Title
The Dynamics of Nonlinear Inference

Author
Kadakia, Nirag

Publication Date
2017

UNIVERSITY OF CALIFORNIA, SAN DIEGO

The Dynamics of Nonlinear Inference

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

in

Physics

by

Nirag Kadakia

Committee in charge:

Professor Henry D. I. Abarbanel, Chair
Professor Tim Q. Gentner
Professor Michael Holst
Professor Kenneth Intriligator
Professor Hans Paar

2017

The Dissertation of Nirag Kadakia 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 Bijal.

EPIGRAPH

Do you ever just get down on your knees and thank God
that you know me and have access to my dementia?
– Costanza

TABLE OF CONTENTS

Signature Page iii

Dedication iv

Epigraph v

Table of Contents vi

List of Figures ix

List of Tables xi

Acknowledgements xii

Vita xiii

Abstract of the Dissertation xv

Chapter 1 Introduction 1
- 1.1 Invitation: the simple oscillator and a naive procedure 1
- 1.2 The ubiquity of inverse problems in scientific inquiry 3
- 1.3 Characteristics and challenges 4
- 1.4 The scope of this thesis 6

Chapter 2 Inference in nonlinear dynamical systems 9
- 2.1 Bayesian recursion as a starting point 9
  - 2.1.1 The Bayesian recursion equations 10
- 2.2 Kalman-based filtering and smoothing 12
  - 2.2.1 Kalman filter: linear dynamics and Gaussian errors 12
  - 2.2.2 Nonlinear extensions: extended and unscented Kalman filters 13
  - 2.2.3 Smoothing: backwards passes 14
  - 2.2.4 Ensemble Kalman filter: A partially statistical approach 15
- 2.3 Sequential Monte Carlo 16
  - 2.3.1 The bootstrap filter 19
  - 2.3.2 Resampling 20
- 2.4 The limitations of the sequential filter 20
  - 2.4.1 Filtering in neural estimation 20
  - 2.4.2 Breakdown of the Gaussian assumption 21
- 2.5 Beyond filtering: the variational approximation 26
  - 2.5.1 The simultaneous conditional expectation 26
  - 2.5.2 The ubiquity of local minima in $A(\mathbf{X}|\mathbf{Y})$ when model linearity is broken . 28
  - 2.5.3 The Gaussian error approximation 30
  - 2.5.4 Manifold structure of $A(\mathbf{X}|\mathbf{Y})$ 31

Chapter 3 Variational annealing 35
- 3.1 The variational annealing algorithm 35
  - 3.1.1 Variational annealing as homotopy continuation 37
- 3.2 Inference in chaotic systems with variational annealing 38
  - 3.2.1 Annealing plots 38
  - 3.2.2 Reliable estimation of sparsely observed Lorenz 1996 model 38
  - 3.2.3 Estimation quality of globally minimizing paths 43

3.3 The local vs. the global perspective 46
3.4 Parameter estimation in variational annealing 47
3.4.1 Two-variable Fitzhugh-Nagumo neuron model 48

Chapter 4 Nonlinear estimation in neural systems 52
4.1 The Hodgkin-Huxley neuron 53
4.1.1 Hopf bifurcations in the HH model 56
4.2 Projection neurons in the zebra finch HVC 58
4.2.1 Modeling nonlinear biophysical neurons and the role of inference 58
4.2.2 Zebra finch HVC 60
4.2.3 Sparse, robust bursting in HVC projection neurons 61
4.3 Modeling the $\mathrm{HVC}_{\mathrm{RA}}$ neuron 62
4.3.1 Known anatomical features 62
4.3.2 Calcium channels and the GHK equation 63
4.3.3 Calcium-gated potassium channels and calcium dynamics 65
4.3.4 Multiple timescales and the origin of bursting 66
4.3.5 Compartmental modeling and robust bursting 68
4.3.6 Full HVC model description and parameterization 70
4.3.7 Qualitative experimental predictions 71
4.4 Inferring the $\mathrm{HVC}_{\mathrm{RA}}$ model parameters 77
4.4.1 Information from transients 77
4.4.2 Timescale separation and an experimental protocol 78
4.4.3 Estimation of slow subsystem: computational details 79
4.4.4 Estimation of slow subsystem: results 82
4.4.5 Full model with slow parameters fixed 84
4.5 Degeneracies in the fast-slow $\mathrm{HVC}_{\mathrm{RA}}$ model: A bifurcation study 89
4.6 Inference in neural networks and the challenges ahead 91

Chapter 5 Continuous time and the path integral of nonlinear inference 93
5.1 Stochastic path integrals and the Onsager-Machlup function 94
5.2 The Onsager-Machlup Euler-Lagrange equations 97
5.2.1 Deterministic dynamics: The traceless Jacobian 99
5.2.2 Instability in the classical motion 100
5.3 The dynamical formulation of nonlinear inference 102
5.3.1 The electromagnetic analogy 106
5.3.2 E-M dynamics of the two-dimensional linear system 107

Chapter 6 A Hamiltonian formulation of nonlinear dynamical inference 109
6.1 Boundary value problems in optimal control 109
6.1.1 Instabilities in Hamiltonian dynamics and the shooting method 110
6.1.2 Multiple shooting 111
6.1.3 Collocation 114
6.2 Toward a Hamiltonian formulation of nonlinear inference 114
6.3 Symplectic invariants in Hamiltonian systems 116
6.3.1 Symplectic integrators 118
6.3.2 Exploiting symplecticity in nonlinear dynamical inference 123
6.4 Symplecticity from discretized classical actions 124
6.4.1 Discrete Euler-Lagrange equations and symplectic flows 124
6.4.2 From discrete actions to variational integrators 127
6.4.3 Discrete variational Hamiltonian mechanics 129
6.5 A Hamiltonian algorithm for nonlinear estimation 130
6.5.1 Symplecticity, boundary conditions, and degeneracies 133
6.6 Numerical experiments of a chaotic attractor 134

6.6.1 Implementation details 135
6.6.2 Results 139
6.6.3 Discussion 144

Chapter 7 Final thoughts 146
7.1 Hamiltonian variational annealing and chaotic synchronization 147
7.2 Chaos and correlations in Monte Carlo sampling 149

Bibliography 151

LIST OF FIGURES

Figure 1.1: Observations of an harmonic system. 2

Figure 2.1: Estimated trajectories of a chaotic model with dense observations. 22
Figure 2.2: Estimated trajectories of a chaotic model with sparse observations. 24
Figure 2.3: Posterior conditional distributions in the estimation of a chaotic attractor with a strong prior. 25
Figure 2.4: Posterior conditional distributions in the estimation of a chaotic attractor with a weak prior. 26
Figure 2.5: The ubiquity of local minima in the cost surface of an estimation of a chaotic dynamical system. 29
Figure 2.6: Visualization of the measurement and model terms in the cost surface manifold $A(\mathbf{X}|\mathbf{Y})$ in the estimation of a chaotic system. 32
Figure 2.7: Lifting of degeneracy in $A(\mathbf{X}|\mathbf{Y})$ manifold as a function of model precision. 34

Figure 3.1: Traces of observed data for variational annealing numerical experiments on the Lorenz 1996 attractor. 39
Figure 3.2: Annealing plots for variational annealing on Lorenz 1996 attractor. 41
Figure 3.3: Visualization of degenerate optima in early stages of variational annealing. 42
Figure 3.4: Visualization of the movement of $A(\mathbf{X}|\mathbf{Y})$ minima as degeneracy lifts in variational annealing. 42
Figure 3.5: Estimation quality in variational annealing for Lorenz 1996 system, as a function of system observability. 44
Figure 3.6: Visualization of optimal and sub-optimal estimates of chaotic Lorenz 1996 system as the variational annealing algorithm progresses. 45

Figure 4.1: Visualization of fixed points and limit cycles of the Hodgkin-Huxley neuron model. . . 57
Figure 4.2: Illustrative bifurcation diagrams of the Hodgkin-Huxley model. 59
Figure 4.3: Sparse $\mathrm{HVC}_{\mathrm{RA}}$ projection neuron bursting and dense $\mathrm{HVC}_{\mathrm{I}}$ inhibitory bursting in zebra finch auditory pathway during song motif. 61
Figure 4.4: Comparison of channel currents using the gating-ohmic expression in HH-model versus the constant field (GHK) expression. 65
Figure 4.5: Bifurcation diagram for generic HH-neuron with added Ca currents and Ca-gated K currents. 68
Figure 4.6: Bursting behavior in an HH-neuron model with slow Ca dynamics and Ca and Ca-gated K channels, in response to steady injected currents. 69
Figure 4.7: Bursting behavior of HH-nueron with calcium dynamics, in response to noisy injected currents. 69
Figure 4.8: Stereotyped bursting in two-compartment $\mathrm{HVC}_{\mathrm{RA}}$ model with calcium dynamics. . . . 72
Figure 4.9: Effect of calcium agonists on burst incidence in two-compartment $\mathrm{HVC}_{\mathrm{RA}}$ model. . . . 74
Figure 4.10: $\mathrm{HVC}_{\mathrm{RA}}$ model response to current pulses into the cell body. 75
Figure 4.11: $\mathrm{HVC}_{\mathrm{RA}}$ model response to current pulses into the cell body, with cells bathed in a calcium agonist. 76
Figure 4.12: Injected somatic current used to generate the data for the numerical experiments on the $\mathrm{HVC}_{\mathrm{RA}}$ model. 78
Figure 4.13: Synthetic observations for numerical experiments in the reduced $\mathrm{HVC}_{\mathrm{RA}}$ neuron model. 79
Figure 4.14: Annealing plot for the variational annealing estimation of the reduced two-compartment $\mathrm{HVC}_{\mathrm{RA}}$ model. 83
Figure 4.15: Estimates of a few select state variables of the reduced two-compartment $\mathrm{HVC}_{\mathrm{RA}}$ model. 84
Figure 4.16: Trace of the synthetic voltage observations for the estimation of the full $\mathrm{HVC}_{\mathrm{RA}}$ model. 85
Figure 4.17: Time traces of estimated variables in the full $\mathrm{HVC}_{\mathrm{RA}}$ model. 87
Figure 4.18: Comparison of forward prediction of the full $\mathrm{HVC}_{\mathrm{RA}}$ model using estimated parameters. 89

Figure 4.19: Comparison of forward prediction of the two-compartment $\mathrm{HVC_{RA}}$ model using estimated parameters, in response to a noisy dendritic current. 89
Figure 4.20: Bifurcation diagrams for the $\mathrm{HVC_{RA}}$ model description using either either assumed parameters or estimated parameters. 90

Figure 5.1: Phase plot and time trace of the nonlinear Van der Pol oscillator. 101
Figure 5.2: Phase plot and time trace of the nonlinear Van der Pol oscillator in Lagrangian coordinates. 102
Figure 5.3: Three-dimensional plot showing the unstable manifold of the Van der Pol oscillator when integrated in Lagrangian coordinates. 103

Figure 6.1: Estimates of chaotic system system by solving boundary problem with multiple shooting. 113
Figure 6.2: Evolution of phase space areas in an autonomous Hamiltonian system using a non-symplectic numerical scheme. 117
Figure 6.3: Evolution of phase space areas in Hamiltonian systems using a symplectic integrator. . . 119
Figure 6.4: Lowest action levels and averaged errors for estimation of the Lorenz 1996 system using variational annealing versus Hamiltonian variational annealing, with observation timestep $\Delta t = 0.05$. 141
Figure 6.5: Lowest action levels and averaged errors for estimation of the Lorenz 1996 system using variational annealing versus Hamiltonian variational annealing, with observation timestep $\Delta t = 0.01$. 142
Figure 6.6: Annealing plots for an outlier dataset in the estimation of the Lorenz 1996 system using variational annealing versus Hamiltonian variational annealing. 143

LIST OF TABLES

Table 3.1: Algorithmic parameters for variational annealing of 10-D Lorenz96 model. 38
Table 3.2: Number of unique minima among 10,000 initializations in the estimation of the chaotic Lorenz 1996 model, using direct optimization versus variational annealing. 40
Table 3.3: Likelihood of finding the global minimum of $A^{q,\beta_{\max}}$ among 10,000 initializations in the estimation of the chaotic Lorenz 1996 model, using direct optimization versus variational annealing. 40
Table 3.4: Algorithmic parameters for unscented Kalman filter applied to the Fitzhugh-Nagumo model. 49
Table 3.5: Parameter estimates of the Fitzhugh-Nagumo model using the unscented Kalman filter versus variational annealing. 50

Table 4.1: Parameters of Hodgkin-Huxley model. 56
Table 4.2: Constant field model parameters for various ions. 64
Table 4.3: Parameters of Hodgkin-Huxley model neuron with traditional currents, in addition to Ca currents, Ca dynamics, and Ca-gated potassium currents. 67
Table 4.4: Parameters of two-compartment $\mathrm{HVC_{RA}}$ neuron model. 73
Table 4.5: Upper and lower bounds on the parameters for the constrained optimization in the reduced $\mathrm{HVC_{RA}}$ model. 81
Table 4.6: Estimated parameters of reduced $\mathrm{HVC_{RA}}$ model, with only the somatic voltage measured. 85
Table 4.7: Upper and lower bounds on the parameters for the constrained optimization in the full $\mathrm{HVC_{RA}}$ model. 86
Table 4.8: Estimations of remaining parameters of full $\mathrm{HVC_{RA}}$ model, with only observations of the somatic voltage. 88

Table 6.1: Algorithmic parameters for comparison of variational annealing with symplectic and non-symplectic Hamiltonian variational annealing. 136
Table 6.2: Likelihood of finding the global minimum, averaged over all p, for the three annealing methods, for $L = 4$ measured variables. 141
Table 6.3: Likelihood of finding the global minimum, averaged over all p, for the three annealing methods, for $L = 5$ measured variables. 142

ACKNOWLEDGEMENTS

Thanks to Vivek Sharma for recognizing early on my love of teaching, and for pushing me to continue evolving as an educator and communicator. Thanks especially for your advice and criticisms, and for pushing me to think outside the norm.

Thanks to the UC San Diego Pysics Department staff – Sharmila, Catherine, Patti, Toni, Brad – for pairing me with professors from whom I would learn and grow as an educator, for scheduling my classes to accommodate my demands, for preparing and perfecting demonstrations, and for always advocating for the students.

Thanks to my group colleagues for colorful conversations and inspiration. Thanks especially to Paul for all of the feedback and insight, and to Dan Rey for setting the gold standard in scientific rigor and professionalism.

Thanks to Daniel Walsh for the winding conversations on physics, teaching, and pedagogy, and for peering with no judgment into even the most elementary of physical truths. The pristine and wondrous curiosity you exude in your pursuit of scientific understanding is a welcome reminder to those of us who occasionally forget that we too chose this field for the purest of reasons.

Thanks to my students for inspiring me, making me laugh, and challenging me to improve as a teacher.

Thanks to my advisor Henry Abarbanel for his encouragement, his ability to recognize and foster talent and interest, and for helping me through work I am sincerely proud of.

Thanks to my parents for their endless encouragement, motivation, and for feeding my curiosity as a child, for pushing me to strive in school, and for rejecting mediocrity. Thanks for the countless sacrifices you made so that I would have the opportunity to freely pursue what I love. Thanks to my sister Ashaini for her levity, and to all my family and friends for their encouragement and support.

Finally, thanks most of all to my wife Bijal, without whose unwavering support I would have long ago left graduate school, got a real job, and made considerably more money by this time.

VITA

2005	B. S. in Mathematics, University of Texas at Austin
2005	B. A. in Plan II Liberal Arts Honors, University of Texas at Austin
2005-2007	Mathematics teacher, Camden Academy Charter High School, Camden, NJ
2007-2009	Graduate research assistant, University at Albany
2009	M. S. in Nanoscale Science, University at Albany
2009-2012	Research scientist, Ion Beam Laboratory, University at Albany
2012-2017	Graduate teaching assistant, University of California, San Diego
2017	Ph. D. in Physics, University of California, San Diego

PUBLICATIONS

H. D. I. Abarbanel, E. Armstrong, D. Breen, N. Kadakia, D. Rey, S. Shirman, and D. Margoliash. A unifying view of synchronization for data assimilation in complex nonlinear networks. *Chaos*, in review.

N. Kadakia, D. Rey, J. Ye, and H. D. I. Abarbanel. Symplectic structure of statistical variational data assimilation. *Quarterly Journal of the Royal Meteorological Society*, **143**(703): 756-771 (2017).

D. Breen, S. Shirman, E. Armstrong, N. Kadakia, and H.D.I. Abarbanel. $\mathrm{HVC_I}$ neuron properties from statistical data assimilation. *arXiv*:1608.04433 [q-bio.NC] (2016).

N. Kadakia, E. Armstrong, D. Breen, U. Morone, A. Daou, D. Margoliash, and H. D. I. Abarbanel. Nonlinear statistical data assimilation for $\mathrm{HVC_{RA}}$ neurons in the avian song system. *Biological Cybernetics*, **110**(6): 417-434 (2016).

N. Kadakia. Hybrid Monte Carlo with chaotic mixing. *arXiv*:1604.07343 [physics.data-an] (2016).

J. Ye, D. Rey, N. Kadakia, M. Eldridge, U. Morone, P. Rozdeba, H. D. I. Abarbanel, and J. C. Quinn. Systematic variational method for statistical nonlinear state and parameter estimation. *Physical Review E*, **92**(5): 052901 (2015).

J. Ye, N. Kadakia, P. J. Rozdeba, H. D. I. Abarbanel, and J. C. Quinn. Improved variational methods in statistical data assimilation. *Nonlinear Processes in Geophysics*, **22**(2): 205-213 (2015).

P. Akhter, M. Huang, W. Spratt, N. Kadakia, and F. Amir. Tailoring the optical constants in single-crystal silicon with embedded silver nanostructures for advanced silicon photonics applications. *Journal of Applied Physics*, **117**(12): 123102 (2015).

P. Akhter, M. Huang, N. Kadakia, W. Spratt, G. Malladi, and H. Bakhru. Suppressing light reflection from polycrystalline silicon thin films through surface texturing and silver nanostructures. *Journal of Applied Physics*, **116**(11): 113503 (2015).

P. Akhter, N. Kadakia, G. Malladi, and M. Huang. Reduced reflection in polycrystalline silicon thin films through ion implantation induced surface textures and deposition of silver nanostructures. *MRS Conference Proceedings*, **1391**, DOI: 10.1557/opl.2012.1011 (2012).

N. Kadakia. Nanoparticle scattering, absorption, and interface effects for surface plasmon enhanced thin silicon solar cells: theory, past findings, and future directions. *Handbook of Research on Solar Energy Systems and Technologies*. Ed. Sohail Anwar. IGI Global (2012).

N. Kadakia, M. Huang, and H. Bakhru. Embedded silver nanoparticle fabrication for surface plasmon-enhanced silicon photovoltaics. *SPIE Conference Proceedings*, **8111**: 811101 (2011).

N. Kadakia, M. Huang and H. Bakhru. Fabrication of subsurface metallic nanoparticles for enhanced carrier generation in silicon-based photovoltaics. *MRS Conference Proceedings*, **1322**, DOI: 10.1557 /opl.2011.1104 (2011).

N. Kadakia, S. Naczas, H. Bakhru, and M. Huang. Method for silicon surface texturing using ion implantation. *AIP Conference Proceedings*, **1336**: 337–340 (2011).

N. Kadakia, S. Naczas, H. Bakhru, and M. Huang. Fabrication of surface textures by ion implantation for antireflection of silicon crystals. *Applied Physics Letters*, **97**(19): 191912 (2010).

ABSTRACT OF THE DISSERTATION

The Dynamics of Nonlinear Inference

by

Nirag Kadakia

Doctor of Philosophy in Physics

University of California, San Diego, 2017

Professor Henry D. I. Abarbanel, Chair

The determination of the hidden states of coupled nonlinear systems is frustrated by the presence of high-dimensionality, chaos, and sparse observability. This problem resides naturally in a Bayesian context: an underlying physical process produces a data stream, which – though noisy and incomplete – can in principle be inverted to express the likelihood of the underlying process itself. A large class of well-developed methods treat this problem in a sequential "predict-and-correct" manner that alternates information from the presumed dynamical model with information from the data.

One might instead formulate this problem in a temporally global, non-sequential manner, which suggests new avenues of approach within an optimization context, but also poses new challenges in numerical implementation. The *variational annealing* (VA) technique is proposed to address these problems by leveraging an inherent separability between the convex and nonconvex contributions of the resulting functional forms. VA is shown to reliably track unobservable states in sparsely observed chaotic systems, as well as in minimally-observed biophysical neural models.

Second, this problem can be formally cast in continuous time as a Wiener path integral, which then suggests classical solutions derived from Hamilton's principle. These solutions come with their own difficulties in that they comprise an unstable boundary-value problem. Accordingly, a further technique called *Hamiltonian variational annealing* is proposed, which again exploits an existing separability of convexity and nonlinearity, this time in a an enlarged manifold constrained by underlying symmetries.

A running theme in this thesis is that the optimal estimate of a nonlinear system is itself a dynamical system that lives in an unstable, symplectic manifold. When this system is recast in a variational context, instability is manifested as nonconvexity, the central idea being that when this nonconvexity is incorporated in a systematic and gradual way, the classical solutions can be tracked reliably.

Chapter 1

Introduction

1.1 Invitation: the simple oscillator and a naive procedure

Imagine observing a physical system whose dynamics are approximately dictated by Hooke's Law, so the position x and velocity v of the system evolve according to the system of ordinary differential equations:

$$\begin{aligned} \dot{x} &= -\omega^2 v \\ \dot{v} &= x. \end{aligned} \tag{1.1}$$

Suppose we want to precisely determine the character of the dynamics – the oscillation frequency ω – by observing this system in motion, but the only devices available to us are a ruler and a stopwatch. We could use these to make measurements $Y = \{y_i\}$ of x at equally spaced times $\{t_i\}$, as in Fig. 1.1a. Being either supremely optimistic or supremely naive, we now put our full faith in the accuracy of our measurements, whereby

$$y_i \equiv x(t_i). \tag{1.2}$$

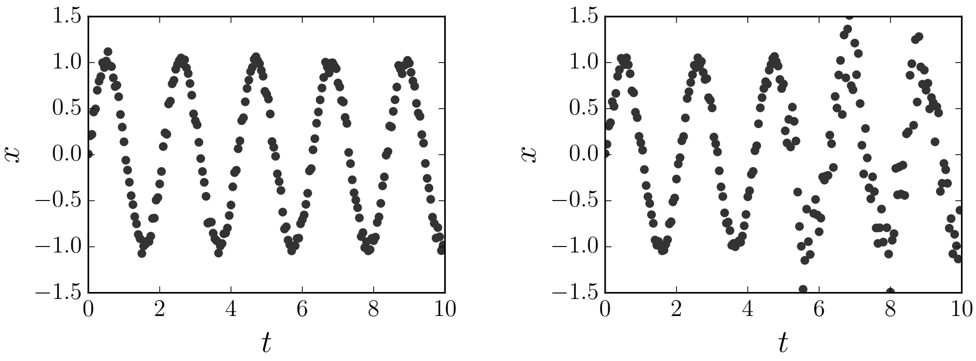

(a) (b)

Figure 1.1: (a) "Observations" of the position of a weight attached at the end of a bobbing spring. (b) The same system, interrupted halfway by an earthquake.

Using Eq. 1.2, we can then incorporate our measurements into our model by substituting the data into a discretization of $\dot{x}$,

$$-\omega^2 v_i = \dot{x} \approx \frac{x(t_{i+1}) - x(t_i)}{\Delta t} \rightarrow \frac{y_{i+1} - y_i}{\Delta t} \tag{1.3}$$

and then inserting this into a discretization of $\dot{v}$, thereby expressing ω only in terms of the measurements:

$$\frac{y_{i+2} - 2y_{i+1} + y_i}{\Delta t^2} = -\omega^2 \frac{v_{i+1} - v_i}{\Delta t} = -\omega^2 x_i \sim -\omega^2 y_i. \tag{1.4}$$

Of course, this equation may give different values of ω, ω_i, for each t_i. So we further assume that a reasonable estimate of ω can be found by averaging over ω_i:

$$\begin{aligned} \omega_i^2 &\sim \frac{-y_{i+2} + 2y_{i+1} - y_i}{y_i \Delta t^2} \\ \langle \omega \rangle &= \frac{1}{N-2} \sum_i \omega_i. \end{aligned} \tag{1.5}$$

Eq. 1.5 gives a prescription for *estimating the parameter of a dynamical system, in closed form, using a batch of measured data*.

1.2 The ubiquity of inverse problems in scientific inquiry

This crude procedure, peppered with some reasonable assumptions and many extraordinary ones, is essentially the problem of interest in this thesis: How does one effectively incorporate measured data to accurately and reliably infer unknown variables of a dynamical model? Think of this as an *inversion* of the underlying physical process. That is, with a dynamical model and initial conditions in hand, we could i) generate a series of observations by forward integration to generate a "true" trajectory $\mathbf{X}$, ii) corrupt $\mathbf{X}$ with noise to generate "data" $\mathbf{Y}$, and finally iii) prune $\mathbf{Y}$ to only the measurements $\mathbf{Y}'$ actually accessible with existing experimental tools. The inverse problem runs this process in reverse, determining the original $\mathbf{X}$ from the observed but sparse $\mathbf{Y}'$.

The generality of such dynamical inverse problems cannot be understated. Mathematical descriptions underlying physical, biological, and chemical processes are proposed on the basis of physical law – these interactions are modeled by differential or algebraic equations that fully but parsimoniously describe how an input in one section of the system generates an output elsewhere. Yet these functional forms are parameterized, sometimes quite heavily, by constants independent of the nature of these forces, and must be either chosen heuristically or determined from experiment. Further, only a subset of the dynamical state space can usually be observed at all. Finally, when the dynamics are nonlinear, the situation is even more delicate in that entirely distinct qualitative behavior can be produced by differing choices of these constants alone – this is the theory of dynamical bifurcations. The hope of the inference problem in nonlinear dynamical systems is that, despite these complications, information from observations may be systematically combined with the dynamical model to infer all this unknowns in an accurate and reliable manner. As an illustration of the ubiquity of this problem, consider the following examples:

- The Hodgkin-Huxley (HH) model of a neural cell exhibits well-defined voltage spikes in response to injected current stimuli, mimicking the voltage response of actual biological neurons [Hodgkin and Huxley, 1952]. The ease with which spikes are elicited, along with spiking characteristics such as waveform and frequency, depend sensitively on the model parameters [Dayan and Abbott, 2005, Ermentrout and Terman, 2010]. Of the four variables in the dynamical state space, only the voltage is measurable directly with electrophysiology, so the remaining state variables, along with the static parameters, must be inferred from voltage traces alone. Finally, in actual neural systems, the complexity of this scenario is magnified by at least a billion.

- Earth's atmosphere is a remarkably complex fluid which can nevertheless be modeled under some reasonable assumptions by a set of six nonlinear differential equations, the atmospheric primitive equations, whose state space of dynamical variables consists of wind velocity, pressure, temperature, and moisture content at each point around the globe [Bennett, 1992, Kalnay, 2002]. Obviously this is an enormous number of variables of which only a tiny fraction are accessible at any time. This small set of measurements, taken from sparsely located weather centers, must somehow be incorporated into the highly nonlinear primitive equations to generate reliable daily weather forecasts.

Though this thesis will focus on inference in *dynamical* models – and those that are nonlinear and chaotic in particular – inference in some form or another drives the work of every scientist. After all, even the straightforward task of fitting data to a straight line is essentially an inverse problem.

1.3 Characteristics and challenges

The linear oscillator example in the first section of this thesis proposed a simple procedure to estimate ω by incorporating the measured data into equations of motion. Since the dynamics are linear, the problem is far simpler (and the estimation far more naive) than in practical scenarios of interest, such as those encountered in neuroscience and meteorology. Still, the example illustrates some universal features common to inference in dynamical systems, both linear and nonlinear:

- *A presumed dynamical model* How this description will be chosen is the result of other physical considerations, and consists of some combination of ordinary and partial differential equations or discrete maps specifying the evolution of the system in time. It is a result of both physical considerations and reflects some compromise between reality, versatility, and tractability.

