal of computational neuroscience*, 16(2):139–157, 2004.

[88] Bryan A Toth, Mark Kostuk, C Daniel Meliza, Daniel Margoliash, and Henry DI Abarbanel. Dynamical estimation of neuron and network properties i: variational methods. *Biological cybernetics*, 105(3-4):217–237, 2011.

[89] Roger D Traub, Robert K Wong, Richard Miles, and Hillary Michelson. A model of a CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *Journal of Neurophysiology*, 66(2):635–650, 1991.

[90] Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.

[91] J. Wang, D. Breen, A. Akinin, H. D. I. Abarbanel, and G. Cauwenberghs. Data assimilation of membrane dynamics and channel kinetics with a neuromorphic integrated circuit. In *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 584–587, Oct 2016.

[92] Kenji Yamamoto, Yoshifumi Ueta, Li Wang, Ryo Yamamoto, Naoko Inoue, Kaoru Inokuchi, Atsu Aiba, Hideto Yonekura, and Nobuo Kato. Suppression of a neocortical potassium channel activity by intracellular amyloid-β and its rescue with homer1a. *The Journal of Neuroscience*, 31(31):11100–11109, 2011.

[93] J Ye, N Kadakia, PJ Rozdeba, HDI Abarbanel, and JC Quinn. Improved variational methods in statistical data assimilation. *Nonlinear Processes in Geophysics*, 22(2):205–213, 2015.

[94] Jingxin Ye, Daniel Rey, Nirag Kadakia, Michael Eldridge, Uriel I Morone, Paul Rozdeba, Henry DI Abarbanel, and John C Quinn. Systematic variational method for statistical nonlinear state and parameter estimation. *Physical Review E*, 92(5):052901, 2015.

[95] Jingxin Ye, Paul J Rozdeba, Uriel I Morone, Arij Daou, and Henry DI Abarbanel. Estimating the biophysical properties of neurons with intracellular calcium dynamics. *Physical Review E*, 89(6):062714, 2014.

[96] T. Yu and G. Cauwenberghs. Analog VLSI biophysical neurons and synapses with programmable membrane channel kinetics. *Biomedical Circuits and Systems, IEEE Transactions on*, 4(3):139–148, 2010.

[97] T. Yu, T. J. Sejnowski, and G. Cauwenberghs. Biophysical neural spiking, bursting, and excitability dynamics in reconfigurable analog vlsi. *IEEE Trans Biomed Circuits Syst*, 5(5):420–9, 2011.

[98] Hua Zhang, Jie Liu, Suya Sun, Ekaterina Pchitskaya, Elena Popugaeva, and Ilya Bezprozvanny. Calcium signaling, excitability, and synaptic plasticity defects in a mouse model of alzheimer's disease. *Journal of Alzheimer's Disease*, 45(2):561–580, 2015.

[99] Xin Zou, Damien Coyle, KongFatt Wong-Lin, and Liam Maguire. Computational study of hippocampal-septal theta rhythm changes due to beta-amyloid-altered ionic channels. *PLoS One*, 6(6):e21579, 2011.