- *Regularization* Even though data exists for only a few variables, each variable may still (in theory) be unambiguously estimated, since they are all coupled together through the dynamical equations. The equations may therefore be thought of as a sort of regulator, precluding the sparsely observed system from being underdetermined. Regulators play a central role in the theory of inverse problems by permitting the computation of parameters from otherwise ill-posed problems [Vogel, 2002, Hansen, 2010]. As will be implied later in this thesis, the reliability and accuracy of the estimation relies critically on the choice and nature of the regulator, whose selection is much an art as a science.

- *A discrete formulation* Observations are made in discrete time, so continuous dynamics must be approximated. The choice of this discretization is problem-dependent, driven by considerations such as desired accuracy and computational cost. One of the themes of contributions of this thesis will be a new formulation of the dynamical inference problem in a symplectic manifold, where we will see that this choice will have nontrivial consequences.

On the other hand, the linear model and the estimation procedure presented in the opening section of this thesis contained several amenable features that are all but absent in complex, realistic scenarios of interest. Among them are:

- *Model linearity* As the dynamics in Eq. 1.1 are linear, v was straightforwardly removed in favor of x. Even with a less naive estimation procedure, such as least squares, the variables could still be estimated in closed form with linear algebra techniques [Björck, 1996, Hansen, 2010, Särkkä, 2013]. For any semi-interesting model, the dynamics are nonlinear, and this inversion cannot be done in closed form.

- *Linear parameterization* The static parameter ω also entered the equation linearly, and so was expressible explicitly in terms of the data. In most cases, parameters do not admit closed form expressions, and numerically calculating them may be computationally prohibitive.

- *Dense observability* The system is two-dimensional, half of the dynamical states are observed, and there is only a single parameter to be estimated. The anticipated difficulties in scaling are obvious: even if the estimate could be algebraically determined, matrix inversion scales at least quadratically with dimension [Cormen et al., 2001]. Further, nonlinearities in either parameters or variables would preclude a closed-form expression and require iterative techniques such as the Newton-Raphson method [Nocedal and Wright, 2006]. Finally, for non-convex cost functions (which we will see are fundamental to nonlinear systems), the solution to this inversion is far from unique.

- *Model specificity* The estimated parameters may even live in a degenerate manifold, that is, several distinct data sets produce estimates of near-equal accuracy. This is often the case when the data is sufficiently sparse and noisy. To take an extreme case, an intracellular voltage trace of a single neuron in a 10^6-neuron system would be entirely inadequate in determining the $\mathscr{O}(10^{12})$ connection strengths of the network. Any inference procedure would return several estimates of near-equal validity for a particular dataset.

- *Perfect models and clean data* Errors in both models and measurements must be taken into account. Suppose, for example, that a light earthquake disrupted the latter half of the experiment, producing a measured trajectory as in Fig. 1.1b. Any reasonable procedure would take this variance into account by weighting the estimate more heavily toward the earlier data.

While all of these considerations will enter the scope of this thesis, we will be concerned mainly with the ramifications of the first three – the presence of chaotic behavior, of sparse observability and of nonlinear parameterization. This emphasis is driven by the focus of this thesis on the inference problem in neuroscience in particular, with the intention of developing estimation techniques that respect, as much as possible, the known biophysical features of actual neurons that lead to action potential propagation – ion channels, dynamical synapses, multiple timescales. Admittedly, this is somewhat in contrast to the present emphasis of computational neuroscience research, which favors increased complexity and dimensionality in what are essentially *static* input-output machines [Bishop, 1995]. This emphasis is not without good reason, driven by the recent success of deep neural networks in machine learning and artificial intelligence [LeCun et al., 2015]. The central assumption of this thesis is that these state machines, though versatile and powerful, are fundamentally limited in their neglect of the complex nonlinear interactions at the cellular and synaptic levels. The hope is that the methods presented here will lay a foundation for transcending these limitations when they are one day realized.

1.4 The scope of this thesis

The deepest roots of this thesis lie in Rudolf Kalman's analytical solution to the linear tracking problem more than a half century ago [Kalman, 1960, Särkkä, 2013]. The algorithm, now known eponymously as the *Kalman filter*, prescribes in closed form the optimal estimate of a state vector subject to normally-distributed noisy observations and discrete-time linear dynamics. Chapter 2 will begin with a discussion of the Kalman filter, its extensions, and its limitations for the problems of interest in this thesis. The primary focus of this chapter will be the breakdown of the Gaussian-linear assumption when dealing with sufficiently nonlinear models, and the ramifications therein. Also discussed is the precarious status of parameter estimation in the filtering context: indeed, one of the key differences between object tracking and meteorology – for which the Kalman filter is the primary workhorse – and neural systems is that while unknown parameters are rare in the former they are pervasive in the latter [Auger et al., 2013, Houtekamer and Zhang, 2016].

Chapter 3 will then propose an algorithm that dispenses entirely with the iterative framework of filters, instead casting the estimation as a variational approximation to a high-dimensional Bayesian integral [Ye et al., 2015a, Ye et al., 2015b]. The variational approximation brings its own complications, and a systematic method to mitigate these difficulties, predicated on the particular mathematical structure of these integrals, will be presented. The benefits of this *variational annealing* technique will then be illustrated via numerical experiments in sparsely measured chaotic systems. This calculation will set the stage for a more involved application to detailed neuron models in Chapter 4.

Chapter 4 is a lengthy study of the estimation problem applied to the vocal pathway in zebra finch songbirds. Male zebra finch express a single, short, stereotyped mating call throughout their adult life, and the neural circuit involved in learning, repeating, and maintaining this song motif has been extensively studied [Kubota and Taniguchi, 1998, Brainard and Doupe, 2002, Fee et al., 2004, Mooney and Prather, 2005]. This circuit is characterized by some peculiar features. It has been found, for example, that neurons projecting from the vocal command center to motor neurons in the larynx – those purportedly initiating the vocalization – burst ultra-sparsely and ultra-reliably during the song motif, while inhibitory neurons fire in patterns far more erratic and dense [Hahnloser et al., 2002, Kosche et al., 2015]. A neural model for the projection neurons incorporating known experimental features and responses will be presented and analyzed via dynamical bifurcation theory. Variational annealing will be applied to estimate the unknown parameters, and we will find interesting complicating features of the estimation procedure that stem from i) the presence of multiple timescales as well as ii) model degeneracies. A timescale splitting protocol will be proposed in response to these complications.

Chapter 5 begins the second half of the thesis, which illuminates some novel mathematical features of the estimation problem. Moving formally to continuous time, the Bayesian integrals encountered in Chapter 2 can now be viewed as some combination of a continuous stochastic process and an incoming observation process. The formal equivalence between stochastic differential equations and Wiener path integrals is highly suggestive [Chaichian and Demichev, 2001, Chow and Buice, 2015]; in our case, it leads to a fully continuous path integral formulation of the estimation problem, with associated Lagrangian function and action functional. The variational formulation presaged in Chapters 2 and 3 has now opened up a wealth of analogies with classical and quantum mechanics, including the Euler-Lagrange and Hamiltonian descriptions [Goldstein et al., 2002]. The notion that the *estimate* of the state is itself a well-defined dynamical system is what gives this thesis its name. What's more, we will see that the Lagrangian function that describes this system has a direct correspondence to classical electromagnetism. This connection, still nascent, will be presented mainly as a divergence point for future work in this field.

The mechanical analogy appears to imply that the estimate may be solvable by simply integrating dynamics numerically in time to generate an optimal estimate. But we will find that while the traditional mechanics problem arising from Hamilton's principle is an initial-value formulation – well-defined and unique, if not fully integrable – the estimation dynamics is a boundary value problem since the initial state, being sparsely observed, is not fully characterized. A key motif in Chapter 5 is that this boundary value problem is intractable since the dynamics of the estimation are persistently unstable, so a reliable estimation procedure would require a strong initial guess that is rarely possible.

This notion, reminiscent of similar issues in optimal control theory, will set the stage for a further look at the estimation dynamics in a Hamiltonian formulation in Chapter 6. Integral curves of Hamilton's equations preserve symplectic invariants, and it is proposed that by casting the estimation dynamics into canonical coordinates, these symmetries can be leveraged to stabilize the estimate. The notion of symplecticity, its destruction in discrete time, and reinstatement by particular numerical integration schemes, as well as the robust formulation of symplecticity in *discrete* Lagrangians, will be reviewed in this chapter. Applying this idea to the estimation problem, the novel contribution of this chapter is an inference algorithm akin to variational annealing, but formulated entirely in canonical coordinates [Kadakia et al., 2017].

The conclusion will be brief, and mostly open-ended. In it will be proposed an (as yet unexplored) extension of the annealing methods to chaotic system synchronization. Finally, we conclude with an optimistic but sober reminder of the gap between the breadth of these techniques and the challenges presented by the supremely complex neural systems actually found in nature.

Chapter 2

Inference in nonlinear dynamical systems

2.1 Bayesian recursion as a starting point

Though not stated outright, implicit in the oscillator example in Sec. 1.1 is that the estimation problem is fundamentally probabilistic: the data is noisy, the model is inherently approximate. Further, since the measurements are sequential, arriving at discrete points in time, it is natural to treat the problem in a Bayesian setting, where an initial prior distribution is iteratively updated with incoming data and model tendencies to produce updated posterior distributions [Särkkä, 2013]. In the context of neuroscience and other biological models, the ultimate goal is the determination of unknown static parameters of the system, yet this will rely on estimating, simultaneously or in alternating fashion, the time course of the dynamical variables – in particular, those variables which cannot be observed. So as a first step, we seek expressions for the distributions of dynamical states, assuming the parameters are known.

Before delving into the Bayesian formulation of the estimation problem, note that the forthcoming results can be derived equivalently without explicitly drawing on Bayesian inference. In Kalman's original paper, the optimal estimate and accompanying optimal covariance were found using a correspondence between random variable correlations ("projections") and optimality [Kalman, 1960]. The derivation in this manner is elegant and clean. Working with the full distributions themselves, however, is a bit more transparent, and importantly will smooth the transition to non-Gaussian distributions which are central to this thesis.

2.1.1 The Bayesian recursion equations

The estimation problem is set in discrete time, where the state of the D-dimensional system at time t_n is denoted by a vector $\mathbf{x}_n \in \mathbb{R}^D$. Working for now in complete generality, we assume only that the state evolves as a Markov process [Gardiner, 1985]. Simultaneously, measurements $\mathbf{y}_n$ are generated by passing $\mathbf{x}_n$ through some *observation operator* $\mathbf{H}_n$:

$$\mathbf{y}_n = \mathbf{H}_n(\mathbf{x}_n) + \mathbf{r}_n \tag{2.1}$$

The parenthesis indicate that $\mathbf{H}_n$ may be a nonlinear function and its subscript that it may not be stationary, while the addition of random variable $\mathbf{r}_n$ reflects the reality that measurements are never pure. For all of the examples in this thesis, $\mathbf{H}_n$ is essentially a constant linear projection and $\mathbf{r}_n$ is a homogeneous Gaussian process. For example, the Hodgkin-Huxley model dynamics are described by three unmeasurable gating variables and one observed voltage variable: $\mathbf{H_n} \equiv \mathbf{H}$ is a projection matrix of unity rank. Similarly, current clamp electrophysiology is clean up to the addition of some Gaussian white noise, whereby $\mathbf{r}_n \equiv \mathbf{r} \sim \mathcal{N}(0, \sigma^2)$.

The presence of the random process $\mathbf{r}_n$ linking states and observation naturally suggests a probabilistic dependence of the measurements upon the true state, for example, $p(\mathbf{y}_n|\mathbf{x}_n) \sim \mathcal{N}(\mathbf{H}_n(\mathbf{x}_n), \mathrm{Var}[\mathbf{r}_n])$. We need to invert this expression for a distribution of $\mathbf{x}_n$, having measured $\mathbf{y}_n$: $p(\mathbf{x}_n|\mathbf{y}_n)$. These distributions are of course linked by Bayes' Rule.

Actually, a more useful quantity is the distribution of $\mathbf{x}_n$ given not just the present observation, but all those prior:

$$p(\mathbf{x}_n|\mathbf{Y}_n), \tag{2.2}$$

where upper-case $\mathbf{Y}_n$ is the set of all observations; $\mathbf{Y}_n = \{\mathbf{y}_0, ..., \mathbf{y}_n\}$. This expression can be broken up into a product of individual factors by applying the Markov property and Bayes' rule in alternating fashion. The Markov assumption allows this expression to be decoupled into factors representing information gain timepoint by timepoint. For Markovian dynamics, we can always use the the Chapman-Kolmogorov equation [Gardiner, 1985],

$$p(f_3|f_1) = \int df_2 p(f_3|f_2)p(f_2|f_1), \tag{2.3}$$

the content of which is that the probability of transitions from state f_1 to f_3 is the sum of all possible transitions also passing through an intermediate state f_2. Applying both this and Bayes' Rule to the final dynamical state $\mathbf{x}_n$ in Eq. 2.2 gives:

$$\begin{aligned} p(\mathbf{x}_n|\mathbf{Y}_n) = p(\mathbf{x}_n|\mathbf{y}_n, \mathbf{Y}_{n-1}) &\propto p(\mathbf{y}_n|\mathbf{x}_n, \mathbf{Y}_{n-1})p(\mathbf{x}_n|\mathbf{Y}_{n-1}) \\ &= p(\mathbf{y}_n|\mathbf{x}_n)p(\mathbf{x}_n|\mathbf{Y}_{n-1}) \\ &= p(\mathbf{y}_n|\mathbf{x}_n)\int dx_{n-1} p(\mathbf{x}_n|\mathbf{x}_{n-1})p(\mathbf{x}_{n-1}|\mathbf{Y}_{n-1}). \end{aligned} \tag{2.4}$$

In the second line, the disappearance of the measurement history in the first factor is due to the presumed conditional independence of $\mathbf{y}_n$ upon previous measurements $\mathbf{Y}_{n-1}$, given $\mathbf{x}_n$, as a consequence of Eq. 2.1. The final factor in the last line of Eq. 2.4 is identical to the analogous distribution, just one timestep earlier. This suggests an obvious iterative recursion. Repeating the steps in Eq. 2.4 from the initial to final time therefore gives:

$$p(\mathbf{x}_N|\mathbf{Y}_N) \propto p(\mathbf{y}_N|\mathbf{x}_N)\left[\int \prod_{n=0}^{N-1} d\mathbf{x}_n p(\mathbf{x}_{n+1}|\mathbf{x}_n)p(\mathbf{y}_n|\mathbf{x}_n)\right] p(\mathbf{x}_0) \tag{2.5}$$

In general, these integrals do not admit closed form expressions, but they may be estimated with Markov Chain Monte Carlo sampling or of its many variants [Metropolis et al., 1953, Duane et al., 1987, Neal, 1996, Gelman et al., 2013]. Though Monte Carlo approaches are powerful, useful, and widely studied, they will for the most part not be considered in this thesis; rather the focus will be on variational approximations and their relation to sequential filtering. Let us therefore begin with the various manifestations of the latter.

2.2 Kalman-based filtering and smoothing

2.2.1 Kalman filter: linear dynamics and Gaussian errors

Since products of Gaussians are also Gaussian, then Eq. 2.5 can be evalauted iteratively in closed form when the transition and measurement probabilities, $p(\mathbf{x}_n|\mathbf{x}_{n-1})$ and $p(\mathbf{y}_n|\mathbf{x}_n)$, are normally distributed. This is the basis of the Kalman filter, which expresses exact recursion relations for the moments of these normal distributions at each point in time [Kalman, 1960, Grewal and Andrews, 2014]. The Kalman filter relies on the observation that a Gaussian transition probability arises when a linear map generating $\mathbf{x}_{n+1}$ from $\mathbf{x}_n$ is combined with perturbing "process" noise that is normally distributed. Specifically, we can write

$$\begin{aligned} \mathbf{x}_{n+1} &= \mathbf{F}_n\mathbf{x}_n + \mathbf{q}_n \\ \mathbf{q}_n &\sim \mathscr{N}(0, \mathbf{Q}_n) \\ \rightarrow p(\mathbf{x}_{n+1}|\mathbf{x}_n) &\sim \mathscr{N}(\mathbf{F}_n\mathbf{x}_n, \mathbf{Q}_n). \end{aligned} \tag{2.6}$$

The measurements follow Eq. 2.1 with linear observation operators and Gaussian corruption:

$$\begin{aligned} \mathbf{y}_n &= \mathbf{H}_n\mathbf{x}_n + \mathbf{r}_n \\ \mathbf{r}_n &\sim \mathscr{N}(0, \mathbf{R}_n) \\ \rightarrow p(\mathbf{y}_n|\mathbf{x}_n) &\sim \mathscr{N}(\mathbf{H}_n\mathbf{x}_n, \mathbf{R}_n). \end{aligned} \tag{2.7}$$

When these assumptions – linear dynamics and Gaussian noise – are used in iterative fashion from a Gaussian prior, $p(\mathbf{x}_0) \sim \mathscr{N}(\mathbf{m}_0, \mathbf{P}_0)$, the integrals in Eq. 2.4 are exact. After some lengthy but straightforward integration and algebraic manipulations, what results are the Kalman filtering distributions at each time step, which are most transparently written in two distinct steps reflecting the alternating applications of the state evolution by the Markov dynamics (prediction) and incorporation of new observations by Bayes' Rule (updating). These distributions are:

$$\begin{aligned} p(\mathbf{x}_n|\mathbf{Y}_{n-1}) &\sim \mathscr{N}(\mathbf{m}_n^-, \mathbf{P}_n^-) \\ p(\mathbf{x}_n|\mathbf{Y}_n) &\sim \mathscr{N}(\mathbf{m}_n, \mathbf{P}_n), \end{aligned} \tag{2.8}$$

where the moments of the prediction step are

$$\begin{aligned}\mathbf{m}_n^- &= \mathbf{F}_{n-1}\mathbf{m}_{n-1}\\ \mathbf{P}_n^- &= \mathbf{F}_{n-1}\mathbf{P}_{n-1}\mathbf{F}_{n-1}^{\mathrm{T}} + \mathbf{Q}_{n-1},\end{aligned} \tag{2.9}$$

and the moments of the updating step are

$$\begin{aligned}\mathbf{m}_n &= \mathbf{m}_n^- + \mathbf{K}_n(\mathbf{y}_n - \mathbf{H}_n\mathbf{m}_n^-)\\ \mathbf{P}_n &= \mathbf{P}_n^- - \mathbf{K}_n(\mathbf{H}_n\mathbf{P}_n^-),\end{aligned} \tag{2.10}$$

where $\mathbf{K}_n = \mathbf{P}_n^-\mathbf{H}_n^{\mathrm{T}}(\mathbf{H}_n\mathbf{P}_n^-\mathbf{H}_n^{\mathrm{T}} + \mathbf{R}_n)^{-1}$ is the gain factor quantifying the strength with which the data corrects the anticipated evolution of the dynamical step alone. By applying these equations to the *a priori* assumed values of $\mathbf{m}_0$ and $\mathbf{P}_0$, the updated moments $\mathbf{m}_1, \mathbf{P}_1$, can be calculated in closed form, and so on. Since the filtering distributions are Gaussian, the *optimal* estimate after all the observations are incorporated is simply the final mean value $\mathbf{m}_N$, with accompanying variance $\mathbf{P}_N$.

The expressions for the optimal estimate and covariance are explicit. Any anticipated difficulties will arise entirely from issues of ill-conditioning and cost in performing the multiplication, inversions, and storage required for $\mathbf{K}_n$. The theory and practical methods to perform these linear algebraic manipulations effectively has been developed in quite some detail, and can be found for example in [Grewal and Andrews, 2014].

2.2.2 Nonlinear extensions: extended and unscented Kalman filters

The Kalman filter equations are of course predicated on the fairly limiting assumptions of linearity and Gaussian statistics. There are a few ways to relax this assumption and still produce quality estimates. In the case of nonlinear dynamics, $\mathbf{x}_{n+1} = \mathbf{f}(\mathbf{x}_n) + \mathbf{q}_n$, we can opt to linearize around the current mean [1] :

$$\mathbf{F}_n = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_{\mathbf{u}_n}, \tag{2.11}$$

---

[1] Here and throughout, lower-case **f** will be used to denote nonlinear functions, while upper-case **F** will denote matrix operators.

and use this approximated evolution operator in Eqs. 2.9. This is the extended Kalman Filter (EKF), and is a first-order approximation of the Kalman filter for nonlinear systems. EKF works well for weak nonlinearities and small timesteps, but can quickly become inaccurate – and possibly even diverge – for many realistic nonlinear models of interest [Julier and Uhlmann, 1997, Julier et al., 2000, Särkkä, 2013].

Nonlinear evolutions operators destroy the Gaussian nature of the probability distributions, even if the process noise $\mathbf{q}_n$ is itself still normally distributed – the argument of the exponential in $p(\mathbf{x}_{n+1}|\mathbf{x}_n)$ is $\sim \mathbf{f}(\mathbf{x}_n)^{\mathrm{T}}\mathbf{Q}_n^{-1}\mathbf{f}(\mathbf{x}_n)$, which is no longer quadratic. EKF solves this problem by linearizing the dynamics: $\mathbf{f} \rightarrow \mathbf{F}$. But one could adopt the alternate viewpoint – retain the exact dynamics but break the assumption of Gaussianity. That is, we dispense with the probability distributions altogether, instead *representing* the distribution by a set of discrete D-dimensional points whose statistics are chosen to match the first and second moments at t_n. These points can then be propagated through the full nonlinear dynamics, with $\mathbf{u}_{n+1}$ and $\mathbf{P}_{n+1}$ then recomputed by sampling. This is the essence of the *unscented Kalman filter*, which represents the distribution by a prescribed set of $2D+1$ *sigma points*, $\mathscr{X}_n^m$, which are dilated, shifted copies of the $2D+1$ so-called *canonical* sigma points. $\mathscr{X}_n^m$ are individually evolved through the nonlinear model $\mathbf{f}$, the result being used to calculate the sample mean and covariance $\mathbf{m}_n^-$ and $\mathbf{P}_n^-$. In the Kalman filter, these moments are given analytically, whereas here they are statistical.

These sampled moments are used in Eqs. 2.10 to produce the updated moments $\mathbf{m}_n$ and $\mathbf{P}_n$, and the cycle is restarted by again producing new sigma points by dilating and shifting the canonical sigma points to reflect the new statistics given by $\mathbf{m}_n$ and $\mathbf{P}_n$. Thus, in UKF, the Gaussian assumption is relaxed in the update step (allowing full model nonlinearity), but reinstated in the analysis step since the new sigma points are prescribed only by the first two moments. Since the dynamics are not linearized, Jacobian calculations are avoided ;in fact, the dynamical map need not even be differentiable. The UKF is is third-order accurate compared to the first-order EKF, remarkably for similar computational cost [Julier and Uhlmann, 1997, Julier et al., 2000, Wan and van der Merwe, 2002].

2.2.3 Smoothing: backwards passes

The Kalman filtering equations produce an estimate of $\mathbf{x}_n$ given prior observations. This is quite natural in, for example, meteorology: estimates are updated in real-time on a 24-hour cycle, given newly gathered daily weather patterns. In estimating neural systems, the situations is less restricted. Results are analyzed after an entire batch of experimental data is taken, and the time for analysis and estimation is not constrained by a morning forecast. It is reasonable to ask about the neural activity early in the experiment, at time t_n, having knowledge not only of observations before t_n, but those throughout the entire observation window $[t_0, t_N]$,

$$p(\mathbf{x}_n|\mathbf{Y}_n) \to p(\mathbf{x}_n|\mathbf{Y}_N). \tag{2.12}$$

Though the optimal estimates produced by these two distributions coincide at the end of the observation window, we would expect that for the optimal estimates in the latter distribution are more accurate for earlier $\mathbf{x}_n$, having been informed by more datapoints. One method for producing the distributions $p(\mathbf{x}_n|\mathbf{Y}_N)$ essentially requires running the Kalman filter twice – once forward, once backwards. As in Eq. 2.4, the derivation requires only Bayes' Rule and the Markov property, resulting in :

$$p(\mathbf{x}_n|\mathbf{Y}_N) = p(\mathbf{x}_n|\mathbf{Y}_n) \int \left[\frac{p(\mathbf{x}_{n+1}|\mathbf{x}_n) p(\mathbf{x}_{n+1}|\mathbf{Y}_N)}{\int p(\mathbf{x}_{n+1}|\mathbf{x}_n) p(\mathbf{x}_n|\mathbf{Y}_n) d\mathbf{x}_n} \right] d\mathbf{x}_{n+1}. \tag{2.13}$$

This "backward" recursion relation generates $p(\mathbf{x}_n|\mathbf{Y}_N)$ by combining $p(\mathbf{x}_{n+1}|\mathbf{Y}_N)$ with the output of the forward Kalman filter, $p(\mathbf{x}_n|\mathbf{Y}_n)$. When the dynamics are normal and the noise is Gaussian, we expect closed form recursion relations for the optimal estimates and covariances, analogous to Eqs. 2.8-2.10. These closed form expressions comprise what is now known as the Rauch-Tung-Striebel smoother, named after its developers [Rauch et al., 1965].

2.2.4 Ensemble Kalman filter: A partially statistical approach

The Kalman filter and its nonlinear extensions require the manipulation and storage of large covariance matrices, costly matrix multiplication and inversions, and limiting assumptions on allowable nonlinearities in the dynamical evolution. We saw that in UKF we could alleviate the issue of nonlinearities by representing the distribution by set of discrete points and propagating these points through the exact dynamics $\mathbf{f}$. In UKF, $2D+1$ distinct sigma points are needed, so scaling may become prohibitive for the high-dimensional models prevalent in geophysics and neuroscience.

Can we adequately represent the distribution by fewer discrete samples? This idea is the basis of the *ensemble Kalman filter* (EnKF), which represents the distribution at each timepoint t_n by an ensemble of M discrete, *sampled* D-dimensional points $\{\mathbf{x}_n^1, ..., \mathbf{x}_n^M\}$ [Evensen, 1994, Evensen, 2003, Evensen, 2009]. The idea is similar to UKF, except that M is arbitrary, and is in practice often smaller than the state space dimension D. The EnKF proceeds as follows. In the prediction step, each sample is passed through the nonlinear dynamics to produce the updated samples $\tilde{\mathbf{x}}_n^m$, i.e. $\tilde{\mathbf{x}}_n^m = \mathbf{f}(\mathbf{x}_n^m) + \mathbf{q}_n$ (recall that $\mathbf{q}_n$ is an instantiation of the process noise random variable, Eq. 2.7). Now, rather than calculating the covariance matrix analytically as in Eq. 2.9, the moments are calculated statistically from the updated ensemble members, as in UKF:

$$\begin{aligned} \mathbf{m}_n^- &= \langle \tilde{\mathbf{x}}_n^m \rangle = \frac{1}{M} \sum_m \tilde{\mathbf{x}}_n^m \\ \mathbf{P}_n^- &= \langle (\tilde{\mathbf{x}}_n^m - \mathbf{m}_n^-)^2 \rangle = \frac{1}{M-1} \sum_m (\tilde{\mathbf{x}}_n^m - \mathbf{m}_n^-)^{\mathrm{T}} (\tilde{\mathbf{x}}_n^m - \mathbf{m}_n^-), \end{aligned} \tag{2.14}$$

where $1/(M-1)$ is Bessel's correction on the sample covariance, since the mean is unknown [Johnson and Wichern, 2007]. With these moments in hand, the update equations are equivalent to Eqs. 2.10 with one difference: the observations cannot be used directly, but must themselves be sampled from a distribution with mean $\mathbf{y}_n$ and variance $\mathbf{R}_n$, otherwise the resulting updated covariances $\mathbf{P}_n$ are underestimated [Evensen, 2003].

The EnKF samples are initialized from some prior distribution, or background, $\mathbf{x}_0^m \sim p_{\mathrm{b}}(\mathbf{x})$, which is chosen to reflect some prior knowledge of the system state before observations have been incorporated. It is then run forward by iteratively applying, to each sample $\mathbf{x}_n^m$, the predict and update steps just described. When D is large, even a moderate M may produce adequate statistics and remain stable. See [Mitchell et al., 2002, Natvik and Evensen, 2003] for applications to atmospheric and ocean models in particular; an overview of applications and extensions can be found in [Evensen, 2003].

2.3 Sequential Monte Carlo

In implementing the filtering equations Eq. 2.9 and 2.10, both UKF and EnKF dipped a cautious toe in the waters of sampling methods. Recall that the initial impetus of Gaussian-linear assumption was that it could allow the computation of the Bayesian conditional distributions Eq. 2.5 in closed form. While UKF and EnKF both relax the Gaussian assumption during the prediction step – at the cost of an exact expression for the moments – they effectively reinstate it during the analysis step by updating only the first two moments of the sample distribution via the optimal Kalman gain correction $\mathbf{K}$. In UKF, the intent was that a non-Gaussian

$p(\mathbf{x}_{n+1}|\mathbf{x}_n)$ might be represented more accurately by a collection of samples, while in EnKF, the focus was on reducing the cost of matrix operations in high D by representing the full distributions by a smaller collection of stochastic points.

We might instead dive headlong into a sampling approach, dispensing entirely with any of Kalman's closed form expressions for $\mathbf{u}_n$ and $\mathbf{P}_n$. Could this be done effectively, in a way that mitigates the curse of dimensionality? Certainly, a naive application of the Metropolis-Hastings Monte Carlo method to the full conditional distribution Eq. 2.5 might be slow and inefficient for appreciable D [Metropolis et al., 1953, Gilks et al., 1995, Gelman et al., 2013]. It might be more advantageous to continue working in the iterative context suggested by the Kalman filter. This is the approach taken by one of the many variants of the *particle filter*, which essentially performs the recursion in Eq. 2.4 by sampling from a surrogate distribution iteratively in time, weighted appropriately by the measurement and model functions, $p(\mathbf{y}_n|\mathbf{x}_n)$ and $p(\mathbf{x}_n|\mathbf{x}_{n-1})$ [Gordon et al., 1993, Doucet et al., 2001, Arulampalam et al., 2002, Doucet and Johansen, 2009]

Though particle filters arise in many forms, including genetic-type algorithms, we will focus here on one widespread variant amenable to Gaussian observation and process noise, where neither the observation operator nor the the dynamics are required to be linear. This approach is known in the literature as *sequential Monte Carlo* (SMC), though may sometimes also be referred to generically as a particle filter or sequential importance sampling (SIS) [Doucet and Johansen, 2009, Doucet et al., 2001].

To begin, SMC approximates the conditional distributions at all steps by a collection of M samples, $\mathbf{x}_n^m$, as in EnKF. To each sample at time t_n is also associated a *weight* w_n^m, with $\sum_m w_n^m = 1$. The goal of SMC is to update both these samples and their weights iteratively, but since samples cannot always be easily drawn from the integrand of Eq. 2.4, the idea is to instead draw them from a feasible *proposal distribution* with compensating weights.

Specifically, the algorithm at time t_n proceeds as follows. The conditional distribution from the previous step, $p(\mathbf{x}_{n-1}|\mathbf{Y}_{n-1})$, is represented by M discrete weighted samples $\mathbf{x}_{n-1}^m$:

$$p(\mathbf{x}_{n-1}|\mathbf{Y}_{n-1}) \approx \sum_{m=1}^{M} w_{n-1}^m \delta(\mathbf{x}_{n-1} - \mathbf{x}_{n-1}^m), \tag{2.15}$$

which when inserted into the Bayesian recursion Eq. 2.4, gives

$$p(\mathbf{x}_n|\mathbf{Y}_n) \propto p(\mathbf{y}_n|\mathbf{x}_n) \int d\mathbf{x}_{n-1} p(\mathbf{x}_n|\mathbf{x}_{n-1}) \sum_{m=1}^{M} w_{n-1}^{m} \delta(\mathbf{x}_{n-1} - \mathbf{x}_{n-1}^{m}). \tag{2.16}$$

The integrand here cannot be sampled, so again the goal is to incorporate a feasible proposal distribution. In SMC, it is assumed that the proposal distribution can be separated into time-adjacent factors – this is what allows us to sample iteratively: $q(\mathbf{X}_n|\mathbf{Y}_n) = \prod_{n=1}^{N} q(\mathbf{x}_n|\mathbf{X}_{n-1}, \mathbf{Y}_n)$. Inserting the nth factor of this decomposition into Eq. 2.16:

$$p(\mathbf{x}_n|\mathbf{Y}_n) \propto \sum_{m=1}^{M} w_{n-1}^{m} \int d\mathbf{x}_{n-1} \frac{p(\mathbf{y}_n|\mathbf{x}_n) p(\mathbf{x}_n|\mathbf{x}_{n-1})}{q(\mathbf{x}_n|\mathbf{X}_{n-1}, \mathbf{Y}_n)} q(\mathbf{x}_n|\mathbf{X}_{n-1}, \mathbf{Y}_n) \delta(\mathbf{x}_{n-1} - \mathbf{x}_{n-1}^{m}). \tag{2.17}$$

Since delta functions similar to the last factor in this expression have propagated from previous times, we anticipate that not only $\mathbf{x}_{n-1} \to \mathbf{x}_{n-1}^{m}$, but the stronger statement $\mathbf{X}_{n-1} \to \mathbf{X}_{n-1}^{m}$. Then, carrying out the $\int d\mathbf{x}_{n-1}$ integral gives

$$p(\mathbf{x}_n|\mathbf{Y}_n) \propto \sum_{m=1}^{M} w_{n-1}^{m} \frac{p(\mathbf{y}_n|\mathbf{x}_n) p(\mathbf{x}_n|\mathbf{x}_{n-1}^{m})}{q(\mathbf{x}_n|\mathbf{X}_{n-1}^{m}, \mathbf{Y}_n)} q(\mathbf{x}_n|\mathbf{X}_{n-1}^{m}, \mathbf{Y}_n). \tag{2.18}$$

We are now in a position to approximate $p(\mathbf{x}_n|\mathbf{Y}_n)$ by sampling from $q(\cdot)$. Technically, we could choose any number of samples L, independent of the existing particles M,

$$p(\mathbf{x}_n|\mathbf{Y}_n) \propto \sum_{m=1}^{M} \sum_{l=1}^{L} w_{n-1}^{m} \frac{p(\mathbf{y}_n|\mathbf{x}_n^{l}) p(\mathbf{x}_n^{l}|\mathbf{x}_{n-1}^{m})}{q(\mathbf{x}_n^{l}|\mathbf{X}_{n-1}^{m}, \mathbf{Y}_n)} \delta(\mathbf{x}_n^{l} - \mathbf{x}_n), \tag{2.19}$$

but in this scheme the number of samples would increase exponentially in n. The idea of the particle filter is that a *single* sample is chosen for each existing particle, $\mathbf{x}_n^{m} \sim q(\mathbf{x}_n|\mathbf{X}_{n-1}^{m}, \mathbf{Y}_n)$, giving:

$$p(\mathbf{x}_n|\mathbf{Y}_n) \propto \sum_{m=1}^{M} w_{n-1}^{m} \frac{p(\mathbf{y}_n|\mathbf{x}_n^{m}) p(\mathbf{x}_n^{m}|\mathbf{x}_{n-1}^{m})}{q(\mathbf{x}_n^{m}|\mathbf{X}_{n-1}^{m}, \mathbf{Y}_n)} \delta(\mathbf{x}_n^{m} - \mathbf{x}_n), \tag{2.20}$$

Thus, we have the following picture: the conditional distribution at time t_n is represented by M

particles (consider them "propagated" particles from the previous step), with corresponding weights:

$$\tilde{w}_n^m = w_{n-1}^m \frac{p(\mathbf{y}_n|\mathbf{x}_n^m)p(\mathbf{x}_n^m|\mathbf{x}_{n-1}^m)}{q(\mathbf{x}_n^m|\mathbf{X}_{n-1}^m,\mathbf{Y}_n)} \tag{2.21}$$

$$w_n^m = \frac{\tilde{w}_n^m}{\sum_{m=1}^M \tilde{w}_n^m} \tag{2.22}$$

SMC propagates particles by sampling from the proposal distribution $q(\cdot)$ at each time step, compensating for this surrogate distribution with iteratively updated weights w_n^m, which in turn only require direct evaluations of the (unnormalized) measurement and transition distributions. In this way, the filtered distribution represents the full Bayesian integrals in a tractable recursive manner that makes no assumptions about linearity, unimodality, or Gaussianity. Estimators can then be calculated by appropriately weighted sample means, variances, etc., for example:

$$\mathbf{u}_n = \langle \mathbf{x}_n \rangle = \sum_{m=1}^M w_n^m \mathbf{x}_n^m. \tag{2.23}$$

2.3.1 The bootstrap filter

One common choice for proposal distribution is the so-called "bootstrap filter," in which the transition density is used as the proposal density [Gordon et al., 1993],

$$q(\mathbf{x}_n|\mathbf{X}_{n-1}^m,\mathbf{Y}_n) = p(\mathbf{x}_n|\mathbf{x}_{n-1}^m), \tag{2.24}$$

giving weight updates of

$$\tilde{w}_n^m = w_{n-1}^m p(\mathbf{y}_n|\mathbf{x}_n^m). \tag{2.25}$$

The bootstrap filter is simple to implement and broadly useful for nonlinear Gaussian Markov models in which the transition density $p(\mathbf{x}_n^m|\mathbf{x}_{n-1}^m)$ is straightforward to evaluate. The bootstrap filter was the first appearance of the particle filter methodology in the literature.

2.3.2 Resampling

A central issue of SMC is that after a few time steps, all but one of the samples will carry negligible weight – there is no distribution whatsoever. This can be solved fairly efficiently with a further *resampling* step, in which, at all (or a select number of) timepoints, new particles $\mathbf{x}_n^{m,\mathrm{r}}$ are regenerated by sampling from the distribution itself:

$$\mathbf{x}_n^{m,\mathrm{r}} \sim \sum_{l=1}^{M} w_n^l \delta(\mathbf{x} - \mathbf{x}_n^l) \tag{2.26}$$

The resampled particles $\mathbf{x}_n^{m,\mathrm{r}}$ are then used as the propagated particles for the next filtering step at t_{n+1}, but now with the weights reset:

$$w_n^m \to 1/M \tag{2.27}$$

$$\mathbf{x}_n^{m,\mathrm{r}} \to \mathbf{x}_n^m. \tag{2.28}$$

Though resampling increases the estimate variance, it removes particles of low weight and instead reshuffles the existing particles to regions of higher probability density, preventing the exponential suppression of particles outside the distribution peaks.

2.4 The limitations of the sequential filter

2.4.1 Filtering in neural estimation

A number of studies, having recognized the similarity of high-dimensional and complex meteorological models to neural systems, have employed Kalman-based filters or SMC to estimate unknowns in neurons and networks [Huys and Paninski, 2009, Ullah and Schiff, 2009, Vogelstein et al., 2009, Meng et al., 2011, Vavoulis et al., 2012]. These works can produce estimates of acceptable accuracy but are narrow in scope. First, the estimated models are often limited to reduced Hodgkin-Huxley models with few state variables, and though SMC has been applied to more detailed compartmental models, estimation accuracy is poor without strong priors on (or even complete knowledge of) nonlinear parameters such as voltage thresholds in gating kinetics [Vavoulis et al., 2012]. None of the models exhibit chaos, operating mainly in the regime between statis and periodicity.

To the point of parameters, recall that state-space filtering – basically an iterative predict-and-correct method for *dynamical* systems – is an unnatural setting for the estimation of static quantities; parameter estimation must therefore be incorporated in some other way. One method appropriate to SMC is the expectation-maximization algorithm (E-M) [Dempster et al., 1977, MacKay, 2003, Kantas et al., 2015]. E-M, in conjunction with SMC, alternates i) particle filtering (with fixed parameters) with ii) optimization of expected joint likelihoods over the unknown parameter space. As it is only guaranteed to converge to local minima, it is generally only useful for parameters that enter HH dynamics linearly, such as maximal conductances and some timescales. Further, since state variable estimates are themselves dependent on the parameter estimate $\hat{\theta}_{n-1}$ at the prior time, then transitions likelihoods $p(\mathbf{x}_n|\mathbf{x}_{n-1}) = p(\mathbf{x}_n|\mathbf{x}_{n-1}, \hat{\theta}_{n-1})$ conditioned on inaccurate $\hat{\theta}_{n-1}$ may cause the estimate to rapidly diverge.

Alternatively, parameters may be estimated by artificially promoting them to model states with trivial dynamics ($\mathbf{p}_{n+1} = \mathbf{p}_n + \mathbf{r}_n$), and treating them on the same footing as state variables, $\{\mathbf{x}\} \to \{\mathbf{x}, \theta\}$. It is known that this introduces bias and increases the estimate variance; as well, it may fail to adequately probe the parameter space by subjecting the parameters to highly correlated Gaussian walks [Kantas et al., 2015]. Naturally, these downsides manifest sharply with strong model nonlinearities, precluding the determination of key neural mechanisms such as gating kinetics and synaptic dynamics.

2.4.2 Breakdown of the Gaussian assumption

While estimate divergence compounds with nonlinearities, sparse and noisy data, and the presence of unknown static parameters, it could be managed with sufficient *a priori* knowledge via $p(\mathbf{x}_0, \theta_0)$. A well-informed prior places the estimate in the vicinity of the underlying trajectory; this is critical because nonlinearities may render the transition probabilities not only non-normal but multimodal entirely. Weak priors may place the estimate in an auxiliary mode of the distribution, subjecting Kalman filters to linearizing approximations around a *local* maxima of $p(\mathbf{x}_n|\mathbf{Y}_n)$ that may never allow return to the true trajectory. Thus, estimate "divergence" equates to tracking peripheral modes of the conditional distribution. As we will soon see, these modes are prevalent in the estimation of nonlinear systems and ubiquitous in chaotic ones.

Since SMC works outside of the linearizing and Gaussian assumptions, one might contend that it is more immune to estimate divergence. Its efficacy is tempered by dimensionality, as the variance of the

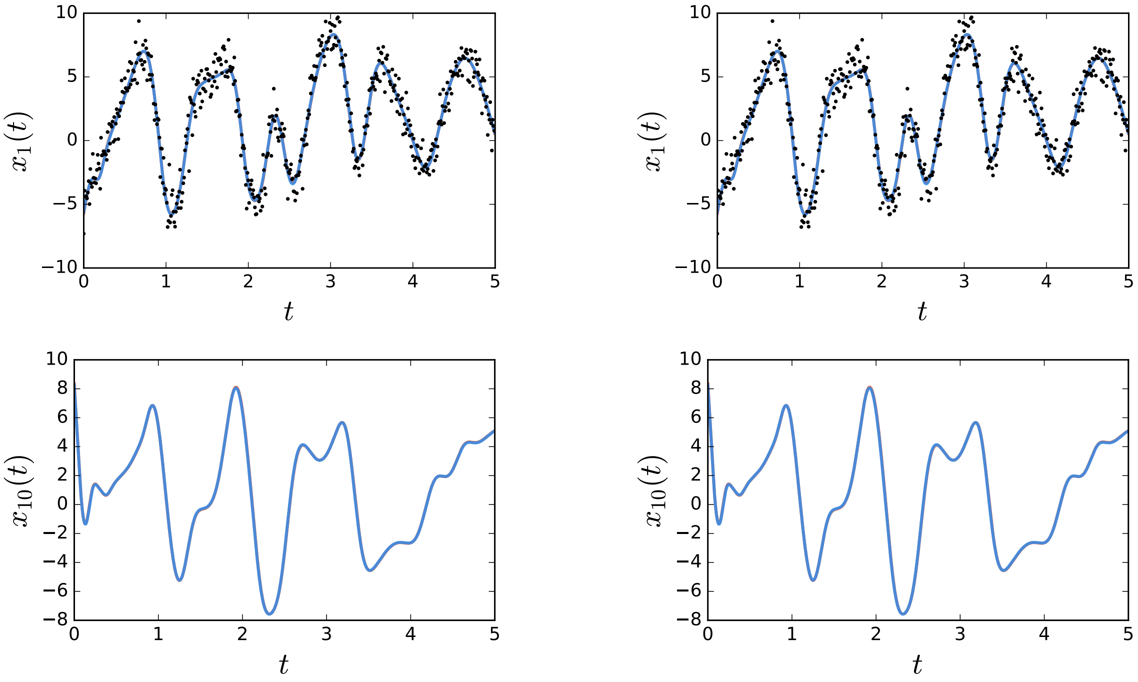

(a)

(b)

Figure 2.1: (a) Estimated trajectories of an observed component (top) and unobserved component (bottom) of the chaotic Lorenz 1996 system with 9 of 10 components observed. Here the mean of the unobserved component's prior distribution is chosen from a normal distribution of deviation $\sigma_{10} = 5.0$ around its true value. (b) The same estimated trajectories, but with $\sigma_{10} = 10.0$.

estimate increases quadratically in timesteps N, and there is evidence that the number of necessary particles to maintain accuracy below observation noise scales exponentially with a parameter that depends collectively on i) model dimension, ii) variance of the prior, and iii) observation characteristics [Snyder et al., 2008]. Applying a variant of SMC to the chaotic Lorenz model in 40 dimensions with 20 measured variables, for example, it was observed that 500-1000 particles were needed to bound the variance below measurement noise [Lorenz, 2006, Nakano et al., 2007]. This bodes poorly for higher-dimensional neural networks, when observations are sparse.

To illustrate how weak priors and sparse observations may conspire with multimodal conditionals to produce diverging estimates, consider the unscented Kalman smoother (UKF with an analogous backward pass) applied to the $D = 10$-dimensional chaotic Lorenz 1996 model, whose continuous dynamics are given by

$$\frac{dx_d}{dt} = f_d(\mathbf{x}) = -x_{d-2}x_{d-1} + x_{d-1}x_{d+1} - x_d + f \tag{2.29}$$

where $f = 8.0$, the subscripts are state variable indices, $\mathbf{x} = \{x_d\}$, and the states are subject to periodic boundary conditions: $x_{d+D} = x_d$ [2]. Synthetic observations are generated by numerically integrating Eqs. 2.29 forward from a given initial condition $\mathbf{x}(0)$ for 501 timesteps in intervals of $\Delta t = 0.01$ – this is the true trajectory $\mathbf{x}^*(t)$ – and then adding iid Guassian noise of zero mean and variance $\sigma^2 = 0.1$. We first assume that $x_1(t), ..., x_9(t)$ are measured, i.e $\mathbf{H}$ is a rank-9 projection operator.

Since the only information at $t = 0$ is the observations of the measured subspace, let us assume a highly peaked prior for these components, centered on the data:

$$p(x_d(0)) = \mathcal{N}(y_d(0), \sigma_d \ll 1) \qquad d = 1, ..., 9. \tag{2.30}$$

On the other hand, there is no *a priori* information about $x_{10}(0)$, it is initialized somewhere within a given range. Specifically, it is chosen from a normal distribution centered on its true value with a given deviation:

$$p(x_{10}(0)) = \mathcal{N}(x^*_{10}(0), \sigma_{10}). \tag{2.31}$$

The results of these estimations are shown in Fig. 2.1, for $\sigma_{10} = 5$ and $\sigma_{10} = 10$. The information passed from the observed variables $\{x_1 ... x_9\}$ via the dynamics is sufficient to align the estimates to of both observed and unobserved states to their true trajectories, *despite* a virtually uninformative prior on the latter.

If on the other hand, only $\{x_1, ... x_5\}$ are observed, then the filter may end up tracking a distinct mode of the system far afield of the true underlying trajectory. Examples of this are shown in Fig. 2.2, for $\sigma_6, ..., \sigma_{10} = 2.0$ and $\sigma_6, ..., \sigma_{10} = 7.0$, respectively. In the former case, the smoothed UKF estimate matches the true trajectory, but when the initialization is outside of this neighborhood as in the latter case, a distinct mode is tracked and the unobserved states (and even the observed trajectories), are poorly estimated. These other modes arise as an interplay of missing measurements – equivalently, large covariances in the initial state – and nonlinear dynamics. The Gaussian assumption has broken down.

The role of the prior distribution $p(\mathbf{x}_0)$ in revealing the non-Gaussian nature of the conditional distributions is illustrated explicitly in Figures 2.3 and 2.4. These figures show histograms representing

---

[2] A notational note: from here forward, timepoints may be indicated by an index as in $\mathbf{x}_n$, *or* in parentheses such as $\mathbf{x}(5.26)$ and $x_3(5.26)$. When t is represented in parentheses, it notates an exact time; when it is a subscript, it notates the nth of a discrete set of timepoints. The choice should be obvious depending on the context. Further, the Roman index d will always refer to model component, while m and n will always refer to time.

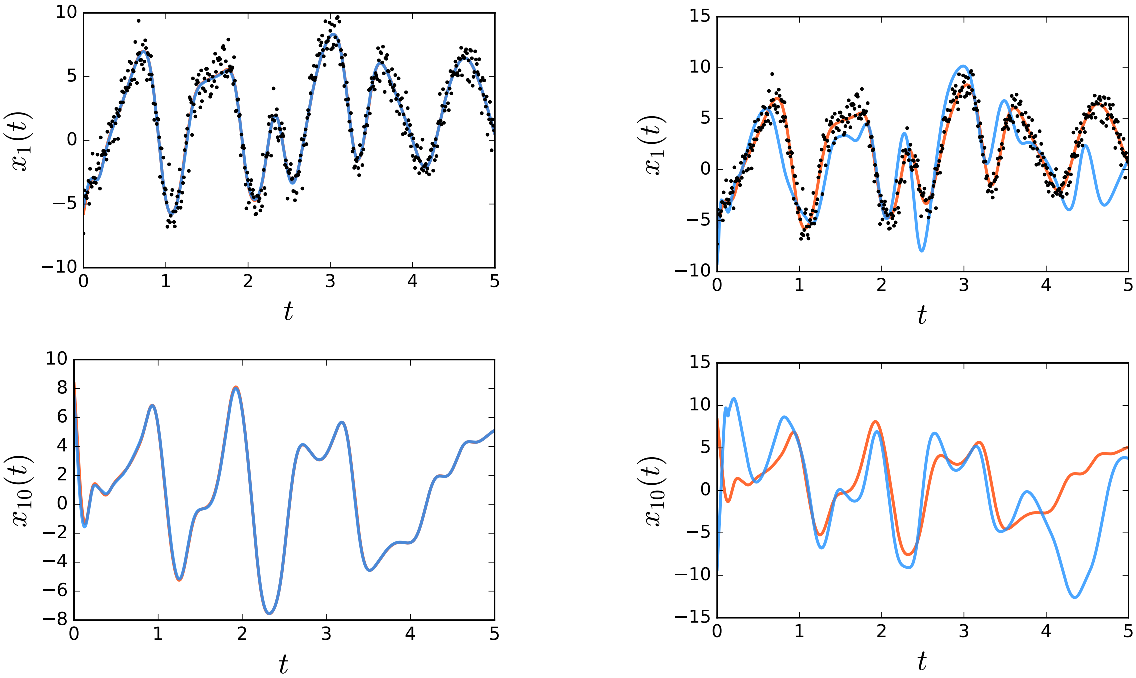

(a) (b)

Figure 2.2: (a) Estimated trajectories of an observed component (top) and unobserved component (bottom) of the chaotic Lorenz 1996 system with 5 of 10 components observed. Here the means of the unobserved components' prior distributions are chosen from a normal distribution of deviation $\sigma_d = 2.0$ around their true values. (b) The same estimated trajectories, but with $\sigma_d = 7.0$, indicating that with sparser observations, a weak prior can lead to filter divergence.

the full conditional distributions $p(\mathbf{x}_n|\mathbf{Y}_n)$ at various times throughout the estimation window, again for the estimation of the Lorenz96 model under the same condition as above. Here, only the variables $\{x_1, ..., x_5\}$ are measured. The histograms are generated using SMC with a bootstrap filter and one million particles. When the prior is centered tightly about the true initial state ($\sigma_d = 0.1$ for the unmeasured variables), the distributions are nearly Gaussian for both measured and unmeasured variable alike (Figure 2.3). Further, the peak of these distributions lies near the true state, as the estimate accurately tracks the true trajectory. When the initial variance of the unmeasured components is increased to $\sigma_d = 2.0$, however, the distributions generate several new peaks and quickly lose their Gaussian character (Fig. 2.4). The multimodality is more pronounced for the unmeasured states (lower panel).

In the context of EKF and UKF, the Gaussian approximation suppresses the other sections of the distribution in the next filtering step, leading the estimate to track a local mode of the distribution. This is suggested by the rightmost panels of Figure 2.4: the true state (yellow line) lies near a peripheral mode, yet

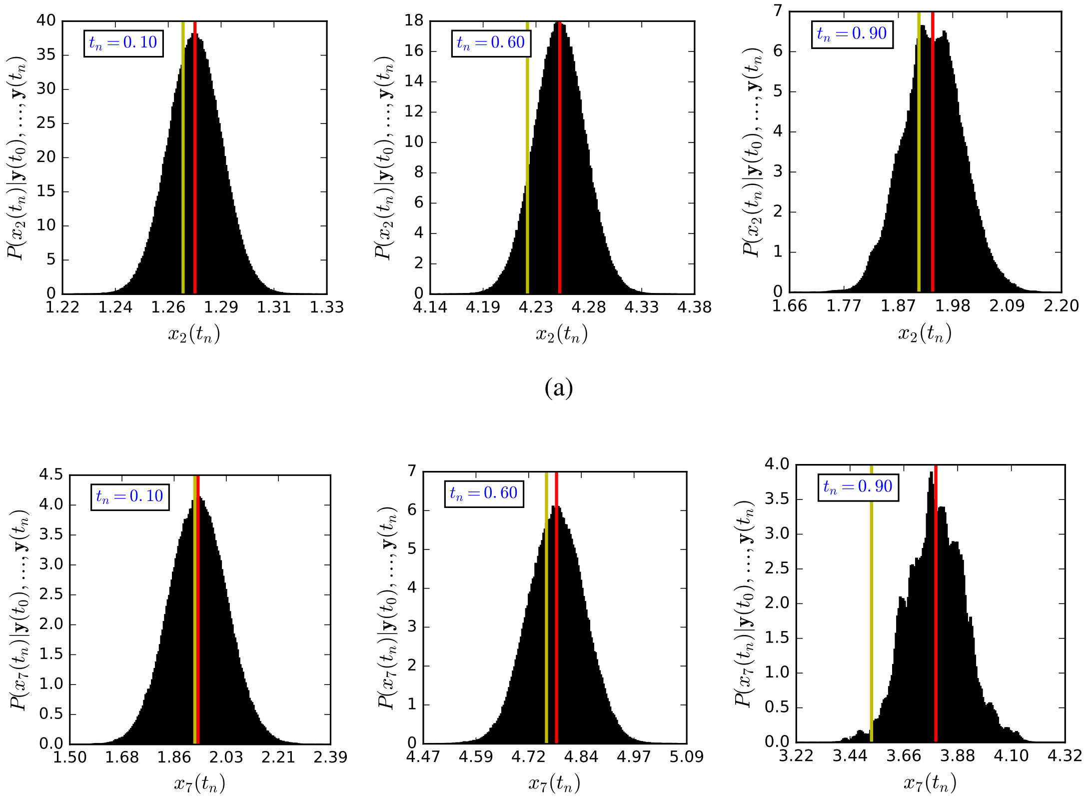

(a)

(b)

Figure 2.3: (a) The posterior conditional distributions in the estimation of the chaotic Lorenz 1996 attractor with five measured variables, at various times $t_n = 0.10, 0.60$, and 0.90. The prior distribution for the unmeasured variables is centered tightly with a deviation of just $\sigma_d = 0.1$ around the true state $x_d^*(0)$. The red line indicates the mean of the distribution at the given time, while the yellow line indicates the true state. (b) The same distributions for an unobserved component.

propagating forward a Gaussian estimate about the statistical mean (far away, in red) would fail to account for that contribution in the subsequent filtering step. Further, recall that these distributions are generated with SMC using 10^6 particles, yet the model is only 10-dimensional. In the context of SMC, more reasonably-sized ensembles would still eventually lose track of all the distribution modes, particularly as resampling shuffles the particles away frmo those that are narrow and highly-peaked, as is the mode representing the true state in the bottom-right panel of Figure 2.4.

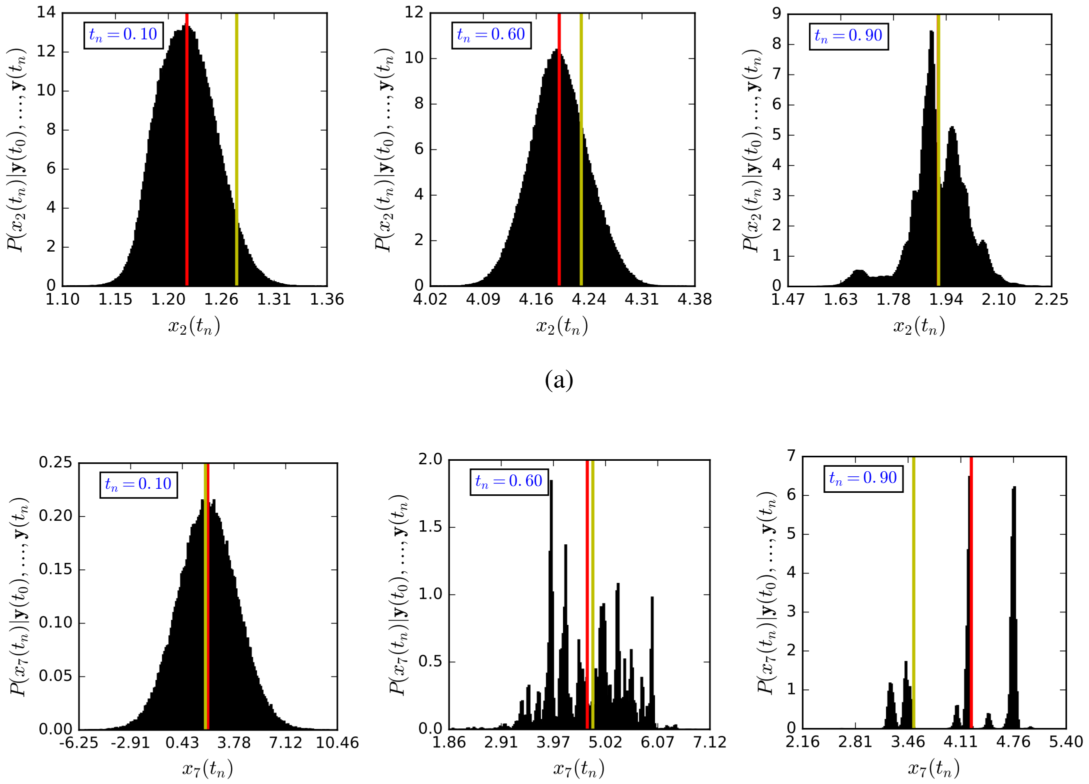

(a)

(b)

Figure 2.4: (a) The posterior conditional distributions in the estimation of the chaotic Lorenz 1996 attractor with five measured variables as in Figure 2.3. The prior distribution for the unmeasured variables are now centered more loosely – with a deviation of just $\sigma_d = 2.0$ around the true states at t_0. (b) The same distributions for an unobserved component.

2.5 Beyond filtering: the variational approximation

2.5.1 The simultaneous conditional expectation

To see how we might partially address the issue of multimodality, let us begin again with the Bayesian recursion equations, Eqs. 2.5:

$$p(\mathbf{x}_N|\mathbf{Y}_N) \propto p(\mathbf{y}_N|\mathbf{x}_N)\left[\int \prod_{n=0}^{N-1} d\mathbf{x}_n p(\mathbf{x}_{n+1}|\mathbf{x}_n) p(\mathbf{y}_n|\mathbf{x}_n)\right] p(\mathbf{x}_0). \qquad \text{(2.5 revisited)}$$

This conditional probability on the left-hand side is also the joint distribution of $\{\mathbf{x}_0, ...\mathbf{x}_N\}$, marginalized over the intermediate variables:

$$p(\mathbf{x}_N|\mathbf{Y}_N) = \int \prod_{n=0}^{N-1} d\mathbf{x}_n p(\mathbf{X}_N|\mathbf{Y}_N). \tag{2.32}$$

Equating Eqs. 2.5 and 2.32 expresses the joint conditional probability as

$$p(\mathbf{X}_N|\mathbf{Y}_N) \propto \exp[-A(\mathbf{X}|\mathbf{Y})] \tag{2.33}$$

where

$$A(\mathbf{X}|\mathbf{Y}) = -\sum_{n=0}^{N-1} \log p(\mathbf{x}_{n+1}|\mathbf{x}_n) - \sum_{n=0}^{N} \log p(\mathbf{y}_n|\mathbf{x}_n) - \log p(\mathbf{x}_0), \tag{2.34}$$

and we have adopted the shorthand $\mathbf{X}_N = \mathbf{X} = \{\mathbf{x}_0, ...\mathbf{x}_N\}$ and similarly for $\mathbf{Y}_N$. The joint conditional can be used to give estimators (e.g. expected values) of quantities $G(\mathbf{X})$ defined on the entire observation time window; for example, when G is the identity, then $\langle G(\mathbf{X})\rangle$ is the expected value of the trajectory itself. In the filtering picture, these expected values are updated sequentially in time, subject to the approximation scheme relevant to that particular procedure. In Kalman filtering, this approximation is the pruning of higher order moments (the Gaussian assumption); in SMC, it is the representation of a continuous function by discrete samples. Here, we instead represent these expected values in exact form and simultaneously in time, dispensing with the iterative context:

$$\langle G(\mathbf{X})|\mathbf{Y}\rangle = \frac{\int d\mathbf{X} G(\mathbf{X}) p(\mathbf{X}|\mathbf{Y})}{\int d\mathbf{X} p(\mathbf{X}|\mathbf{Y})} = \frac{\int d\mathbf{X} G(\mathbf{X}) \exp[-A(\mathbf{X}|\mathbf{Y})]}{\int d\mathbf{X} \exp[-A(\mathbf{X}|\mathbf{Y})]} \tag{2.35}$$

Of course, the whole point of filtering is that high-dimensional Bayesian integrals of this sort are intractable. On the other hand, the exponential form of the joint probabilities suggests that this integral may be at least approximated in the saddle point approximation, *if* the integral were dominated by one or a few such extrema of $A(\mathbf{X}|\mathbf{Y})$ [Laplace, 1986, MacKay, 2003, Abarbanel, 2013]. That is, applying Laplace's approximation to the integral Eq. 2.35, we get

$$\langle G(\mathbf{X})|\mathbf{Y}\rangle \approx \frac{\sum_q G(\mathbf{X}^q)|\partial_x^2 A(\mathbf{X}^q|\mathbf{Y})|^{-1/2}\exp[-A(\mathbf{X}^q|\mathbf{Y})]}{\sum_q |\partial_x^2 A(\mathbf{X}^q|\mathbf{Y})|^{-1/2}\exp[-A(\mathbf{X}^q|\mathbf{Y})]}, \tag{2.36}$$

where $|\partial_x^2 A(\mathbf{X}^q|\mathbf{Y})|$ denotes the determinant of the Hessian matrix of A with respect to $\mathbf{X}$ evaluated at $\mathbf{X}^q$ and the extrema are defined by

$$\left.\frac{\partial A(\mathbf{X}|\mathbf{Y})}{\partial \mathbf{X}}\right|_{\mathbf{X}^q} = 0. \tag{2.37}$$

In other words, Laplace's method [3] requires the minimization of $A(\mathbf{X}|\mathbf{Y})$ over an ND-dimensional manifold housing the full collocated trajectories $\mathbf{X}$, and relies on the ability to find the few most dominant such $\mathbf{X}^q$. Yet the pervasive multimodality we found in the conditional distributions of sequential filtering bodes poorly for the variational method: $A(\mathbf{X}|\mathbf{Y})$ is the negative log likelihood of the joint conditional probability, so a plethora of modes in the latter translate to a plethora of minima in the former. Though we expect that the *true* trajectory $\mathbf{X}$ would give a substantial, possibly the largest, contribution to the expected value integral, it is likely hidden among many others in the $\mathbb{R}^{ND}$ distribution support.

2.5.2 The ubiquity of local minima in $A(\mathbf{X}|\mathbf{Y})$ when model linearity is broken

Let us draw again on the instructive chaotic Lorenz 1996 model in ten dimensions to investigate the density of extrema in the collocated log-likelihood $A(\mathbf{X}|\mathbf{Y})$. As in Section 2.4.2, we again use synthetic data $\mathbf{Y}$ generated by integrating the model equations over 201 timepoints of $\Delta t = 0.01$ and add iid standard Gaussian noise. We can express the discrete transition probabilities $p(\mathbf{x}_{n+1}|\mathbf{x}_n)$ that enter $A(\mathbf{X}|\mathbf{Y})$ by discretizing the continuous time Lorenz model dynamics (Eq. 2.29) with an implicit trapezoidal scheme, adding small process noise of zero mean and small variance $Q =$1e-6 to each variable:

$$-\log p(\mathbf{x}_{n+1}|\mathbf{x}_n) = \sum_{d=1}^{D} \frac{1}{2Q}\left(x_d(t_{n+1}) - x_d(t_n) - \frac{\Delta t}{2}[f_d(\mathbf{x}(t_n)) + f_d(\mathbf{x}(t_{n+1}))]\right)^2. \tag{2.38}$$

One can also think of this as a discrete-time Markov process whose transition probabilities from $\{t_n, x_d(t_n)\}$ to $\{t_{n+1}, x_d(t_{n+1})\}$ are normally-distributed with mean $x_d(t_n) + \frac{\Delta t}{2}[f_d(\mathbf{x}(t_n)) + f_d(\mathbf{x}(t_{n+1}))]$ and variance Q. Fi-

---

[3] With some abuse of terminology, *method of stationary phase*, *Laplace's method*, and *variational method* will be used more or less interchangeably throughout this thesis.

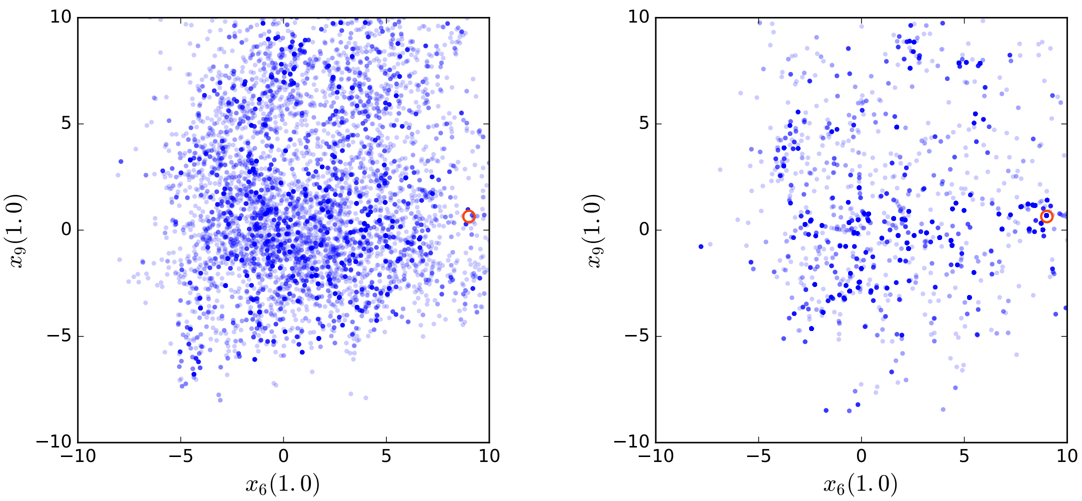

(a)

(b)

Figure 2.5: (a) A two-dimensional projection of the cost surface minima in the estimation of the Lorenz 1996 chaotic attractor with one measured variable. Of 10,000 initializations, more than 4000 distinct minima were found. The true state is indicated by the red circle. (b) The same projection, with 5 measured variables; in this case, nearly 900 distinct minima were found.

nally, the observations are assumed normally distributed, centered on the states with variance unity, whereby

$$-\log p(\mathbf{y}_n|\mathbf{x}_n) = \sum_{d=1}^{L} \frac{1}{2}(y_d(t_n) - x_d(t_n))^2, \tag{2.39}$$

where it should be noted that the sum over states only extends to the number of observations L. Finally, assuming a uniform prior on the system at t_0, $p(\mathbf{x}_0)$ is a constant. $A(\mathbf{X}|\mathbf{Y})$ is thus:

$$A(\mathbf{X}|\mathbf{Y}) = \sum_{n=0}^{N-1} \text{Eq. } 2.38 + \sum_{n=0}^{N} \text{Eq. } 2.39 + \text{constant} \tag{2.40}$$

To visualize the preponderance of extrema along the surface of $A(\mathbf{X}|\mathbf{Y})$, let us minimize the function several times over using a local minimizer, specifically the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [Fletcher, 1987, Broyden, 1970, Fletcher, 1970]. To hit as many optima as possible, we will perform 10,000 distinct optimizations by uniformly sampling an initializing point in the model's approximate dynamical range (a ND-hypercube of side length 30 centered at the origin), letting the optimization proceed from there. To visualize the density of local minima in this high-dimensional manifold, we then project the $ND = 2010$ dimensions down to only two arbitrary coordinates $x_6(t_n = 1.0)$ and $x_9(t_n = 1.0)$.

These $2D$ projections are shown in Figures 2.5a and 2.5b, for $L = 1$ and $L = 5$ measured variables, respectively. When $A(\mathbf{X}|\mathbf{Y})$ contains observations from only the single component x_1, the 10,000 optimizations return more than 4000 distinct estimates (Figure 2.5a). If the five components $\{x_1,...x_5\}$ are observed, the optimization returns nearly 900 distinct estimates. In the first case, finding the lowest minima that contribute to the conditional integral Eq. 2.35 is hopeless. In fact, even with relatively high observability of 50%, the number of minima far exceeds the model dimension, a scenario that does not portend well for even the tiniest biophysical neural network. And yet, the extended Kalman smoother could estimate this same model, reliably and accurately, with 50% observability. It is unclear what the variational formulation buys us, if anything at all.

2.5.3 The Gaussian error approximation

To proceed, let us first make a few simplifying assumptions about the observation noise and model stochasticity, which will allow the conditional log-likelihood to be written in a more concrete form, and suggest a way to rectify the problem of multimodality confronted in the previous sections. Specifically, we first assume that L components of the full state can be directly observed, but these observations are corrupted by normally-distributed noise uncorrelated in space and time. Secondly, we assume as before that the model is a discrete-time stochastic Markov process; we can think of the process as the discretization of a continuous time Langevin equation:

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}) + \boldsymbol{\xi}(t) \quad \rightarrow \quad \mathbf{x}_{n+1} = \mathbf{g}(\{\mathbf{x}_m\},\{t_m\},\Delta t) + \mathbf{q}_n, \tag{2.41}$$

where $\mathbf{g}(\cdot)$ is some numerical discretization of the continuous deterministic dynamics $\mathbf{f}(\cdot)$, and the set notation $\{\mathbf{x}_m\},\{t_m\}$ indicates that $\mathbf{g}$ could in theory depend on several $\mathbf{x}_m$ and/or several timepoints t_m prior. The discrete time forward map $\mathbf{g}$ is still a Markov process, though possibly of higher order if the discretization is beyond second-order. The fluctuating components $\boldsymbol{\xi}(t)$ and $\mathbf{q}_n$ are Gaussian stochastic variables uncorrelated in both space and time, with

$$\begin{aligned} \langle \mathbf{q}_n \rangle &= 0 \\ \langle \mathbf{q}_m \mathbf{q}_n \rangle &= \mathbf{Q}_n \delta_{mn}, \end{aligned} \tag{2.42}$$

where $\mathbf{Q}_n$ is a diagonal square matrix of dimension D. Finally, we assume an agnostic prior density on the system at t_0, whereby $\log p(\mathbf{x}_0)$ is constant which we can ignore. We may then write $A(\mathbf{X}|\mathbf{Y})$ explicitly as

$$A(\mathbf{X}|\mathbf{Y}) = \sum_{n=1}^{N} (\mathbf{y}_n - \mathbf{H}\mathbf{x}_n)^2\big|_{\mathbf{R}_n^{-1}} + \sum_{n=0}^{N} (\mathbf{x}_{n+1} - \mathbf{g}(\{\mathbf{x}_m\},\{t_m\},\Delta t))^2\big|_{\mathbf{Q}_n^{-1}}, \tag{2.43}$$

where $\mathbf{R}_n$ is the L-dimensional covariance matrix of the observations at time t_n, and $\mathbf{H}$ is a D-dimensional projection matrix of rank L. For ease of notation, the following component form of these quantities will also be used in this thesis:

$$A(\mathbf{X}|\mathbf{Y}) = \sum_{n=1}^{N}\sum_{d=1}^{L} \frac{R_{\mathrm{m}}(n,d)}{2}(y_d(t_n) - x_d(t_n))^2 + \sum_{n=0}^{N}\sum_{d=1}^{D} \frac{R_{\mathrm{f}}(n,d)}{2}(x_d(t_{n+1}) - g_d(\{\mathbf{x}_m\},\{t_m\},\Delta t))^2 \tag{2.44}$$

where

$$\mathbf{Q}_n^{-1} = \begin{bmatrix} R_{\mathrm{f}}(n,1) & & \\ & \ddots & \\ & & R_{\mathrm{f}}(n,D) \end{bmatrix} \tag{2.45}$$

and

$$\mathbf{R}_n^{-1} = \begin{bmatrix} R_{\mathrm{m}}(n,1) & & & \\ & \ddots & & \\ & & R_{\mathrm{m}}(n,L) & \\ & & & \mathbb{0}_{D-L} \end{bmatrix}. \tag{2.46}$$

Refer to the expressions Eq. 2.43 and Eq. 2.44 as the *Gaussian error approximation*, emphasizing that the errors are Gaussian, even though the distributions themselves are not.

2.5.4 Manifold structure of $A(\mathbf{X}|\mathbf{Y})$

The structure of the $A(\mathbf{X}|\mathbf{Y})$ manifold arises from the competing interplay of two sums in Eq. 2.43. The first "measurement term" localizes the estimate about the observed data; it is quadratic in the measured components and flat in the unmeasured ones, a two-dimensional projection of which is shown in Figure 2.6a. The second sum consisting of "model terms" removes this degeneracy along the unmeasured directions in

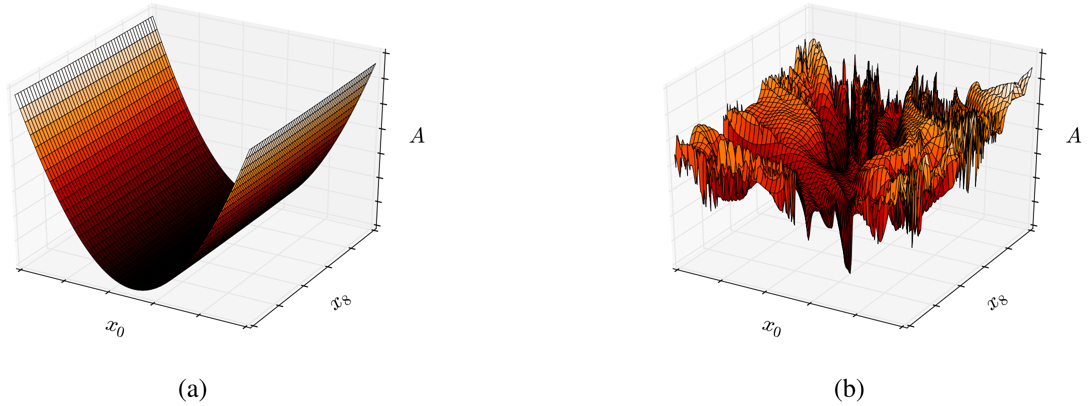

Figure 2.6: (a) The contribution of the measurement term to the cost surface $A(\mathbf{X}|\mathbf{Y})$, projected to two variables. The surface is quadratic in the measured direction x_9 and degenerate in the unmeasured x_6. (b) Same projections for the model term in $A(\mathbf{X}|\mathbf{Y})$, showing an approximately fractal surface arising from the chaotic dynamics.

a most dispiriting way: local minima permeate the $A(\mathbf{X}|\mathbf{Y})$, portending sever issues for local optimization routines. As we would expect a fully quadratic $A(\mathbf{X}|\mathbf{Y})$ with linear dynamics, indeed the fractal-like surface of the model term contributions (Figure 2.6b) arises from the chaotic nature of the Lorenz model, a connection that has been noted in the context of parameter estimation in engineering and biological applications, as well as in chaotic control and synchronization [Abarbanel et al., 2009, Zelinka et al., 2014].

The interplay of the two competing contributions – one tractable but under-determined, one unambiguous but convoluted – depends sensitively on the balance of the inverse noise covariances R_f and R_m, or more generally, between the spectra of $\mathbf{Q}_n$ and $\mathbf{R}_n$. $\mathbf{R}_n$ is determined by experimental setup – the precision of the measurement apparatus – but the choice of $\mathbf{Q}_n$ is often heuristic and unknown. Typically, it is tuned in a way to most adequately smooth the estimated trajectory from noisy data or produce faster convergence [Bavdekar et al., 2011, Karasalo and Hu, 2011, Feng et al., 2014]. Filter performance is quite sensitive to model covariance; while high covariance can compromise estimation accuracy, small $\mathbf{Q}_n$ will (in SMC for example) produce *sample impoverishment* – strong confidence in the dynamics renders the filtered distribution highly peaked, leading to an ensemble consisting of many copies of only a few distinct samples [Särkkä, 2013]. In the Kalman filter, unnaturally high confidence in the dynamics can heavily bias the estimate and lead to divergence [Karasalo and Hu, 2011].

We might expect that the variational approximation to the full simultaneous conditional distribution Eq. 2.5 would be less susceptible to these issues, the iterative context having been dispensed with. But the consequences of low process noise manifest in a a related way. To illustrate this, Figure 2.7 shows the progression of $2D$ projections of $A(\mathbf{X}|\mathbf{Y})$ from $R_{\mathrm{f}} = 1\text{e-}5$ to $R_{\mathrm{f}} = 1\text{e}7$ in the vicinity of the true path of the $10D$ Lorenz 1996 model. In the variational context, increasing R_{f} sharpens the minima of A, much like what occurs in the filter. The expected value integral is dominated by the needle-like valley in the last panel, but the basin of attraction for this minimum narrows considerably with increasing R_{f}. Among the many local minima in $A(\mathbf{X}|\mathbf{Y})$ – as we saw in Figure 2.5 – the true path becomes more difficult to locate as the model dynamics are enforced. Any local optimization routine seeking the stationary points of $A(\mathbf{X}|\mathbf{Y})$ would converge to a spurious minimum in much the same way as the iterative filter. Though the symptoms differ, the upshot is the same.

Yet the process noise, as mentioned, is often a heuristic parameter informed less by physical considerations than by the need for filter accuracy and performance. Much as it can be tuned in the Kalman filter, it could likewise be promoted to an algorithmic parameter in the variational framework. Indeed, the gradual deformation of $A(\mathbf{X}|\mathbf{Y})$ – from one defined by a valley of degenerate minima to one nearly fractal – suggests an iterative way to do this within the structure of the Gaussian error approximation.

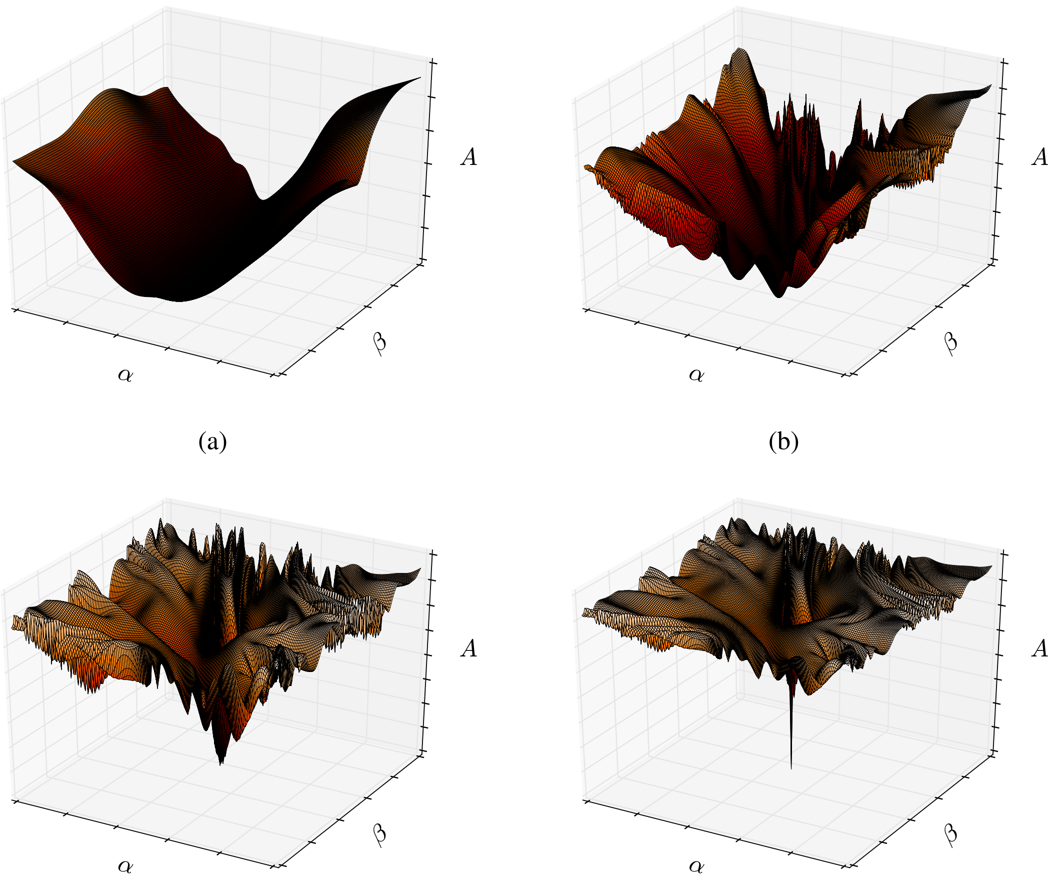

(a) (b)

(c) (d)

Figure 2.7: (a) Projections onto a measured variable α and an unmeasured variable β of the cost surface $A(\mathbf{X}|\mathbf{Y})$ for low model precision, $R_{\rm f} = $ 1e-5. The surface is nearly convex in the measured direction and highly degenerate in the measured direction. (b) Same projections for higher model precision, $R_{\rm f} =$1e-1; the degeneracy in the unmeasured directions has partially lifted (c) $R_{\rm f}$ = 1e3 (d) $R_{\rm f}$ = 1e7. Not only has the degeneracy lifted, but the global minimum is not hidden in a small region with a tiny basin of attraction, hindering its detection in the full high-dimensional state space.

Chapter 3

Variational annealing

3.1 The variational annealing algorithm

The challenges posed by the structure of the Gaussian error approximation manifold $A(\mathbf{X}|\mathbf{Y})$ were illustrated quite vividly by the surface plots in Section 2.5.4. It was suggested that this challenge might be tackled by exploiting the heuristic character of the process noise $\mathbf{Q}_n$. To this end, let us begin with the observation that $A(\mathbf{X}|\mathbf{Y})$ simplifies when the process noise is maximal: when $R_\mathrm{f}(n,d) \approx 0$ (where ≈ 0 is the order of machine precision), then

$$A(\mathbf{X}|\mathbf{Y})_{Q\to\infty} = \sum_{n=1}^{N} (\mathbf{y}_n - \mathbf{H}\mathbf{x}_n)^2|_{\mathbf{R}_n^{-1}} \tag{3.1}$$

$$= \sum_{n=1}^{N}\sum_{d=1}^{L} \frac{R_\mathrm{m}(n,d)}{2}(y_d(t_n) - x_d(t_n))^2. \tag{3.2}$$

$A(\mathbf{X}|\mathbf{Y})_{Q\to\infty}$ contains no information from the model dynamics, and relies solely on information from the observations. The stationary solution is nearly degenerate, characterized by

$$x_d(t) \approx y_d(t) \qquad\qquad d = 1,...,L \tag{3.3}$$

and $x_d(t)$ unconstrained for $d = L+1,...,D$. Assuming that the process noise is time-independent and uniform across state variables, so $R_\mathrm{f}(n,d) \equiv R_\mathrm{f}$, we could now use this path as the seed for a subsequent optimization

Algorithm 1: Variational Annealing (VA)

```
Input  : Observations $\mathbf{Y}$ and annealing parameters $\beta_{\max}$, $\alpha$, $R_{f0}$
Output: Estimated trajectories and parameters $\{\hat{\mathbf{X}}^{q,\beta}\}$
for $q = 1, ..., Q$, in parallel do
    Sample $\mathbf{X}^q_{\text{init}}$ uniformly from presumed dynamical range
end
for $\beta \leftarrow 0$ to $\beta_{\max}$ do
    $R_f \leftarrow R_{f0}\alpha^\beta$
    $A_\beta(\mathbf{X}) \leftarrow A(\mathbf{X}|\mathbf{Y})|_{R_f}$ (Eq. 2.44)
    for $q = 1, ..., Q$, in parallel do
        $\hat{\mathbf{X}}^{q,\beta} \leftarrow$ argmin $A_\beta(\mathbf{X})$ with initial guess $\mathbf{X}^q_{\text{init}}$
        $\mathbf{X}^q_{\text{init}} \leftarrow \hat{\mathbf{X}}^{q,\beta}$
    end
end
```

in which R_f has been slightly increased, the degeneracy in the unmeasured directions will lift. Many such optimizations initialized in parallel from low R_f will return distinct estimates as this alternating optimization-initialization procedure continues. Further, if the increase in R_f is sufficiently gradual, some of these estimates may track the lowest minima whose basins of attraction narrow with higher R_f as seen in Figure 2.7d. In this way, the largest contributions to the conditional expectation Eq. 2.35 value could be systematically followed to the tiny basins that are virtually invisible for high model precision.

With this brief outline, let us propose the following annealing-like algorithm. To broaden the generality of this method, and to emphasize the equal footing of parameters and states, we write $\mathbf{X}$ to include static unknown parameters $\mathbf{P}$:

$$\{\mathbf{X}\} \equiv \{\mathbf{x}_1, ..., \mathbf{x}_N, \mathbf{P}\}. \tag{3.4}$$

Choosing an initially small $R_f = R_{f0}$, perform Q minimizations of $A(\mathbf{X}|\mathbf{Y})$ over the joint state and parameter space $\mathbf{X}$. At this first step, these optimizations are initialized randomly, generally by uniformly sampling the states $x_d(t)$ from their dynamical range and the parameters from their presumed range. This is the first step of the anneal, and the resulting paths $\hat{\mathbf{X}}^{q,0}$ of these optimizations are then used as the initialization for a new minimization with $R_f = R_{f0}\alpha$, where α is fixed at some value larger than one. This is continued for $\beta = 1, ..., \beta_{\max}$, with R_f increasing from $R_{f0}\alpha^\beta$ to $R_{f0}\alpha^{\beta+1}$ at each step. The result is a collection of Q

estimated paths and parameters for each annealing level β, labeled as

$$\{\hat{\mathbf{X}}^{q,\beta}\}, \tag{3.5}$$

with corresponding

$$\hat{A}^{q,\beta} \equiv A(\hat{\mathbf{X}}^{q,\beta}|\mathbf{Y}). \tag{3.6}$$

This procedure, which we call *variational annealing*, is described in Algorithm 1 and was introduced in [Ye et al., 2015a, Ye et al., 2015b] and utilized in [Breen et al., 2016, Kadakia et al., 2016] where further illustrations and discussion can be found.

3.1.1 Variational annealing as homotopy continuation

The graduation deformation of the cost surface resembles a similar method for tracking solutions to nonlinear differential equations, called *homotopy continuation* [Allgower and Georg, 1990, Li, 1997]. The idea of homotopy continuation is that the solution of a nonlinear system of equations might be reliably tracked by beginning with the known solution of a simpler but related system which can be deformed into the original one in a continuous way. For example, $g(x) = x^3 - 8 = 0$ can be deformed into $x^3 - x^2 - 4 = 0$ by solving $g(x,\gamma) = x^3 - 8 + \gamma(-x^2 + 4) = 0$ from $\gamma : 0 \rightarrow 1$. That is, one uses the solution of $g(x,0) = 0 \rightarrow x = 2$ to initialize a search for $g(x,1/N) = 0$, and so on, for N steps until the solution $g(x,1) = 0$ is determined with what is (hopefully) a well-informed initial guess.

Variational annealing exploits a similar notion, that the Gaussian error action with weakly enforced dynamics is not only trivial but nearly degenerate; it corresponds to homotopy continuation with β taking the role of γ. The motivation is that while the lowest minima demanded by Laplace's Method in Eq. 2.36 may exist in extremely narrow basins of attraction (Figure 2.7d), the degenerate solutions at low R_f may first place at least some components of the estimates in their vicinity, before the appearance of the many spurious minima at higher R_f. The connection between variational annealing and homotopy continuation is simply being mentioned here as a matter of theoretical interest, but suggests some fruitful directions for future work.

Table 3.1: Algorithmic parameters for variational annealing of 10-*D* Lorenz 1996 model.

α	2.0	β_{max}	30
R_{f0}	1e-4	L	{1,...,5}
Q	10,000	$\sigma^2 = 1/R_{\text{m}}$	1.0
Δt	0.01	N	200

3.2 Inference in chaotic systems with variational annealing

3.2.1 Annealing plots

Fully visualizing the $A(\mathbf{X}|\mathbf{Y})$ manifold in $ND+P$ dimensions is impossible. One place we can begin is a plot of the estimated minimum values through the annealing progression – $A^{q,\beta}$ versus β – with an eye toward the dispersion of $\hat{A}^{q,\beta}$ as a function of experimental considerations such as observability (L) and noise ($\sigma = 1/\sqrt{R_m}$). Specifically, In these “annealing plots,” we expect that for $\beta \approx$ small, all of the $\hat{A}^{q,\beta}$ are approximately equal, this degeneracy lifting as β increases and the model dynamics are enforced. With sufficient β the estimates will begin to settle into distinct minima, with the hope that many converge to the global minimum corresponding to the true path. Of course, how many do so is a strong function of the number and character of the observed datapoints. In these plots, a particular progression of the estimate for given q will be referred to generically as a “run.”

3.2.2 Reliable estimation of sparsely observed Lorenz 1996 model

Let us use these plots to investigate the performance of variational annealing for the Lorenz 1996 model in ten dimensions (Eq. 2.29), given the model and algorithmic conditions listed in Table 3.1. We saw in Figures 2.5 that the surface of $\mathbf{A}(\mathbf{X}|\mathbf{Y})$ with these chaotic dynamics is riddled with local minima, and our hope is that a large proportion of distinctly initialized runs will track the global minimum corresponding true path despite this fractal character.

We first generate synthetic observations as before by integrating the Lorenz model equations (Eq. 2.29) with a 4th-order Runge-Kutta integrator and infusing the resulting trace with additive normally distributed noise of variance $\sigma^2 = 1.0$. These data are used as $\mathbf{Y}$ in the measurement term in of $A(\mathbf{X}|\mathbf{Y})$, Eq. 2.44, with various numbers of observed components from one to five to compare the algorithm performance as a function of observability. Traces of some of these data are shown in Figure 3.1 to give an idea of the noise

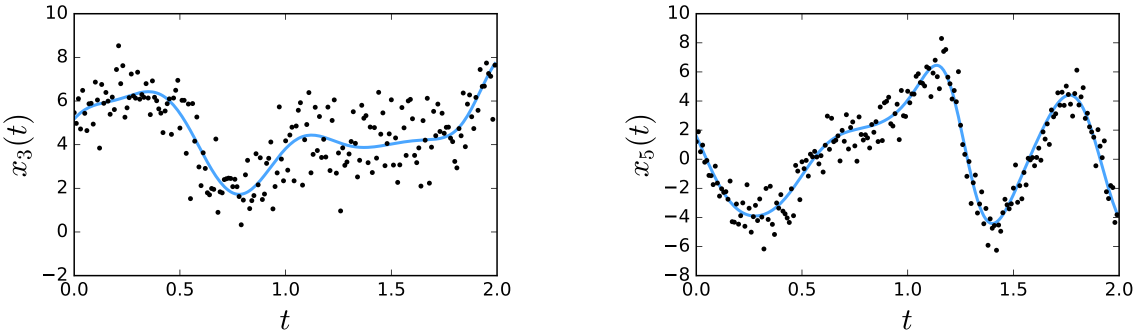

Figure 3.1: Observations (black dots) and true trajectory (blue solid curve) for two variables in the variational annealing numerical experiments on the Lorenz 1996 chaotic attractor.

deviation compared to the dynamical range of the model.

The measurement term of $A(\mathbf{X}|\mathbf{Y})$ is derived from an implicit trapezoidal integrator of the continuous dynamics, whereby $g_d(\cdot)$ in Eq. 2.44 is expressed as

$$g_d(\{\mathbf{x}_m\},\{t_m\},\Delta t) = x_d(t_n) + \frac{\Delta t}{2}[f_d(\mathbf{x}(t_n)) + f_d(\mathbf{x}(t_{n+1}))]. \tag{3.7}$$

For each $q \in Q$, the annealing is initialized by uniformly sampling each component of $\mathbf{X}$ within [-15, 15], which approximately spans the dynamical extent of the strange attractor. Variational annealing is then carried out for each q as described in Section 3.1 and Algorithm 1, where the optimizations themselves are performed with the interior point constrained optimization routine IPOPT [Wächter and Biegler, 2006]. Constraints are enforced to maintain the variables within their given range; it appears that none of the variables hit the bounds, so the constraints are inactive anyhow. In this case, IPOPT effectively performs Newton's method, for which we provide analytical expressions for both first and second derivatives of $A(\mathbf{X}|\mathbf{Y})$ explicitly, so no finite difference approximations are needed. Each run completes in around 10 seconds on an Intel i7 quad core processor.

Figure 3.2 shows the annealing plots for the five cases $L = \{1,...,5\}$. Often, distinct runs will return equivalent $\hat{A}^{q,\beta}$ values, so many of the 10,000 lines may be overlaid. To highlight this, the opacity of the blue lines roughly reflects the multiplicity of a particular run. First observe the degeneracy alluded to earlier for model inverse covariance. Though all 10,000 runs produce the same $\hat{A}^{q,\beta}$ for $\beta \lessapprox 15$, the associated estimates $\hat{\mathbf{X}}^{q,\beta}$ lie in *distinct* minima of the $A(\mathbf{X}|\mathbf{Y})$ manifold. They are truly degenerate, though only in the unmeasured components – this is illustrated explicitly in Figure 3.3a: for $\beta = 10$, the 10,000 estimates

Table 3.2: Number of unique minima $A^{q,\beta_{\max}}$ among 10,000 initializations in the estimation of the chaotic Lorenz 1996 model with sparse observations, using direct optimization versus variational annealing.

$L = ..$	1	2	3	4	5
Direct optimization	4199	1927	1647	1031	868
Variational annealing	124	40	16	3	2

Table 3.3: Likelihood of finding the global minimum of $A^{q,\beta_{\max}}$ among 10,000 initializations in the estimation of the chaotic Lorenz 1996 model with sparse observations, using direct optimization versus variational annealing.

$L = ..$	1	2	3	4	5
Direct optimization	< 0.1	2.04	4.90	11.91	11.46
Variational annealing	< 0.1	9.08	68.12	97.35	99.99

project to a single value along one of the the measured components $x_1(1.0)$, but scatter quite broadly along an unmeasured $x_5(1.0)$. Compare this to a projection along two unmeasured variables, which is dispersed broadly along both axes (Figure 3.3b). At sufficient R_f, the degeneracy begins to break and the measured components begin to veer from the central line and settle in minima $\hat{A}^{q,\beta}$, now of differing values; along the unmeasured axes, further movement occurs (Figure 3.4).

Crucially, this level splitting depends heavily on system observability. We saw in Figure 2.5 that the $A(\mathbf{X}|\mathbf{Y})$ manifold at high R_f is densely occupied by local minima, even when half of the variables are measured. Yet the minima produced by variational annealing with the same observations are a *tiny subset of these* – indeed only the lowest among them. This is seen quite explicitly by comparing the density of returned optimizations in Figs. 3.4b to those in Figure 2.5. For $L = 1$, the number of distinct minima returned by variational annealing is about 100, dropping to less than 20 for $L = 3$-5 (Table 3.2). In the direct optimization, more than 800 distinct estimates were returned even for 50% observability.

More telling are the number of minima that lie at the expected global minimum for $R_\text{f} \to \infty$. To derive this global minimum value, note that in the $R_\text{f} \to \infty$ limit, the model term dominates $A(\mathbf{X}|\mathbf{Y})$, so $\hat{\mathbf{X}}^{q,\beta}$ nearly mimics the true trajectory. Thus the model term in $A(\mathbf{X}|\mathbf{Y})$ is near zero, while the measurement term

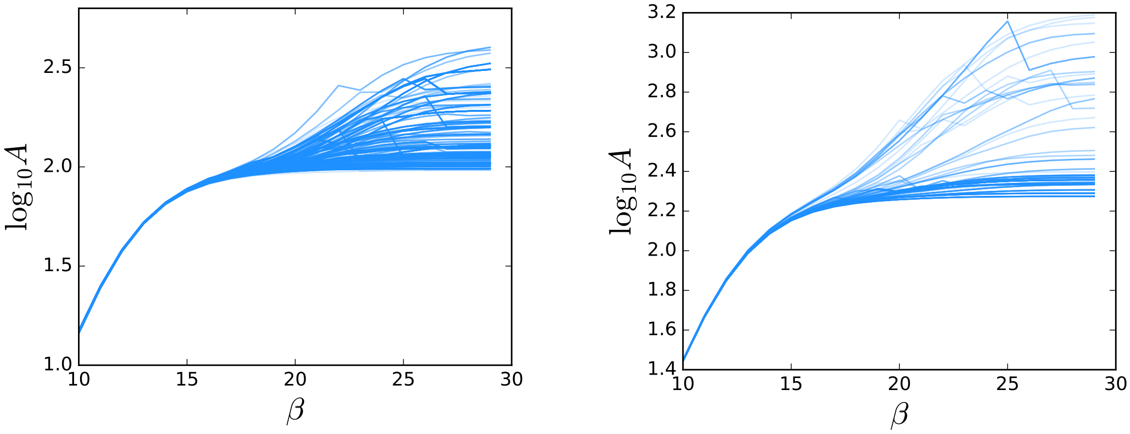

(a) $L = 1$

(b) $L = 2$

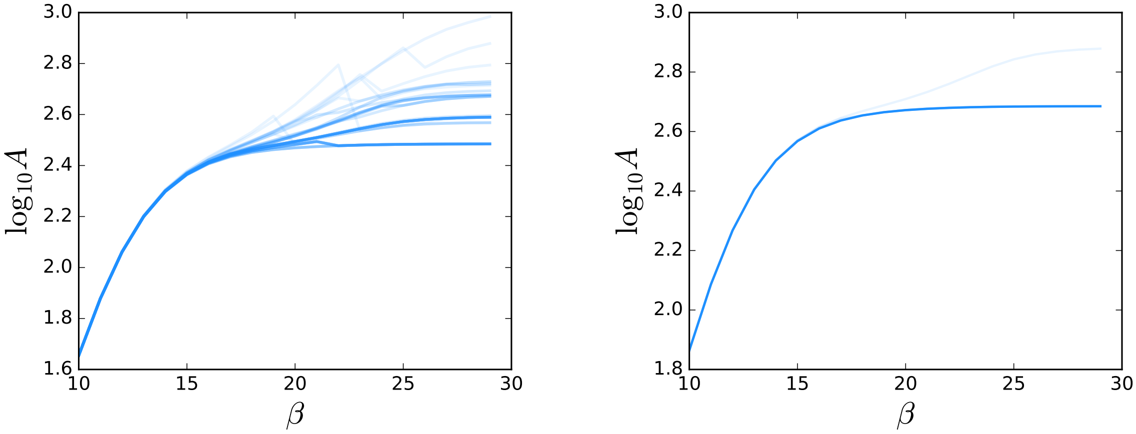

(c) $L = 3$

(d) $L = 5$

Figure 3.2: The annealing plots for the 1e4 parallel estimations of the Lorenz 1996 attractor using variational annealing, for different levels of observability. The number of distinct minima reduces a more variables are observed, but even for low observability, this number is far lower than what was found in direct optimization. For sufficient observability, the global minimum is found with near 100% reliability.

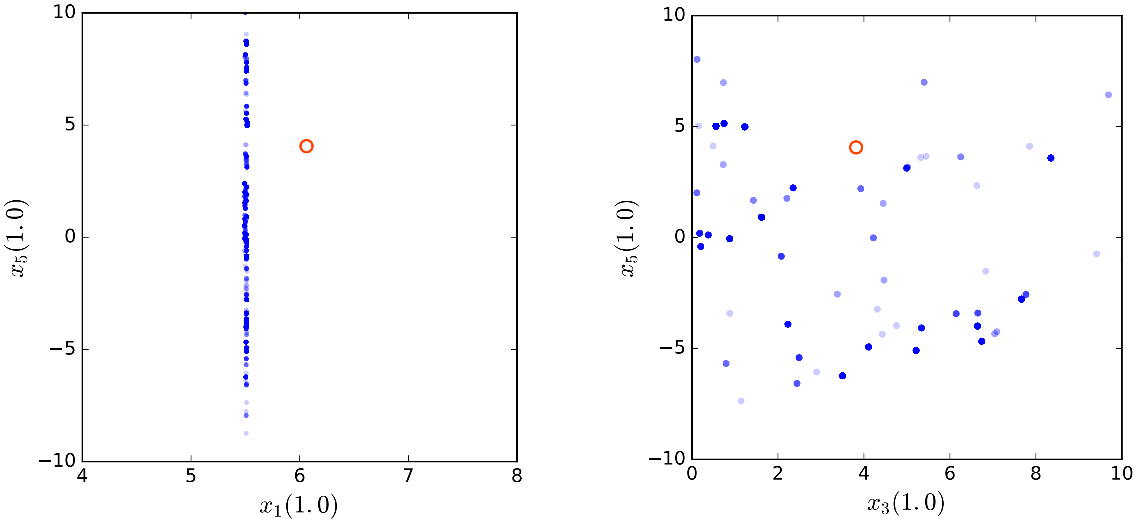

(a) (b)

Figure 3.3: (a) A $2D$ projection of the minima found with variational annealing for low R_{f} ($\beta = 10$) shows that these minima scatter broadly along the unmeasured component $x_5(1.0)$ but are tightly confined in the measured components $x_1(1.0)$. (b) This broad scattering is again seen for two unmeasured components.

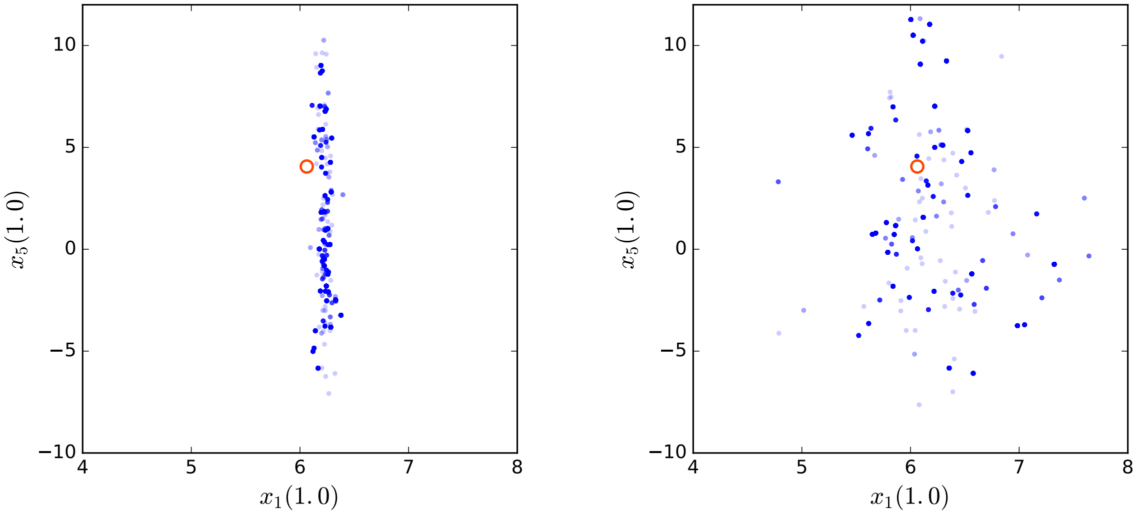

(a) (b)

Figure 3.4: (a) Projecting the estimate of the Lorenz 1996 model along the same axes as in Figure 3.3a, now with higher model precision $\beta = 17$, shows the movement of the estimates along both components as the degeneracy lifts. (b) Same plot, now with $\beta = 29$, shows further movement.

produces an expected error of:

$$E = \left\langle \sum_{n=1}^{N} \sum_{d=1}^{L} \frac{R_{\mathrm{m}}(n,d)}{2} (y_d(t_n) - x_d(t_n))^2 \right\rangle \tag{3.8}$$

$$= \sum_{n=1}^{N} \sum_{d=1}^{L} \frac{R_{\mathrm{m}}(n,d)}{2} \sigma^2 \tag{3.9}$$

$$= \frac{NL}{2}, \tag{3.10}$$

which for this example equals $\{2.00, 2.30, 2.48, 2.60, 2.70\}$ for $L = \{1, .., .5\}$ in log base 10. The percentage of the $Q = 10,000$ runs that achieve this value at $\beta_{\max}$ is shown in Table 3.3, and indicated qualitatively in Fig. 3.2, where the lowest level appears quite dark for $L \geq 3$. The advantages of variational annealing over direct optimization are evident: remarkably, virtually all of the variational annealing runs find the global minimum when $L > 3$, a feat accomplished by barely 10% for direction optimization.

In light of Laplace's method, the global minimum is critical in that it dominates the expected value integral Eq. 2.35 – if it is well separated from the local minima lying above. Consider for example $L = 3$, where the lowest two levels seen in Fig. 3.2c are $\hat{A}^{q_1,\beta_{\max}} = 303.8$ and $\hat{A}^{q_2,\beta_{\max}} = 368.014$. When exponentiated, the contribution of $\hat{A}^{q_1,\beta_{\max}}$ is more than 25 orders of magnitude smaller than that from the global minimum, completely eliminating its relevance to $\langle G(\mathbf{X})|\mathbf{Y}\rangle$. All subsequent examples illustrating the power and versatility of variational annealing will thus center on our ability to find this path.

3.2.3 Estimation quality of globally minimizing paths

All these results notwithstanding, the goal of the estimation problem is inferring the unknowns: How well do the estimated trajectories compare with the true ones? In Figure 3.5 is plotted the true and estimated trajectory $\hat{\mathbf{X}}^{q_1,\beta_{\max}}$ corresponding to the lowest level $\hat{A}^{q_1,\beta_{\max}}$. As indicated above, this estimate dominates the expected integral when $A^{q_1,\beta_{\max}}$ is sufficiently separated from the rest – this is true by at least four orders of magnitude for $L = 2, ..., 5$ and true to within 1% for $L = 1$. In this case, $\hat{\mathbf{X}}^{q_1,\beta_{\max}}$ will alone represent the expected value of the path, $\langle \mathbf{X}|\mathbf{Y}\rangle$. The estimates are poor when a single variable is measured (aside from the measured component $\mathbf{x}_1(t)$), but are quite accurate with only 20% observability, and excellent beyond. This holds true not only for the observed components, but, importantly, the unobserved components as well.

It is the accurate determination of these hidden states that is the true test of variational annealing. We are virtually agnostic about their trajectories and can at best choose them uniformly within a broad dynamic

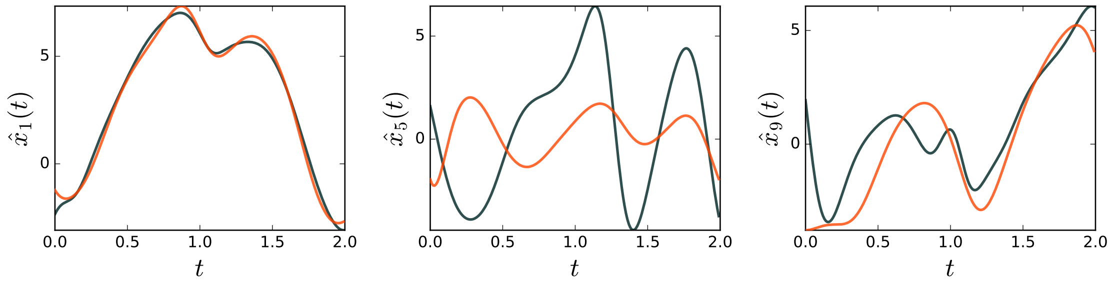

(a) $L = 1$

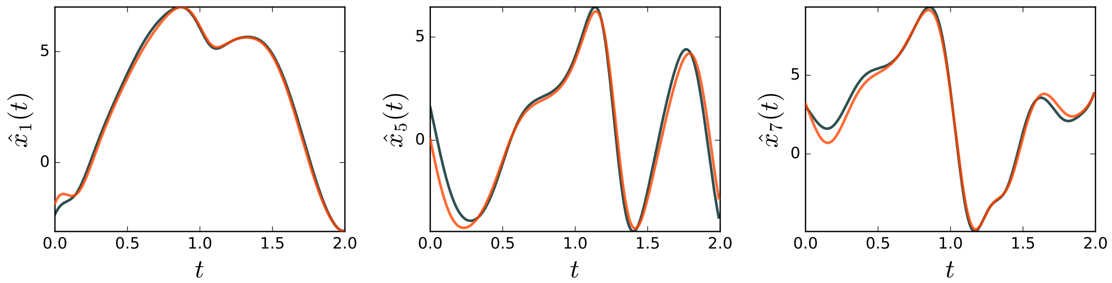

(b) $L = 2$

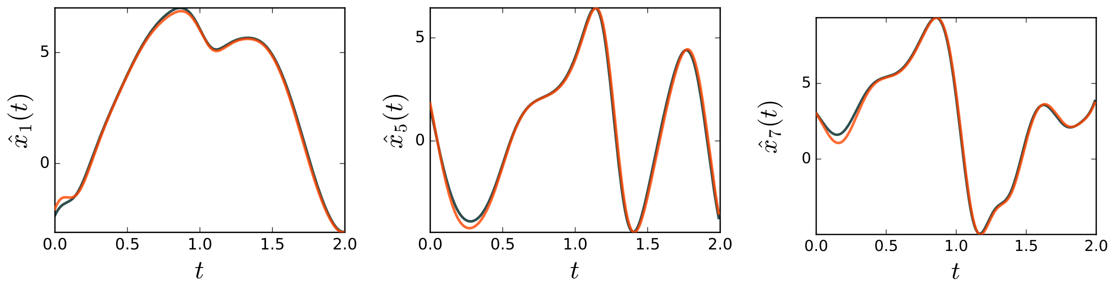

(c) $L = 3$

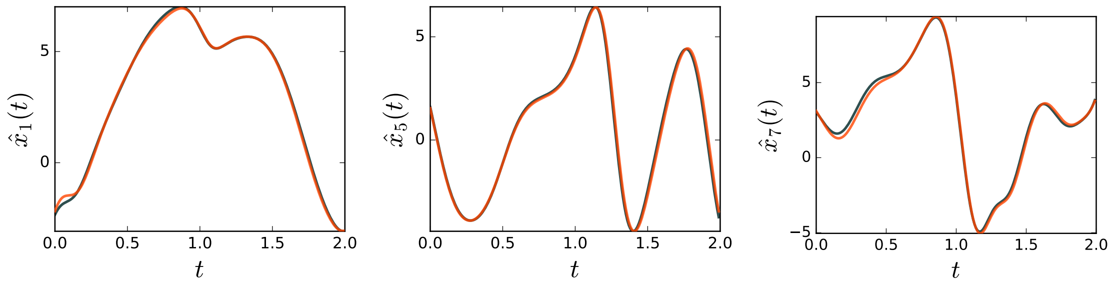

(d) $L = 5$

Figure 3.5: Optimal estimated paths (orange) versus true paths (black) in the Lorenz 1996 system using variational annealing, for various numbers of observed components from one to five.

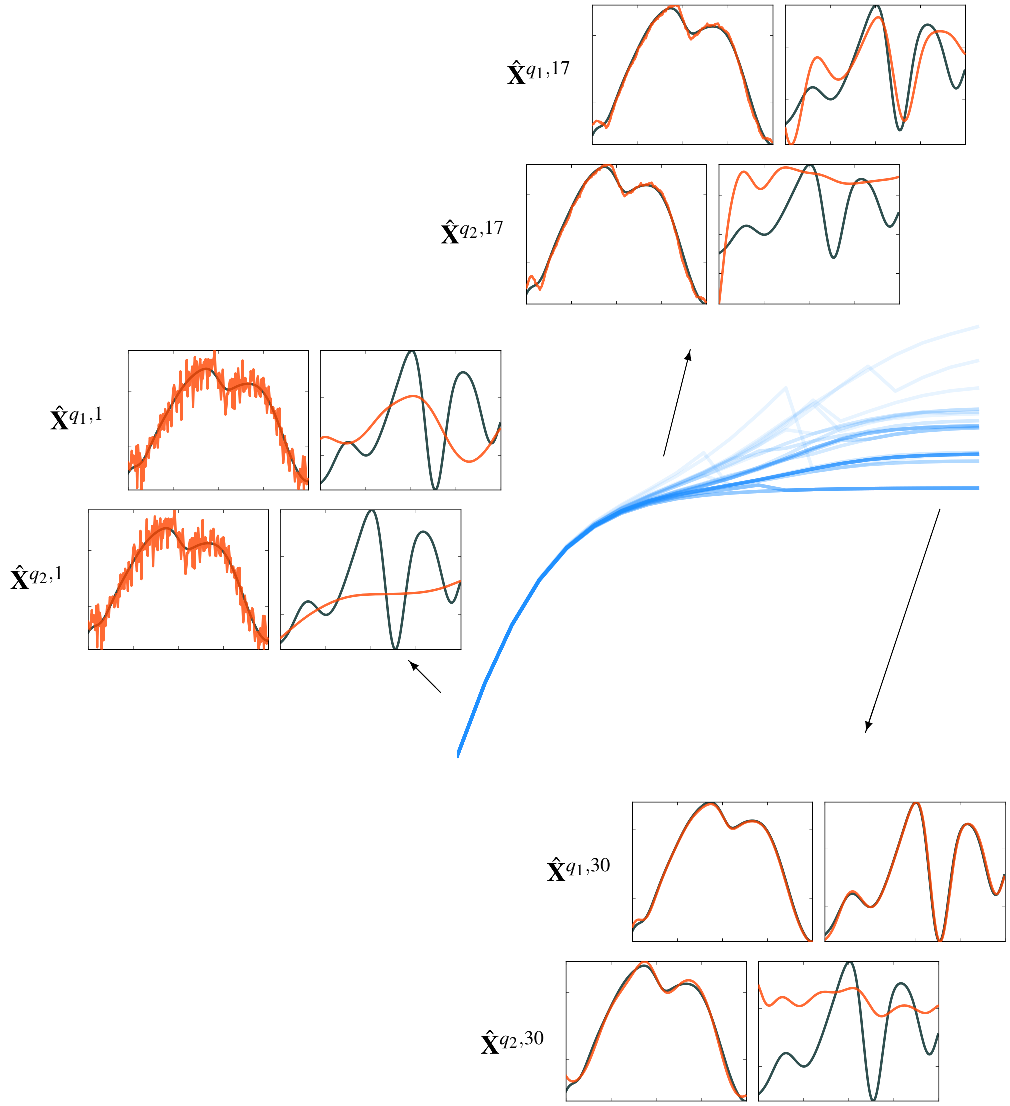

Figure 3.6: Progression of an optimal estimate (top row in each set) and suboptimal estimate (bottom row) as variational annealing progresses, plotted on the corresponding annealing plot for $L = 3$. Each row of plots shows one observed variable (left) and one unobserved variable (right). As β increases, the enforcement of model precision smooths the noisy trajectories found at low β, but even though the two A values of these estimates are fairly close at high β, only the lowest one produces a truly accurate estimate of the unobserved components.

range – yet nearly 7 in 10 attempts generates accurate estimations, even when observability is 30%; further, a slightly higher observability guarantees that we locate these hidden trajectories with near perfect reliability.

Further insight can be gained by plotting the estimates as the annealing progresses. Figure 3.6 shows estimated trajectories for one measured and one unmeasured variable, for two distinct runs at a few values of β. At low β, the estimates match the data closely in the measured components, but differ in the unmeasured variables. At higher β, two things occur: the measured components smooth as the model dynamics are enforced more stringently, and unmeasured components may converge toward the true states. At high β, the estimate corresponding to the lowest $\hat{A}^{q,\beta}$ is excellent, though other estimates (even those of slightly higher $\hat{A}^{q,\beta}$) may be quite poor. The lowest level $\hat{A}^{q_1,29}$ lies near the global minimum (Eq. 3.10), and so we expect the estimate to match the true dynamics *a priori*. Not only is it comforting to see this borne out by the time traces themselves, but also justifies the use of the more compact annealing plots as an informative metric for estimation quality.

3.3 The local vs. the global perspective

Finally, let us return to our primary motivation in moving beyond Kalman filters: the absence of an informative prior. The advantage of the filter was its computational simplicity via a systematic, gradual incorporation of observed information, timepoint by timepoint. The upshot is that it necessarily propagates forward the biases of the prior distribution, producing sensitivities to the presence of heavily non-Gaussian features of the full distribution. Further, low process noise $\mathbf{Q}_n$ compounds these sensitivities, demanding judicious, yet often heuristic, tuning.

In theory this might have been addressed with an ensemble method such as the particle filter, but as mentioned SMC suffers from untenable problems of dimensionality. With this in mind, variational annealing "works" for several reasons. First, even for mildly appreciable R_f, the global minimum of $A(\mathbf{X}|\mathbf{Y})$ is an excellent approximation to the full distribution – it reflects estimates that are accurate by *by orders of magnitude*. This means i) the majority of the conditional distribution modes are a distraction and ii) it is absolutely critical that these spurious modes are avoided. Variational annealing exploits these facts by working entirely within a variational context which deforms the cost manifold in a systematic way. This is not possible in a filtering context.

Second, it addresses the issue of uninformative priors in a different spirit. To be sure, variational annealing still relies on prior information: while in the filtering context, the prior is the actual estimate distribution at the initial time – $p(\mathbf{x}_0))$ – in variational annealing, this information is effectively embedded as the distribution of the optimization initialization at the first stage of the anneal, $\mathbf{X}^0_{\text{init}}$. Since R_f is so small, the prior is broad in unmeasured variables and nearly delta peaked in the measured ones. But unlike in filtering, the *a priori* information is encoded in the *simultaneous* trajectory at all timepoints n, the apparent message being that tight initialization along the fully collocated trajectory of observed states (a "temporally global perspective") is more robust than a broad initialization of the states at a single time, $\mathbf{x}_0$, propagated sequentially in time – the "local perspective."

Third, there is the as yet unaddressed issue of computational complexity. Newton's method requires the storage and manipulation (e.g. inversion) of Hessian matrices, which can be unwieldy with large dimension. Fortunately, the Hessian $\partial^2_{\mathbf{X}} A(\mathbf{X}|\mathbf{Y})$ is block-banded, since only the collocated state variables at neighboring timepoints are coupled together. The relevant manipulations can therefore be accelerated considerably by using dedicated algorithms for sparse matrices [Pissanetzky, 1984]. This is crucial in that the block-diagonal Hessian structure tames the complexity as a function of problem dimension and makes the variational annealing procedure at all viable.

Finally, though not relevant to the example just shown, we can see that parameters in variational annealing lie essentially on equal footing as states – the iterative framework having been dispensed with, parameter estimation in variational annealing needs no artificial dynamics such as those alluded to in Section 2.4.1. While parameter estimation will be explored in considerable detail in Chapter 4, it will be helpful to motivate this particular advantage of variational annealing with an illustrative example of a simple spiking neural model.

3.4 Parameter estimation in variational annealing

The sensitivities of sequential filtering methods to their prior distributions are particularly important in parameter estimation: oftentimes, more is known *a priori* about dynamical trajectories of biophysical systems than the static quantities. Consider the dynamics of a single neuron, the details of which will be discussed at length in Chapter 4. Ion channel conductances, which may vary over five orders of magnitude even for a single class of neurons, could nonetheless produce spiking behavior in which the membrane voltage is bounded tightly between -100 and 50 mV, and the gating variables between 0 and 1 [Dayan and Abbott, 2005, Ermentrout and Terman, 2010]. Further, dynamical bounds for measured variables are constrained by the known range of the observations themselves. Parameter ranges are considerably more lax, compounding the problem of tracking spurious distribution modes.

In variational annealing, parameters p_i and (collocated) states $x_a(t_n)$ appear on the same footing. Parameters can be treated as they really are – constants *a priori* – at the cost of augmenting the dimensionality of the optimization manifold from *ND* to *ND*+*P*. Of course, these extra dimensions will further complicate the manifold of $A(\mathbf{X}|\mathbf{Y})$, likely inducing further local extrema concealing the global minimum and possibly compromising the efficacy of variational annealing. As a first step then, let us carry out an illustrative comparison of the accuracy of variational annealing against a nonlinear Kalman filter in the 2-variable Fitzhugh-Nagumo (FN) model [1] [FitzHugh, 1961].

3.4.1 Two-variable Fitzhugh-Nagumo neuron model

The Fitzhugh-Nagumo model is

$$\frac{dV}{dt} = f_V(V,w) = V(a-V)(V-1) - w + I(t) \tag{3.11}$$

$$\frac{dw}{dt} = f_w(V,w) = bV - cw \tag{3.12}$$

which is parameterized by a, b, c, in addition to the “injected current” $I(t) = I$, which we assume to be constant. V is akin to a membrane voltage, which we assume to be only measured component, with observation noise $\sigma_V = 0.01$. (Further details of the HH model, upon which the Fitzhugh-Nagumo model is based, is described in Chapter 4).

Estimation with unscented Kalman filter

Let us estimate FN model using both the unscented Kalman smoother and variational annealing. To generate the true trajectory and synthetic data, we fix the parameters to the “true” values listed in Table 3.4 and integrate forward the dynamical equations Eq. 3.12 for some time. To apply the unscented Kalman filter, we promote the parameters to states with trivial dynamics $dp_i/dt = 0$ plus associated process noise σ_{p_i}

---

[1] The FN model of spiking neurons was developed in 1961 by R. Fitzhugh as a simplification of the original Hodgkin-Huxley dynamics; an equivalent electrical circuit model was developed a year later [Nagumo et al., 1962]

Table 3.4: Algorithmic parameters for unscented Kalman filter applied to the Fitzhugh-Nagumo model.

I	0.1	a	0.1	b	0.01	c	0.02	σ_V	0.001
σ_I	0.001	σ_a	0.001	σ_b	0.001	σ_c	0.001	σ_w	0.001

(chosen over several test runs to be as small as possible without causing filter divergence), listed in Table 3.4. The state variables $V(t)$ and $w(t)$ are assumed to have process noise $\sigma_V = \sigma_w = 0.001$; the estimation was fairly insensitive to further reduction in these values. Finally, the state and parameter estimates were each discretized with a forward Euler step, which determines the updated means $\mathbf{m}_n^-$ in the filtering equations:

$$\mathbf{m}_n^- = \mathbf{f}(\mathbf{m}_{n-1}) \tag{3.13}$$

where the vector space here is the augmented 6-dimensional state space $[V, w, I, a, b, c]$.

The UKF was run for 2000 timesteps of size $\Delta t = 0.25$. Kalman filters require statistics for the Gaussian prior distributions at t_0, for which we set:

$$\begin{aligned} m_V(0) &= y_V(0) \\ m_w(0) &\sim \mathscr{U}[-3, 3] \end{aligned} \tag{3.14}$$

and

$$\begin{aligned} m_{p_i}(0) &\sim \mathscr{U}[p_i - \gamma, p_i + \gamma] \\ \gamma &= 0.01, 0.1, 1.0, 10.0. \end{aligned} \tag{3.15}$$

That is, the mean of the measured component $V(t)$ was set to at its measured value, the mean of $w(t)$ was chosen uniformly from $[-3, 3]$ (concomitant with its approximate dynamical range), and the mean of the parameter distributions were chosen uniformly from various ranges centered at their true value. Note that these are not the initial covariances of the initial estimate, $\mathbf{P}_0$, rather the initial *means*, treated as random variables. The initial covariance of the state and parameter distributions, $\mathbf{P}_0$, was assumed diagonal and chosen optimally with several test runs (it should be noted that in real numerical experiments, this may not always be possible).

Table 3.5: Parameter estimates of the Fitzhugh-Nagumo model using i) the unscented Kalman filter with various prior distribution strengths and ii) variational annealing with a very weak prior.

y	*I*	*a*	*b*	*c*
True value	0.100	0.100	0.010	0.020
$\gamma = 0.01$	0.102	0.100	0.010	0.021
$\gamma = 0.1$	0.150	0.101	0.010	0.019
$\gamma = 1.0$	-0.133	0.107	0.005	-0.010
$\gamma = 10.0$	-0.547	0.119	0.008	-0.005
Variational, $\gamma = 100$	0.098	0.100	0.998	0.020

For small-to-medium-sized misplacement of the parameter initializations, $\gamma = 0.01$ and 0.1, the estimates at the end of the observation window were quite accurate, as indicated in Table 3.5. On the other hand, when the prior parameter distribution is centered incorrectly, $\gamma = 1$ and 10, the estimates have degraded, some appreciably.

Estimation with variational annealing

The variational annealing estimation requires no statistics for $p(\mathbf{x}_0)$, rather initial guesses for the collocated state-parameter space in the first annealing optimization, $\mathbf{X}^0_{\text{init}}$ in addition to hard bounds during the optimization itself (at all steps). For the states V and w, the states were confined to their dynamical range, $[-3,3]$. Analogously to the UKF estimation, the parameter bounds were chosen as $[-\gamma,\gamma]$ for various γ. The initial guesses at the first anneal step (with $R_{\text{f0}} = 0.001$) were chosen uniformly within these bounds and the anneal was iterated for 30 steps in increments of $\alpha = 2$.

For a broad range of γ, even as high as $\gamma = 100$, the parameter estimates with variational annealing are excellent (Table 3.5). Moreover, the estimates are virtually unchanged beyond a certain value of the annealing parameter $R_{\text{f}} = Q^{-1}$, eliminating the need to tune the process noise for maximum performance. At least in this particular case, variational annealing is far more robust to the type of parameter uncertainties that commonly reflect the state of knowledge of real neural systems.

It was found that a lengthier filtering window can improve estimates for UKF; in particular, for $\gamma = 1.0$, increasing the number of filtering steps from 2000 to 40,000 brought the parameters to errors within the levels of the $\alpha = 0.1$ estimates. Still, this fails to occur for $\alpha = 10.0$; the estimates remain poor even for such long time windows, suggesting that the spurious minimum to which the estimate has settled may be fairly peaked and impervious to further incoming observations.

This model, though a fairly simple one, is the best-case scenario for Kalman filters: the parameters themselves enter the equations linearly and the model is relatively well-observed (50% of the variables are measured). Yet the example still suffices to illustrate that parameter estimation may arise more naturally in the variational context, which relies on a collocated cost function and does not require the introduction of artificial time dependencies. The nonconvexities in the cost function imbued by unknown parameters – particularly nonlinear ones – are well tamed, in contrast to the spurious modes that can trap the sequential filter.

For the neural estimation problem, this is the key point. Filters are often applied to scenarios in which the quantity of interest is the state vector at the final time in an estimation window. This is certainly the case in weather prediction, for example, where knowledge of past temperature and climate conditions is combined with 6-hour observation window and climate models to predict the next day's forecast – all the variables of interest are dynamical, coupled together via the primitive equations. In addition, meteorological systems are sparsely but indefinitely observed.

Data collection in neural systems differs markedly from the meteorological context in that i) observations are not continuous, but gathered in short bouts, ii) background terms are virtually agnostic, and iii) the key quantities of interest are not unobservable dynamical variables but static (or slowly-varying) parameters underlying network connectivity and critical timescales of activity. It is therefore appropriate that a short-window optimization procedure may be more amenable to learning the variables of interest in neural networks, where a few minutes of extracellular voltage traces or calcium fluorescence both contains a wealth of information and exhausts the limits of experimental feasibility.

Chapter 4

Nonlinear estimation in neural systems

Neural networks are among the most complex dynamical systems in nature. Even the behavior of a single neuron requires highly-detailed, nonlinear biophysical descriptions to faithfully describe the transformation of input stimuli to robust action potential propagation [Dayan and Abbott, 2005, Izhikevich, 2007, Ermentrout and Terman, 2010]. It was realized in the middle of the 20th century that a description of neurons as leaky capacitors with stochastically-driven ion "gates" could reproduce excitable, spiking behavior in a remarkably faithful and parsimonious way [Hodgkin and Huxley, 1952, Catterall et al., 2012]. It soon became clear that the transition from quiescence to spiking was a manifestation of a particular class of bifurcations, ubiquitous in nonlinear dynamical systems.

Bifurcations are topological changes in dynamical behavior via changes in parameterization [Strogatz, 2001]. The defining feature is that these change are qualitatively *discontinuous*, as opposed to gradual or predictable. To make this distinction, compare the following two examples:

- A *simple* harmonic oscillator is parameterized by its spring constant k. For any given k, the portrait in the x-v phase plane traces the same curve – an ellipse – merely of differing size and shape. Since there are no qualitative behavioral changes, the system does not undergo a bifurcation with respect to k.
- A *damped* harmonic oscillator is parameterized by both spring constant k and damping coefficient γ. For sufficient γ, the oscillator exponentially decays to its equilibrium; however if γ is sufficiently small, it will oscillate about its equilibrium with diminishing amplitude. These phase portraits are qualitatively distinct: the first is a curve that *connects to* the origin, while the second is a curve that *spirals toward*

the origin. The system therefore undergoes a bifurcation as a function of γ.

The second example, though a simple, linear system, illustrates the defining feature of bifurcations: continuous changes in parameters produce abrupt changes in system response. Nonlinear systems are even richer, exhibiting not only equilibria and periodic orbits, but limit cycles and chaos – indeed it is the bifurcation from a stable equilibrium to a limit cycle that accounts for the generation of action potentials in the spiking neuron.

The message here is that parameters in biophysical neural models fundamentally encode the breadth and qualitative character of neuron response; yet they can rarely be probed directly by experiment. So enters the role of nonlinear inference. In this section we will apply the tools described in the previous chapter, variational annealing, to the determination of neural network parameters in a biophysically-inspired model of neurons that underlie vocal processing in the avian brain. Let us begin with a discussion of the physical basis of the model neuron and an analysis of its dynamics via bifurcation theory.

4.1 The Hodgkin-Huxley neuron

The modern mathematical description of a point neuron as a nonlinear continuous dynamical system owes its development to the work of Alan Hodgkin and Andrew Huxley in 1952. This formulation, known as the Hodgkin-Huxley (HH) model, has been proven remarkably robust and predictive, to the extent that virtually every neuron characterized in the 70 years since can be described within its framework [Johnston and Wu, 1994].

Two success of the HH model was predicated on two breakthroughs: i) a description of the lipid bilayer membrane as a capacitative junction, and ii) a description of ionic motion through permissive ion channels as a nonlinear voltage-mediated decay process. The equivalent electric circuit for this process is one in which the capacitative membrane is in series with the resistive ion channels; by Kirchoff's junction rule, these currents must sum to zero:

$$\frac{dQ}{dt} = C_{\mathrm{m}} \frac{dV_{\mathrm{m}}}{dt} = -\sum_i I_i \tag{4.1}$$

where C_{m} is the membrance capacitance, V_{m} the membrane voltage, and I_i the current through ion channel i.

How do ionic currents drive changes in the membrane voltage? Ions permeate into and out of a cell under the action of two forces. On the one hand, large disparities in concentration across the membrane would produce a diffusive flux; on the other hand, ions are charged, and this flux would soon lead to an opposing electrical force. As we would expect, an equilibrium situation exists when the chemical gradient is balanced by the voltage gradient. The total flux and associated current derived from this collective electrochemical gradient is

$$J_{\mathrm{T}} = J_{\mathrm{diffusive}} + J_{\mathrm{electric}} = -\frac{\mu kT}{q}\frac{\partial [C]}{\partial x} - \mu z [C] \frac{\partial V_{\mathrm{m}}}{\partial x} \tag{4.2}$$

$$I_{\mathrm{T}} = -\mu z RT \frac{\partial [C]}{\partial x} - \mu z^2 F [C] \frac{\partial V_{\mathrm{m}}}{\partial x}, \tag{4.3}$$

where μ is the ion mobility, $[C]$ its concentration, and z is its valence. The first term arises from Fick's Law and Einstein's relation $D = \frac{\mu kT}{q}$; the second is the microscopic Ohm's Law. Equilibrium, $J_{\mathrm{T}} = 0$, results therefore when the voltage drop across the membrane is

$$V_{\mathrm{eq}} = -\frac{RT}{zF} \ln \frac{[C]_{\mathrm{in}}}{[C]_{\mathrm{out}}}. \tag{4.4}$$

In other words, for a single ion channel, when $V_{\mathrm{m}} = V_{\mathrm{eq}}$, Eq. 4.1 is identically zero and the neuron is at rest [1]. The consequence of Eq. 4.4 is that deviations from V_{eq} produce current flow, the magnitude of which depends on a (possibly voltage-dependent) conductance $g_i(V_{\mathrm{m}})$:

$$I_i = g_i(V_{\mathrm{m}}) \cdot (V_{\mathrm{m}} - V_{\mathrm{eq},i}). \tag{4.5}$$

Extending to many ion channels, each channel with associated equilibrium potential $V_{\mathrm{eq},i} \equiv E_i$:

$$C_{\mathrm{m}} \frac{dV}{dt} = -\sum_i g_i \cdot (V - E_i), \tag{4.6}$$

where here forward, $V_{\mathrm{m}} \equiv V$ for notational brevity and the voltage-dependence of g is understood.

The second core feature of the HH model is its specification of the voltage-dependent ion conductances g_i. The idea is to model the ion channel as a collection of "gates" which open and close stochastically,

---

[1] Eq. 4.4 is no doubt familiar to solid state physicists – it is analogous to the expression for the built-in voltage of a p-n junction as a function of electron or hole density. This is of course unsurprising – electrons and holes are also charged species that equilibriate under the combined action of electric and chemical potentials

with voltage-dependent rates, $\alpha(V) \equiv \alpha$ and $\beta(V) \equiv \beta$, respectively. By the mass action law, the rate at which the gates close is proportional to those that are currently open, with proportionality constant given by the voltage-dependent rate β, and vice versa. Denoting the fraction of open gates at time t by $n(t)$, this gives

$$\frac{dn}{dt} = \alpha(1-n) - \beta n = \frac{n_\infty(V) - n}{\tau(V)}, \tag{4.7}$$

where the steady state function n_∞ and decay rate $\tau(V)$ can be written in terms of α, β. The latter expression is more illuminating in that it indicates that, out of equilibrium, $n(t)$ exponentially decays to the baseline $n_\infty(V)$ with rate $\tau(V)^{-1}$. In their original paper, Hodgkin and Huxley determined this baseline function by fitting of experimental data. In fact, it is found that quite often, n_∞ is logistic, which suggests a situation in which the gates are "activated" with increasing (or decreasing) voltage.

Since the gating function $n(t)$ gives the percentage of open gates within the ion channel; the conductance can be expressed by multiplying this function by some maximum constant value, whereby

$$g_i = \bar{g}_i n^p(t), \tag{4.8}$$

where $\bar{g}_i$ is the maximal conductance when the gates are fully opened, and where p is an integer power – the conductance g_i is not constant but depends on time through the gating variable $n(t)$, which in turn depends on the state of $V(t)$. This power is used to indicate the number of independent components of the ion channel, *all* of which must be activated to allow current to fully flow. The original HH model, which modeled the giant squid axon, assumed the existence of potassium and sodium channels, along with a persistent, non-gated "leak" channel that serves more or less as a catch-all for other currents not explicitly modeled. For potassium, $p = 4$, while for the sodium channel, $p = 3$, but it includes also an *inactivating* gate $h(t)$ which must be closed for current to flow. The full dynamical model equations for the squid axon thus assume the form

$$\begin{aligned} C_\mathrm{m}\frac{dV}{dt} &= -\bar{g}_\mathrm{Na} m^3 h \cdot (V - E_\mathrm{Na}) - \bar{g}_\mathrm{K} n^4 \cdot (V - E_\mathrm{K}) - \bar{g}_\mathrm{L}(V - E_\mathrm{L}) + I_\mathrm{inj}(t) \\ \frac{di}{dt} &= \frac{i_\infty(V) - i}{\tau(V)} \end{aligned} \tag{4.9}$$

for gates $i = \{m, n, h\}$, where I_inj indicates additional injected currents. The steady state logistic function can

Table 4.1: Parameters of Hodgkin-Huxley model, Eqs. 4.9.

E_{K}	-80	E_{L}	-55	E_{Na}	50	mV
$\bar{g}_{\mathrm{K}}$	360	$\bar{g}_{\mathrm{L}}$	3	$\bar{g}_{\mathrm{Na}}$	1000	nS
θ_m	-40	θ_n	-55	θ_h	-62.2	mV
σ_m	10	σ_n	20	σ_h	-5	mV
$\tau_{m,0}$	0.05	$\tau_{n,0}$	1	$\tau_{h,0}$	1	ms
$\tau_{m,1}$	0.4	$\tau_{n,1}$	5	$\tau_{h,1}$	7.5	mV
$\tau_{m,2}$	0	$\tau_{n,2}$	0	$\tau_{h,2}$	0	mV
C_{m}	1					pF

be modeled many ways; for the remainder of this thesis we will use

$$i_\infty(V) = \frac{1}{2}\left(1 + \tanh\frac{V - \theta_i}{2\sigma_i}\right). \tag{4.10}$$

The fact that the inactivating gate h must be *closed* to open the sodium channel is indicated by a negative value for σ_h – it is a negative logistic function. The time constant functions $\tau_i(V)$ are generally more heuristic (or even assumed constant); the following model is appropriately versatile:

$$\tau_i(V) = \tau_{i,0} + \tau_{i,1}\left(1 - \tanh^2\frac{V - \theta_i}{2\sigma_i}\right) + \tau_{i,2}\left(1 + \tanh\frac{V - \theta_i}{2\sigma_i}\right). \tag{4.11}$$

The Hodgkin-Huxley neuron is four-dimensional and nonlinear, parameterized by various constants that enter the dynamical equations both linearly (E_i, $\bar{g}_i$, C_{m}) and nonlinearly (everything else). Key among these parameters is the injected current, if assumed constant: the variation of I_{inj} accounts for a critical bifurcation that allows the neuron to transition from a stable resting state to robust spiking [Ermentrout and Terman, 2010].

4.1.1 Hopf bifurcations in the HH model

To investigate the dynamical behavior of the model, let us fix all parameters except I_{inj}, assumed constant, to the values given in Table 4.1. Since the system is 4-dimensional, it is impossible to plot all variables simultaneously, but a phase portrait in the V-n-h manifold will show the salient dynamical behavior.

Figure 4.1 shows phase portraits of the system when I_{inj} is set first to 0 pA, and then to 70 pA. These trajectories were determined by integrating forward the HH equations, Eq. 4.9, for 1 second in timesteps

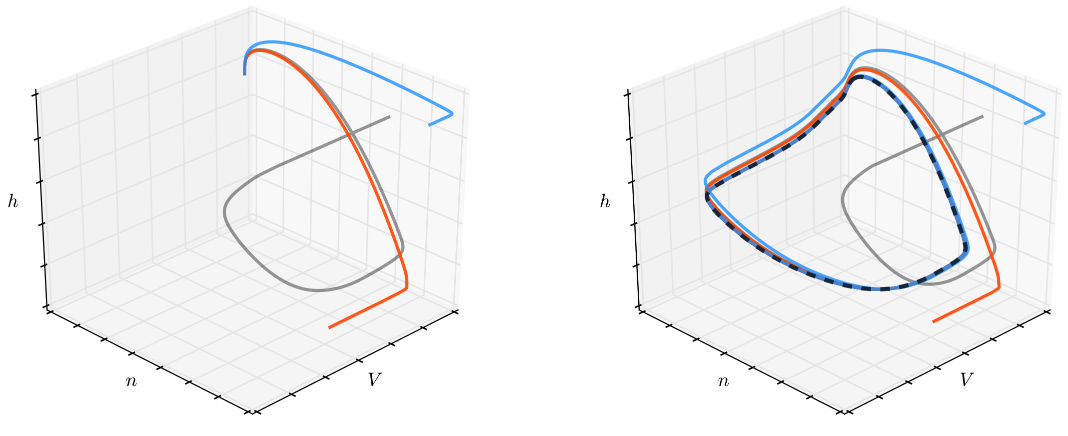

(a) (b)

Figure 4.1: (a) Phase portraits of the Hodgkin-Huxley neuron in the V-n-h manifold, using the parameterization listed in Table 4.1 and with 0 pA of injected current. The three trajectories, initiating at different points, all converge to the same fixed point. (b) Phase portraits of HH model with the same initial points as in (a), but now 70 pA of injected current. All three initializations converge to the same limit cycle (black dotted curve) representing tonic spiking.

of $\Delta t = 0.005$ ms, using a 4th order Runge-Kutta numerical integrator, and initialized from various random points in the approximate dynamical range. It is evident from Figure 4.1a that a stable fixed point with relatively large basin of attraction exists when the bifurcation parameter I_{inj} is set to 0 pA. On the other hand, we see in Figure 4.1b that a stable limit cycle, with large basin of attraction, occurs with $I_{\text{inj}} = 70$ pA. Since these behaviors are qualitatively quite distinct, a bifurcation must occur somewhere between these points. Specifically, stable node loses its stability via the creation of a stable limit cycle, which is known as a *Hopf bifurcation* [Strogatz, 2001,Ermentrout and Terman, 2010]. To see this in detail, a plethora of phase plots like those in Figure 4.1 would need to be examined and compared for various bifurcation parameters. However, the qualitative behavior can be visualized in a single plot known as a bifurcation diagram.

The bifurcation diagram for this HH system is shown in Figure 4.2. The x-axis is the bifurcation parameter, I_{inj} and on the y-axis is the membrane voltage V. Rather than plotting phase plots of the system, we only plot the limiting qualitative dynamical behavior: fixed points, (stable fixed points in blue, unstable ones in red) and limit cycles. The diagram plots the V-component of the fixed points (equilibria) that exist for various values of the bifurcation parameter. Some regions may also contain a stable limit cycle, whose maximum and minimum voltage values are indicated by in grey.

The fixed points were found by setting the HH dynamical equations, Eq. 4.9, to zero; this system of nonlinear algebraic equations was solved with hundreds of initial guesses (for each value of I_{inj} to find all possible such points. The fixed point stability was then found by calculating the eigenspectrum of Jacobian of the HH vector field; the system is locally unstable when the real part of at least one of these eigenvalues is positive.

We see in the bifurcation diagram shown in Figure 4.2a that the neuron spikes for steady injected currents between 56.2 and 599.0 pA. Outside this range, the system will decay to a stable resting voltage somewhere around -60 V. Further, the stability of these low voltage fixed points is lost for injected currents lying between 138.5 pA and 359.5. The combination of these fixed points and limit cycles therefore produces a feature that will we will meet again in our models of the HVC neuron: *bistable* regions in which stable limit cycles and stable fixed points coexist, between 56.2 and 138.5 pA and between 359.5 and 599.0 pA. Depending on the relative sizes of the attractive basins and the initial point, the neuron may either spike or rest for steady injected currents in these ranges. We will soon see that when a spiking neuron is combined with an additional slow subsystem, bistability can be exploited to produce an even richer repertoire of stimulus response.

A similar diagram is shown in Figure 4.2b, except with the gating parameter θ_h now set to the slightly different value of -60.7 mV. Though the diagram is qualitatively similar to that in (a), the regions of stability, spiking, and bistability have shifted. One may imagine that an opposing change in θ_h could annihilate the unstable fixed points altogether. One could imagine even further changes in behavior as any of the other two dozen parameters are adjusted. Even in this relatively small system of just four variables, the behavior is rich and quite varied.

4.2 Projection neurons in the zebra finch HVC

4.2.1 Modeling nonlinear biophysical neurons and the role of inference

Though the Hodgkin-Huxley model is a bare-bones description of a point neuron with no spatial extent or dendritic interaction, it has been enormously applicable in faithfully describing specific neural firing patterns throughout the animal kingdom [Koch and Segev, 1998, Dayan and Abbott, 2005, Izhikevich, 2007, Sterratt et al., 2011]. Further, it is modular and therefore amenable to straightforward extensions – one only need identify the existing ionic channels and add their corresponding functional forms to the sum of

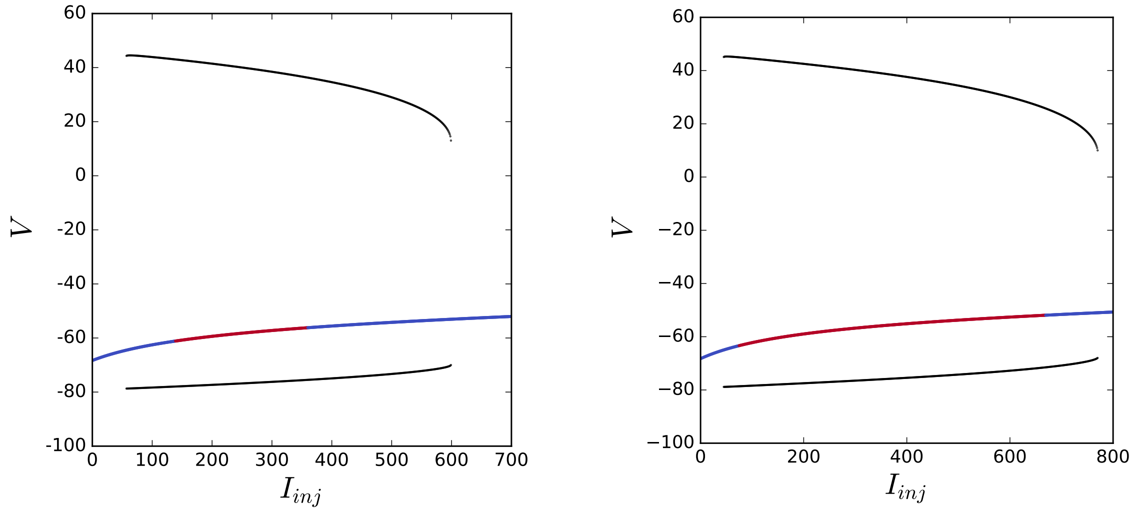

(a)

(b)

Figure 4.2: (a) Bifurcation diagram of the Hodgkin Huxley model with parameters listed in Table 4.1. Stable fixed points are shown in blue, unstable fixed points in red, and stable limit cycle bounds in black. (b) Similar bifurcation diagram of the HH neuron, except now with $\theta_h = -60.7$. The behavior of the system is qualitatively similar to an extent. The bistable regions, for example, have shrunk considerably, and further increases in θ_h would produce a vanishing of the stable regions entirely.

ionic currents in Eq. 4.1.

As an aside, a large lexicon of "reduced" HH-type models has been developed and classified, their behavior systematically understood via bifurcation theory [Izhikevich, 2007]. These models can individually reproduce many desired features of spiking models, and, though not emphasized in this thesis, provide alternate descriptions of reduced dimensionality with minimal compromises in qualitative behavior. We should keep in the back of our mind this that there indeed exists an array of options when choosing biophysical neural models for the purpose at hand.

In developing HH-models of single neurons, a useful starting point is the experimental identification of existing ionic channels using pharmaceutical agents. In practice, this information is binary: the physiologist may be able to identify that an interneuron, for example, contains several distinct Na and K channels in addition to calcium channels, but may not know how strongly these channels contribute to the overall current – for example, $\bar{g}_i$ – nor their kinetic details embodied by the gating parameters and timescales, θ_i, τ_i, etc. Yet the lesson of bifurcations is that these parameters are *absolutely central*, accounting for not only continuously varying quantities such as spiking frequency and spike amplitude, but qualitative distinctions such as the regimes that spiking even occurs at all.

The importance of parameterization in bifurcating nonlinear systems illustrates the power of neural modeling from the "inverse" viewpoint of nonlinear inference: i) known anatomical features are used to identify a general model description (say, a set of differential evolution equations), whose unknown parameters are then ii) found by passing electrophysiological data into the estimation algorithm. With this in mind, let us apply this idea to a particular class of neurons in the avian vocal pathway, by beginning first with what is known about the anatomy and behavior of this widely-studied neural system.

4.2.2 Zebra finch HVC

We will focus here on a particular avian species, *zebra finch*: a colorful songbird that expresses a short, stereotyped mating call throughout its adult lifetime, following a several month *sensorimotor* period of learning as a juvenile [Hahnloser et al., 2002, Fee et al., 2004]. Though the mechanics of this particular species are in no way requisite for the estimation models investigated in this thesis, the system is interesting in its own right as a template for language acquisition, syntax, learning, and synaptic plasticity in higher-level organisms [Brainard and Doupe, 2002]. Indeed, the male zebra finch only learns a single $\sim$1 second song, permitting repeated experimentation on the relevant neural mechanisms involved. This is contrast to other songbirds such as Bengalese finch, which exhibit a still limited but more variable vocal repertoire [Okanoya, 2004].

Vocalization in zebra finch involves the interplay of several distinct brain regions. Primary among these is the HVc (HVC), thought to initiate the string of neural events that ultimately lead to motor commands in the vocalbox [2] [Yu and Margoliash, 1996, Mooney and Prather, 2005]. While the possibility remains that HVC may in turn receive vocal initiation commands from other regions, *nucleus uvaeformis* (Uva) and *nucleus interfacialis of the nidopallium* (NIf), lesion studies suggest that they may play a subsidiary role, if any at all [Cardin et al., 2004, Coleman and Vu, 2005]. HVC in turn forges excitatory synaptic connections with – "projects" to – other regions of the brain that assume distinct roles. In particular, HVC contains projection neurons exciting those in the motor-command center *robust nucleus of arcopallium* (RA), known to be responsible for directly driving the motor neurons in the vocalbox [Mooney, 2000, Mooney and Prather, 2005]. These neurons are denoted as HVC RA-projecting neurons or $\mathrm{HVC_{RA}}$ neurons, and are the primary focus of this chapter.

---

[2] HVC originally stood for "High Vocal Center," a term which has been since deprecated.

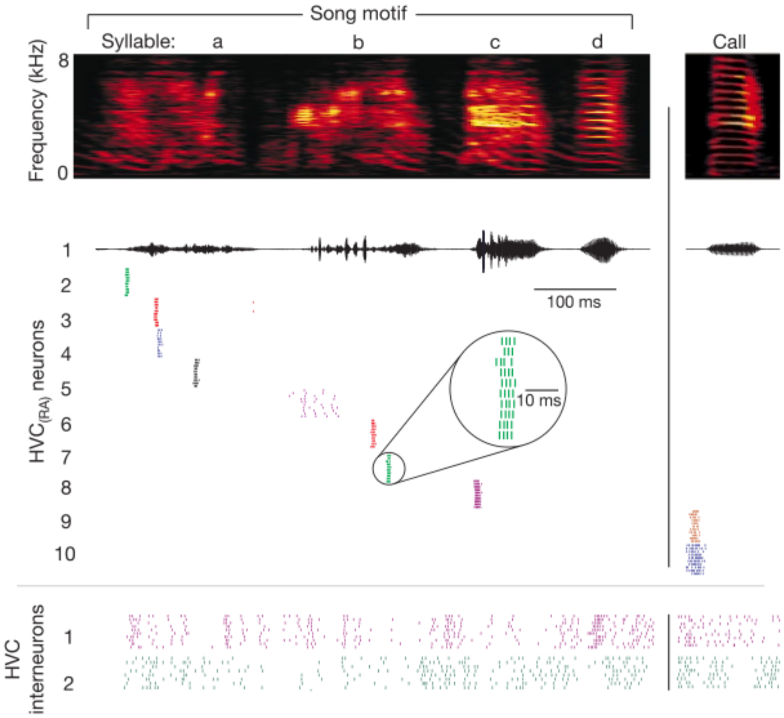

Figure 4.3: (top panel) An auditory spectrum of zebra finch mating call shows distinct syllables comprising a single motif. (middle panel) Spike rasters of nine different HVC_{RA} neurons through several renditions of the song. Each projection neuron bursts at most once, at a tightly time-locked instance in the song motif. (bottom panel) Spike rasters of two HVC inhibitory interneurons. Interneurons spike more broadly in the song, although they appear to exhibit some instances in which they are quiescent, perhaps pointing to the initiation of HVC_{RA} bursts from a release of inhibition [Gibb et al., 2009, Kosche et al., 2015]. [Reprinted by permission from Macmillan Publishers Ltd: *Nature* 419(6902):65–70, copyright 2002]

4.2.3 Sparse, robust bursting in HVC projection neurons

HVC_{RA} neurons have attracted a great deal of attention for the following reason: during a song motif, they exhibit a peculiar firing pattern consisting of only a single short 3-5 ms burst [Hahnloser et al., 2002]. This solitary burst of activity appears to occur repeatably and at the same instance of any given rendition of the motif (Figure 4.3). They are also highly stereotyped, exhibiting little jitter or variation in spike frequency or length across song renditions. That is, they appear to be time-locked somehow to the song itself, and there is an ongoing debate concerning whether these bursting patterns are in correspondence with particular *acoustic features* of the song, or rather to the exact *timing* at which these features occur [Long and Fee, 2008, Goldin et al., 2013, Kosche et al., 2015]. These competing perspectives suggest differing network topologies, from feed-forward synaptic chains to broader, less-directed connection webs [Jin et al., 2007, Gibb et al., 2009, Long et al., 2010, Kosche et al., 2015].

In contrast to RA-projecting neurons, inhibitory *interneurons* fire densely throughout the song, exhibiting a fair amount of variation by rendition, as shown in Figure 4.3 [Mooney and Prather, 2005]. It has been observed that they do exhibit brief pauses, possibly coincident with $\mathrm{HVC}_{\mathrm{RA}}$ bursting, and may therefore play some role in suppressing the projection neurons until it is "their turn" to fire during the song [Gibb et al., 2009, Kosche et al., 2015].

How the interplay of interneurons, projection neurons, and others lead to the sparse, robust pattern found in zebra finch requires a full understanding of both cellular and network properties of HVC circuitry. This is still an open problem, of which the remainder of this chapter will address the following: i) the construction of a comprehensive model of RA-projecting neurons and ii) a method for inferring the full parameterization of the model from current clamp data, *within the constraints of available electrophysiology*. This proposed ideas will generalize readily to the larger features of the system, such as the features of interneurons and connection topologies, and the intent being that a *combination* of systematic methods of inference and clever experimentation can eventually uncover the neural mechanisms responsible for vocalization.

4.3 Modeling the $\mathrm{HVC}_{\mathrm{RA}}$ neuron

4.3.1 Known anatomical features

A number of electrophysiological studies have elucidated the characteristics of various neuron classes present in HVC [Kubota and Taniguchi, 1998, Mooney and Prather, 2005, Long et al., 2010, Daou et al., 2013]. For RA-projecting neurons in particular, there is strong support for the presence of both calcium channels and calcium-mediated channels, such as calcium-gated potassium channels. In addition to these and the usual Na+ and K+ channels, there appear to be *A-type* currents, responsible for spiking delay in response to depolarizing currents.

As we will see, the combined action of calcium channels and calcium-mediated channels will together conspire to produce the robust type of short bursting observed in $\mathrm{HVC}_{\mathrm{RA}}$ neurons. For the purposes of this thesis, we will ignore the action of A-type currents, which do not affect the bursting behavior. We will first begin with a model for calcium and calcium-mediated channels, drawing again on the total channel flux, Eqs. 4.2 and 4.3.

4.3.2 Calcium channels and the GHK equation

Calcium channels require a slightly different description than the conductance-gating description of the Hodgkin-Huxley model, one that respects the vast differences in Ca concentration on either side of the neuron cell body membrane. Consider again the ion current equation, derived by summing the electric and chemical forces acting upon the charged ion species of valence z and mobility μ:

$$I_{\mathrm{T}} = -\mu z R T \frac{\partial [C]}{\partial x} - \mu z^2 F [C] \frac{\partial V}{\partial x}. \tag{4.3 revisited}$$

Recall that the expressions for ion current, $I = \bar{g} n^p (V - V_{\mathrm{eq}})$ were derived by using Eq. 4.3 to determine the equilibrium voltage V_{eq} and assuming that current flows under the combined effect of an Ohmic dependence on $(V - V_{\mathrm{eq}})$ and additional non-Ohmic gates n. Disregarding the action of the gate for a moment, how appropriate is the Ohmic assumption? Certainly it is true when $V \approx V_{\mathrm{eq}}$, but this may be easily violated during an action potential.

It turns out that the assumption is valid when the disparity in concentrations on either side of the membrane, $[C_{\mathrm{out}}]$ and $[C_{\mathrm{in}}]$, is moderate. To see this, let us revisit the channel current, but now model the membrane more realistically as having a finite width l, and include the simplifying assumption that the electric field within the membrane is constant. The latter assumption removes $V(x)$ in favor of only the voltage drop across the membrane $\Delta V \equiv V$:

$$I_{\mathrm{T}} = -\mu z R T \frac{d[C]}{dx} - \mu z^2 F [C] \frac{V}{l}, \tag{4.12}$$

assuming a steady state. This is a first order boundary value equation, where we enforce that the ion concentration on either side of the cell membrane equals the intracellular and extracellular concentrations:

$$\begin{aligned} [C](0) &= [C]_{\mathrm{in}} \\ [C](l) &= [C]_{\mathrm{out}} \end{aligned} \tag{4.13}$$

A solution to the differential boundary problem is

$$I_{\mathrm{T,CF}}(V) = -g_{\mathrm{GHK}} V \frac{[C]_{\mathrm{out}} e^{-V/V_{\mathrm{T}}} - [C]_{\mathrm{in}}}{e^{-V/V_{\mathrm{T}}} - 1}, \tag{4.14}$$

Table 4.2: Constant field (GHK) model parameters for various ions.

	[Na^+]	[K^+]	[Ca^{2+}]
$[C]_{\mathrm{out}}$ (mM)	145	5	2.5
$[C]_{\mathrm{in}}$ (mM)	15	140	1e-4
V_{T} (mV)	25	25	13
V_{eq} (mV)	55	-83	130

where CF denotes the "constant field" approximation, g_{GHK} is composed of various constants, and

$$V_{\mathrm{T}} = \frac{RT}{zF} \tag{4.15}$$

is the "thermal voltage" in analogy to the carrier activation voltage in semiconductors. The notation "GHK" signifies Goldman, Hodgkin, and Katz, after whom this equation is named [Hille, 2001, Ermentrout and Terman, 2010].

The GHK current-voltage relationship is clearly not ohmic. To gain some insight into the relevance of the linear assumption, consider the *I*-*V* characteristics for Na and K, plotted in Figure 4.4a, along with the associated ohmic approximation centered at V_{eq}; typical concentrations for these ions in mammalian species are given in Table 4.2. Since the concentrations of K and Na on either side of the neuron membrane differ by at most 10-30 times, the ohmic approximation remains reasonable in the range of physical membrane voltages, between -80 and 30 mV.

Conversely, calcium concentrations on either side of the neuron membrane vary by as many as four or five orders of magnitude. This pushes the equilibrium voltage up above 100 mV, and the strong rectification renders the Ohmic approximation completely invalid in the spiking regime – the constant field and ohmic expressions differ by several orders of magnitude (Figure 4.4b). For this reason, calcium currents are often modeled by the GHK expression when computationally feasible, with the clear disadvantages of added complexity, parameterization, and nonlinearities in the model description.

Pharmacological experiments have pointed to the presence of these calcium currents in $\mathrm{HVC}_{\mathrm{RA}}$ neurons, in particular a high-threshold "L-type" current which assumes a functional form composed of the

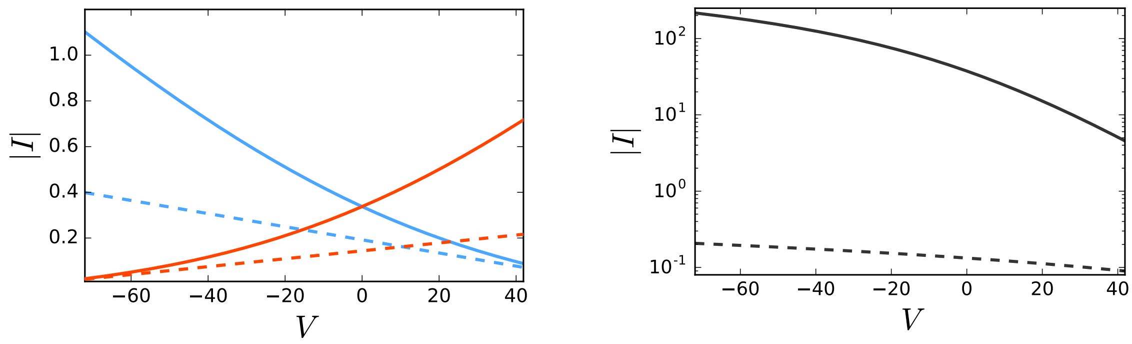

(a)

(b)

Figure 4.4: (a) Magnitude of ion channel current as a function of constant membrane voltage, using the gating-ohmic approximation of the channel current (Eqs. 4.5 and 4.8), shown by the dotted line for Na (blue) and K (orange). The currents given by the more exact GHK expression, Eq. 4.14, is shown by the corresponding solid lines. The gating-ohmic conductance expression is a fair approximation, even far from the equilibrium voltage. (b) A similar comparison for calcium channels. Since calcium concentrations differ markedly on either side of the cell membrane, the gating-ohmic and constant field expressions differ by several order of magnitudes.

maximal conductance indicated by the GHK equation and gating variable r [Daou et al., 2013]:

$$I_{\text{Ca-L}} = \bar{g}_{\text{Ca-L}} r^2 \Phi_{\text{GHK}}(V) \tag{4.16}$$

$$\Phi_{\text{GHK}}(V) = V \frac{[\text{Ca}]_{\text{out}} e^{-V/V_{\text{T}}} - [\text{Ca}]}{e^{-V/V_{\text{T}}} - 1} \tag{4.17}$$

$$\frac{dr}{dt} = \frac{r_\infty(V) - r}{\tau_r(V)}, \tag{4.18}$$

where $[\text{Ca}]_{\text{in}} \equiv [\text{Ca}]$. In modeling HVC_{RA} neurons, we will utilize the realistic concentration-dependent dynamics inherent in the GHK description, keeping in mind that if the parameterization proves unwieldy, this equation could be swapped in favor of a less realistic Hodgkin-Huxley-type gating model.

4.3.3 Calcium-gated potassium channels and calcium dynamics

Due to the large concentration of extracellular calcium relative to that inside the cell, the intracellular calcium can fluctuate substantially during neuron activity. The Ca influx in turn may drive other channels, which are activated by the presence of other ions and are highly selective to calcium [Ermentrout and Terman, 2010]. One such channel whose presence in HVC_{RA} has been suggested pharmacologically is the calcium-gated potassium channel, $I_{\text{K/Ca}}$, which is modeled as an ohmic channel, but with a gating variable that depends on intracellular calcium concentration $[\mathrm{Ca}]_{\mathrm{in}} \equiv [\mathrm{Ca}]$ [Daou et al., 2013]:

$$I_{\mathrm{K/Ca}} = \bar{g}_{\mathrm{K/Ca}} r_{\infty}([\mathrm{Ca}])(V - E_{\mathrm{K}}) \tag{4.19}$$

$$r_{\infty}([\mathrm{Ca}]) = \frac{[\mathrm{Ca}]^2}{[\mathrm{Ca}]^2 + k_s^2}. \tag{4.20}$$

The functional form of $r_{\infty}([\mathrm{Ca}])$ is known as a Hill function and is logistic, in analogy to the limiting voltage gating functions n_{∞}, r_{∞}, etc. Though not shown for notational brevity, the intracellluar calcium concentration [Ca] is explicitly time-dependent. What are its dynamics? The influx of calcium arises from the action of the L-type channel, while a return to baseline is maintained calcium buffers and pumps. Thus:

$$\frac{d[\mathrm{Ca}]}{dt} = \phi I_{\mathrm{Ca\text{-}L}} + \frac{[\mathrm{Ca}]_0 - [\mathrm{Ca}]}{\tau_{\mathrm{Ca}}}, \tag{4.21}$$

where $[\mathrm{Ca}]_0$ is an intracellular equilibrium background concentration. Of particular note is the value of the decay time constant, τ_{Ca} – it is between 20 and 100 ms, an order or two slower than the frequency of spiking. The consequence of these two disparate timescales will soon become apparent.

4.3.4 Multiple timescales and the origin of bursting

The behavior of $\mathrm{HVC_{RA}}$ during a song motif is defined by a single, short series of spikes, time-locked to multiple renditions of the song. One may imagine that these bursts arise from a sustained pulse of synaptic current, lasting for the duration of the burst, about 5 ms. But noise in synaptic currents would then introduce a variability in burst length and firing rate that does not otherwise exist – the bursts exhibit little variation in spike frequency and length over repeated song renditions.

An alternative viewpoint, in line both with this stereotypy and with the experimental evidence for the presence of dominant calcium channels, is that bursting arises from a cellular mechanism alone [Jin et al., 2007]. This occurs fundamentally via i) the interplay of disparate timescales and ii) the presence of coexisting stable states. To see this, let us view a bifurcation diagram of the model Hodgkin-Huxley neuron with usual Na, K, and leak channels, and in addition the calcium channel and calcium-gated potassium channels described by Eqs. 4.16 and 4.20. However, instead of retaining the calcium dynamics (Eq. 4.21), we will fix [Ca], treating it as a bifurcation parameter. The question then is, What are the possible behaviors of the system for a particular value of [Ca], with constant injected current?

Table 4.3: Parameters of Hodgkin-Huxley model neuron with traditional Na, K, and leak currents, in addition to Ca currents, Ca dynamics, and Ca-gated potassium currents.

E_K	-90	E_L	-80	E_Na	55	mV
$\bar{g}_\mathrm{K}$	120	$\bar{g}_\mathrm{L}$	3	$\bar{g}_\mathrm{Na}$	1050	nS
θ_m	-30	θ_n	-35	θ_h	-45	mV
σ_m	9.5	σ_n	1	σ_h	-7	mV
$\tau_{m,0}$	0.01	$\tau_{n,0}$	0.1	$\tau_{h,0}$	0.1	ms
$\tau_{m,1}$	0	$\tau_{n,1}$	0	$\tau_{h,1}$	0	mV
$\tau_{m,2}$	0	$\tau_{n,2}$	0.5	$\tau_{h,2}$	0.75	mV
$[\mathrm{Ca}]_\mathrm{out}$	2500	C_0	0.2	k_s	2.0	μM
θ_r	-40	σ_r	10			ms
C_m	21					pF
$\bar{g}_\mathrm{K/Ca}$	120					nS
ϕ	5e-5					μM/pA/ms
τ_Ca	60					ms
$\bar{g}_\mathrm{Ca\text{-}L}$	8e-3					nS/μM

The parameter values chosen for this illustration are shown in Table 4.3; the injected current is held steady at 200 pA. (Note that for this bifurcation study, $d[\mathrm{Ca}]/dt = 0$, so ϕ and τ_Ca are irrelevant). The corresponding bifurcation diagram for V as a function of the bifurcation parameter [Ca], in Figure 4.5, with stable fixed points plotted in blue and unstable equilibria in red. For low calcium concentrations, less than 0.4 or so, the only stable state is the spiking limit cycle. For a small region between around 0.42 and 0.50, a stable equilibrium at $V \approx -65$ mV coexists with this limit cycle, while above 0.50, the spiking has disappeared, given way only to the stable equilibrium and some other unstable fixed points.

The structure of this diagram gives a clue into how an internal cellular mechanism of calcium could produce bursting. Low Ca concentrations would send the system into tonic spiking – this is the only stable state. If the concentration then slowly rises, spiking will terminate when [Ca] surpasses 0.50, bringing the neuron to its depolarized rest state at $V \approx -60$ mV. If [Ca] then begins to decrease slowly, this stable equilibrium will be tracked for some time. But further decreasing it beyond the turnaround at 0.42 will send the system off of the stable fixed point and back to its only stable option – the limit cycle. If [Ca] begins to rise again, the process repeats.

The interplay of the two distinct timescales is clear: slow Ca modulates the fast Na-K system between quiescence at high [Ca] and spiking at low [Ca], creating well-separated bursts *without* requiring the modulation of the injected current. If the parameterization of the calcium dynamics – ϕ and τ_Ca in Eq. 4.21 – is chosen judiciously, then [Ca] will indeed traverse the bistable region from spiking to resting, producing the

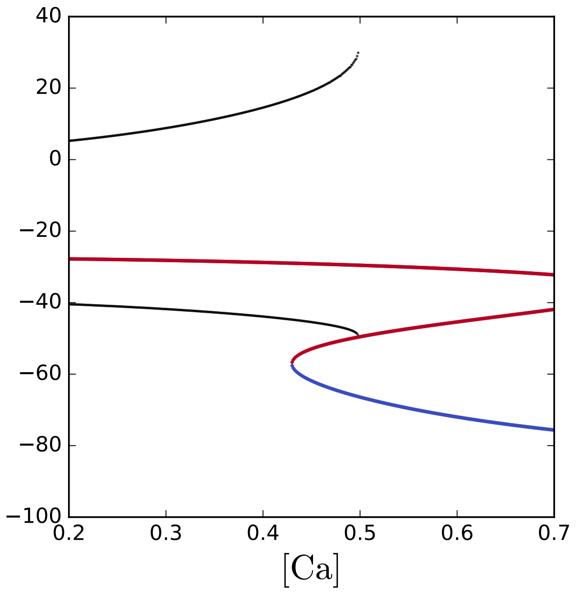

Figure 4.5: Bifurcation diagram for steady-state voltage in a HH-neuron, now with added Ca channels and Ca-gated K channels; parameters indicated in 4.3. The bifurcation parameter is intracellular calcium concentration [Ca], here assumed constant. Stable fixed points are shown in blue, unstable fixed points in red, and stable limit cycle bounds in black. A region of bistability is sandwiched between regions of tonic spiking on the left and hyperpolarized quiescence to the right.

bursting behavior suggested by the bifurcation diagram. Indeed, when the [Ca] dynamics are turned on, with parameters chosen as in Table 4.3, this is indeed what we find (Figure 4.6). The calcium concentration varies in the expected range with the peak concentration ($[\mathrm{Ca}] = 0.51$) coinciding with the transition from spiking to rest predicted by Figure 4.5.

4.3.5 Compartmental modeling and robust bursting

We see that neuron bursting can arise naturally from slow moving calcium ions. The original impetus for searching for this type of cellular mechanism of bursting was the observed stereotypy of $\mathrm{HVC_{RA}}$ bursts: if burst or spike characteristics are sensitive to incoming currents, fluctuations in synaptic currents would quickly destroy this homogeneity.

While calcium dynamics may produce bursting internally, is this enough to retain this robustness to synaptic stimuli? Are the bursts still sensitive to the mean and variance of the injected or synaptic current? To probe this, consider this neuron's response to a pseudo-noisy current, which consists of uniformly chosen I every $\Delta t = 5$ ms between -300 and 300 pA, linearly interpolating between. The resulting voltage and calcium

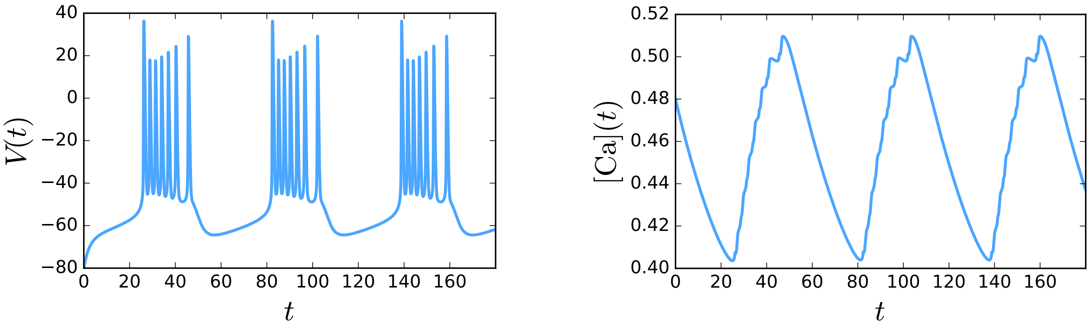

Figure 4.6: (a) Membrane voltage trace of HH-neuron with added Ca channels, Ca-gated K channels, and Ca dynamics (Eq. 4.21). As predicated by the bifurcation diagram in Figure 4.5, the neuron switches between well-separated periods of tonic spiking and quiescence when injected with a steady current. (b) The corresponding trace of intracellular $[\mathrm{Ca}](t)$.

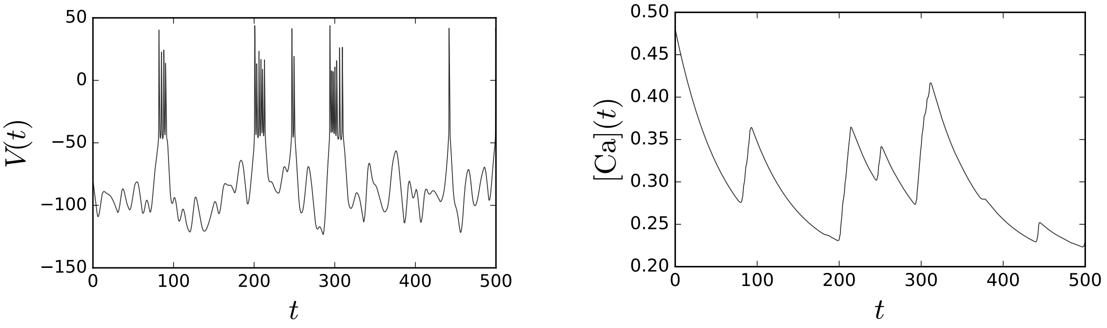

Figure 4.7: Voltage and calcium traces of the HH-neuron with Ca dynamics in response to noisy injected currents shows that the individual burst waveforms, while separated, are still sensitive to variations in the injected current.

traces are shown in Figure 4.7; the bursts are varied, ranging from 1 to 7 spikes or more.

So while the neuron may produce bursts, these bursts waveforms are somewhat sensitive to the injected current. We can reduce this sensitivity by adding an extra mechanism between the calcium and voltage bursts, slightly decoupling these phenomena. This can be accomplished in a biophysically-realistic with some morphological considerations; in particular, we separate the fast spiking behavior and slow calcium and calcium-driven currents into a separate somatic and dendritic “compartments.” These compartments are connected ohmically, the intent being that the fast somatic spiking is now driven by broad dendritic voltage bursts which wash out fast, subthreshold fluctuations and render the somatic spike bursts more robust.

In this scheme, the membrane voltage dynamics assume the following schematic form:

$$C_\mathrm{m}\frac{dV_\mathrm{s}}{dt} = -\sum_i I_i + \bar{g}_\mathrm{SD}(V_\mathrm{d} - V_\mathrm{s}) \tag{4.22}$$

$$C_\mathrm{m}\frac{dV_\mathrm{d}}{dt} = -\sum_i I_i + \bar{g}_\mathrm{SD}(V_\mathrm{s} - V_\mathrm{d}). \tag{4.23}$$

where the somatic currents include K and Na, while the dendritic currents include low-threshold calcium and calcium-gated potassium channels. Either the soma or dendrite could have injected currents, depending on the experimental setup, although somatic injections are more feasible. Finally, the synaptic currents, included, would be placed in the dendritic compartment.

With all this in mind, let us delineate the full model system for the HVC RA-projecting neuron, and then illustrate how this compartmental model smoothes the sensitivity to variances in the injected current. We will also show how the model qualitatively reproduces many experimentally-observed features of HVC RA-projecting neurons in response to calcium blockers and varying injected currents. Finally, we will move to the prodigious task of determining the many unknown parameters of this model from short, noisy, intracellular recordings[3]

4.3.6 Full HVC model description and parameterization

The HVC_RA neuron model is 7-dimensional, described by the following dynamical equations:

$$\begin{aligned}
C_\mathrm{m}\frac{dV_\mathrm{s}}{dt} &= -I_\mathrm{K} - I_{\mathrm{N}a} - I_\mathrm{L} + I_\mathrm{SD} + I_\mathrm{inj,s}(t) \\
C_\mathrm{m}\frac{dV_\mathrm{d}}{dt} &= -I_\mathrm{Ca\text{-}L} - I_\mathrm{K/Ca} - I_\mathrm{SD} + I_\mathrm{inj,d}(t) \\
\frac{di}{dt} &= \frac{i_\infty(V_\mathrm{s}) - i}{\tau(V_\mathrm{s})}, \quad i = m, n, h \\
\frac{dr}{dt} &= \frac{r_\infty(V_\mathrm{d}) - r}{\tau(V_\mathrm{d})} \\
\frac{d[\mathrm{Ca}]}{dt} &= \phi I_\mathrm{Ca\text{-}L} + \frac{[\mathrm{Ca}]_0 - [\mathrm{Ca}]}{\tau_\mathrm{Ca}},
\end{aligned} \tag{4.24}$$

---

[3] This study is also described in detail in its original published form in [Kadakia et al., 2016].

with the following ionic currents

$$\begin{aligned}
I_{\mathrm{K}} &= \bar{g}_{\mathrm{K}} n^4 \cdot (V_{\mathrm{s}} - E_{\mathrm{K}}) \\
I_{\mathrm{N}a} &= \bar{g}_{\mathrm{Na}} m^3 h \cdot (V_{\mathrm{s}} - E_{\mathrm{Na}}) \\
I_{\mathrm{L}} &= \bar{g}_{\mathrm{L}} (V_{\mathrm{s}} - E_{\mathrm{L}}) \\
I_{\mathrm{SD}} &= \bar{g}_{\mathrm{SD}} (V_{\mathrm{d}} - V_{\mathrm{s}}) \\
I_{\text{Ca-L}} &= \bar{g}_{\text{Ca-L}} r^2 \Phi_{\mathrm{GHK}}(V_{\mathrm{d}}) \\
I_{\mathrm{K/Ca}} &= \bar{g}_{\mathrm{K/Ca}} s_\infty([\mathrm{Ca}])(V_{\mathrm{d}} - E_{\mathrm{K}}) \qquad (4.25)
\end{aligned}$$

and static functions

$$\begin{aligned}
i_\infty(V) &= \frac{1}{2}\left(1 + \tanh\frac{V - \theta_i}{2\sigma_i}\right), \quad i = m,n,h,r \\
\tau_i(V) &= \tau_{i,0} + \tau_{i,1}\left(1 - \tanh^2\frac{V - \theta_{\tau,i}}{2\sigma_{\tau,i}}\right) + \tau_{i,2}\left(1 + \tanh\frac{V - \theta_{\tau,i}}{2\sigma_{\tau,i}}\right), \quad i = m,n,h,r \\
s_\infty([\mathrm{Ca}]) &= \frac{[\mathrm{Ca}]^2}{[\mathrm{Ca}]^2 + k_s^2} \\
\Phi_{\mathrm{GHK}}(V_{\mathrm{d}}) &= V_{\mathrm{d}} \frac{[\mathrm{Ca}]_{\mathrm{out}} e^{-V_{\mathrm{d}}/V_{\mathrm{T}}} - [\mathrm{Ca}]}{e^{-V_{\mathrm{d}}/V_{\mathrm{T}}} - 1} \qquad (4.26)
\end{aligned}$$

Note that the τ_i functions in this model are slightly more general than those in Eq. 4.11 in that the thresholds and widths $\theta_{\tau,i}$ and $\sigma_{\tau,i}$ may be distinct than those in $i_\infty(V)$.

The conductance, timescales, threshold voltages, rate constants, and other parameters describing this neuron may vary, sometimes considerably, even within the same class of RA-projecting neurons. The parameters will of course be determined from data, yet as a first indication that the model lives at a proper nexus of sparsity and versatility, let us illustrate a few responses that reflect at least qualitatively what is known from experiment.

4.3.7 Qualitative experimental predictions

To this end, we will fix the parameters of this model to those listed in Table 4.4, which will also be used to generate synthetic data to test the inference methods. The first feature predicted by experiment, of course, are somatic voltage bursts more robust than those found in the single-compartment model, Figure 4.7.

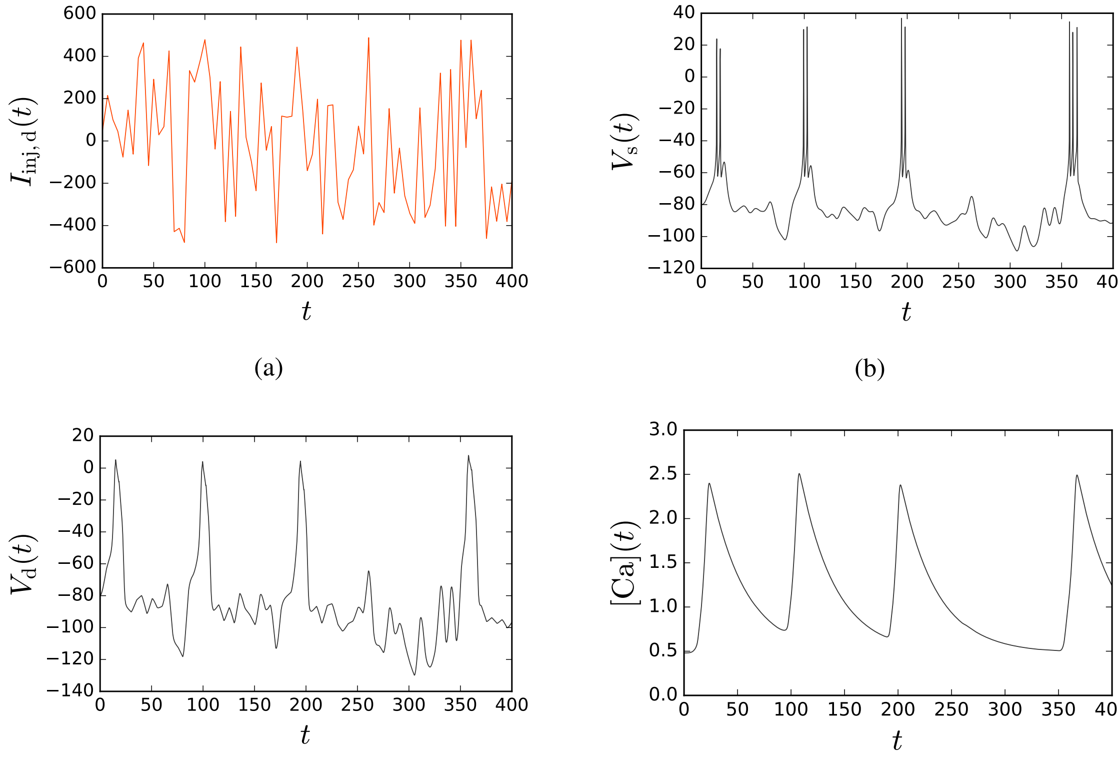

(a)

(b)

(c)

(d)

Figure 4.8: Injected current and responses of some dynamical variables in the two-compartment $\mathrm{HVC_{RA}}$ neuron model, Eqs. 4.24. (a) Injected current into the dendrite, consisting of pseudo-noisy waveform in the excitable range of the neuron model. (b) Voltage response of neuron soma shows that despite variations in the injected current, the bursts are highly stereotyped. (c) Broad dendritic voltage bursts, which elicit the short somatic voltage bursts. (d) Corresponding calcium traces show broad, slow decays before a new burst can be initiated.

Let us record the response to the pseudo-noisy injected current from before, consisting of uniformly chosen $I_{\mathrm{inj,d}}$ every $\Delta t = 5$ ms between -500 and 500 pA, linearly interpolating between (Figure 4.8a). The current is injected into the dendrite, not soma, to mimic synaptic currents which pass from axon to dendrite, then to the cell body. The resulting somatic voltage, Figure 4.8b, exhibits a tighter distribution of burst length and spikes per burst, arising from the stereotyped *dendritic* bursts (Figure 4.8c), in turn caused by the slow motion of calcium, which is now effectively decoupled from the cell body.

It has also been reported that the interburst interval is sensitive to the calcium channel conductance. In particular, [Jin et al., 2007] have found that when calcium channel blockers are applied to the sleeping bird, the time between bursts increases noticeably, and the opposite effect occurs with the application of calcium enhancers. We can test this in the compartmental model by adjusting the calcium channel conductance, $\bar{g}_{\text{Ca-L}}$

Table 4.4: Parameters of two-compartment HVC$_{\mathrm{RA}}$ neuron model.

$\bar{g}_{\mathrm{K}}$	120	$\bar{g}_{\mathrm{L}}$	3	$\bar{g}_{\mathrm{Na}}$	1050	$\bar{g}_{\mathrm{SD}}$	5	nS
$\bar{g}_{\mathrm{K/Ca}}$	240							nS
$\bar{g}_{\mathrm{Ca\text{-}L}}$	0.06							nS/μM
E_{K}	-90	E_{L}	-80	E_{Na}	55	V_{T}	13.5	mV
θ_n	-35	σ_n	10	$\theta_{\tau,n}$	-27	$\sigma_{\tau,n}$	-15	mV
θ_m	-30	σ_m	9.5	$\theta_{\tau,m}$	-	$\sigma_{\tau,m}$	-	mV
θ_h	-45	σ_h	-7	$\theta_{\tau,h}$	-40.5	$\sigma_{\tau,h}$	-6	mV
θ_r	-40	σ_r	10	$\theta_{\tau,r}$	-	$\sigma_{\tau,r}$	-	mV
τ_{n0}	.1	τ_{n2}	.5	τ_{m0}	.01	τ_{m2}	0	ms
τ_{h0}	.1	τ_{h2}	.75	τ_{r0}	1	τ_{r2}	0	ms
$[\mathrm{Ca}]_0$	0.48	k_{s}	3.5	$[\mathrm{Ca}]_{\mathrm{out}}$	2500			μM
ϕ	8.67e-5							μM/pA/ms
τ_{Ca}	33							ms
C_{m}	21							pF

– a reduction would correspond to the presence of an antagonist, and an enhancement to that of an agonist. Figure 4.9 shows the response of the neuron to a pseudo-noisy current with the original conductance $\bar{g}_{\mathrm{Ca\text{-}L}} = 0.06$ nS/μM, a reduced conductance of 0.03 nS/μM, and an inflated value of 0.18 nS/μM; the interburst interval reproduces, at least qualitatively, the experimentally-backed expectations.

The authors in [Jin et al., 2007] also demonstrate a difference in response to somatically and dendritically injected currents. While step currents injected into the dendrite (externally or via synapses) elicit an all-or-nothing burst, increasing step heights into the soma produce instead bursts of gradually increasing spike count. We test this latter prediction by setting $I_{\mathrm{inj,s}}$ to values between 150 and 350 pA for a 20 ms window. As seen in Figure 4.10, the spike count increases steadily with the injected current magnitude, in contrast to an all-or-nothing burst with pre-defined spike count for dendritic current injections. We also see that spikes may occur for low current injections without the presence of the initiating calcium burst.

Interestingly, the same study found that all-or-nothing bursting could be initiated with somatic current injection if done in the presence of a calcium agonist. The model reproduces this finding as well: we again inject a short 20 ms current step of varying magnitude with an enhanced calcium conductance of $\bar{g}_{\mathrm{Ca\text{-}L}} = 0.18$ nS/μM, finding that for current magnitudes above 100 pF or so, a burst of fixed size and spike count results (Figure 4.11). These somatic voltage bursts coincide with a calcium burst in the dendrite, which is now easier to incite due to the calcium agonist.

These four behaviors corroborate the plausibility of the compartmental model of HVC$_{\mathrm{RA}}$ neurons.

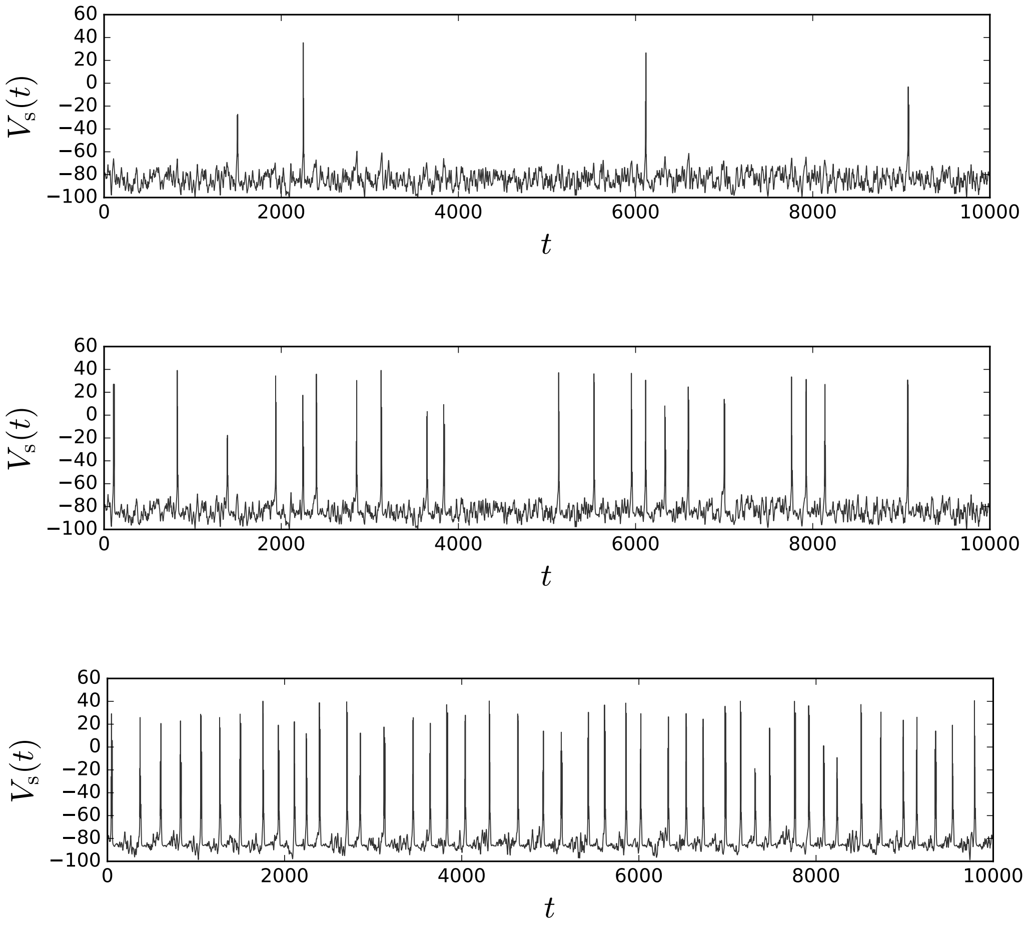

Figure 4.9: Effect of calcium agonists on the HVC$_{\text{RA}}$ neuron are modeled by changes in the calcium conductance $\bar{g}_{\text{Ca-L}}$. The burst incidence increases as $\bar{g}_{\text{Ca-L}}$ is increased from 0.5 (top panel), to 1.0 and 3.0 (middle and bottom panels), in line with experiment [Jin et al., 2007].

But exact metrics such as the waveforms of the spikes, spike count, burst sensitivity, etc., have not been tuned. Since these aspects depend, sometimes quite sensitively, on the model parameterization, we must ultimately leave the determination of the exact parameters to the inverse problem. To this end, we will now see how variational annealing can be used to successfully determine the unknown parameters of our proposed HVC$_{\text{RA}}$ model from realistic observed datasets and realistic prior information.

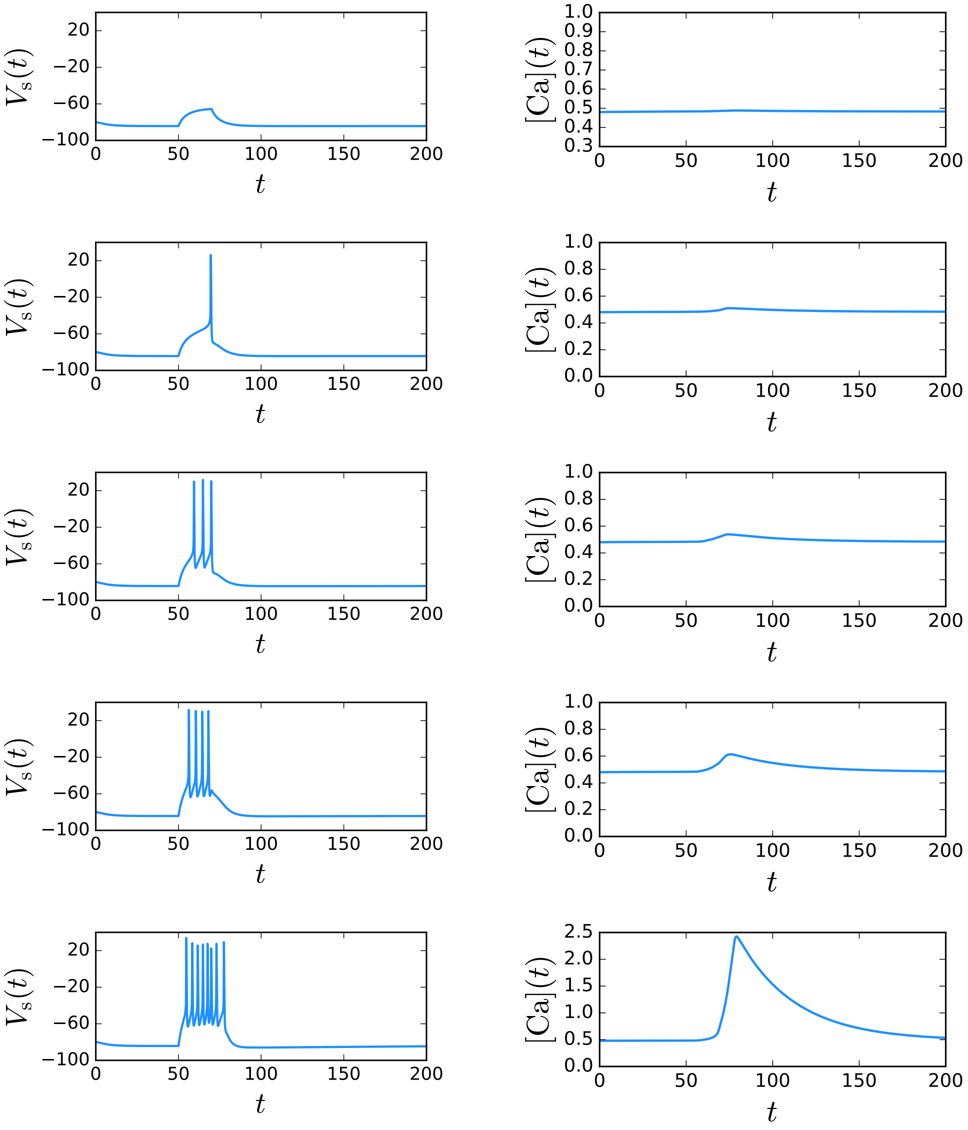

Figure 4.10: The HVC$_{\mathrm{RA}}$ does not exhibit all-or-nothing responses to current injected into the cell body (somatic compartment), as indicated by experiment [Jin et al., 2007]. Somatic voltage responses (left column) and calcium responses (right column) to short, 20 ms somatic current pulses of varying magnitude are shown; the pulse magnitudes range from 100 pA in the top row to 300 pA in the bottom row.

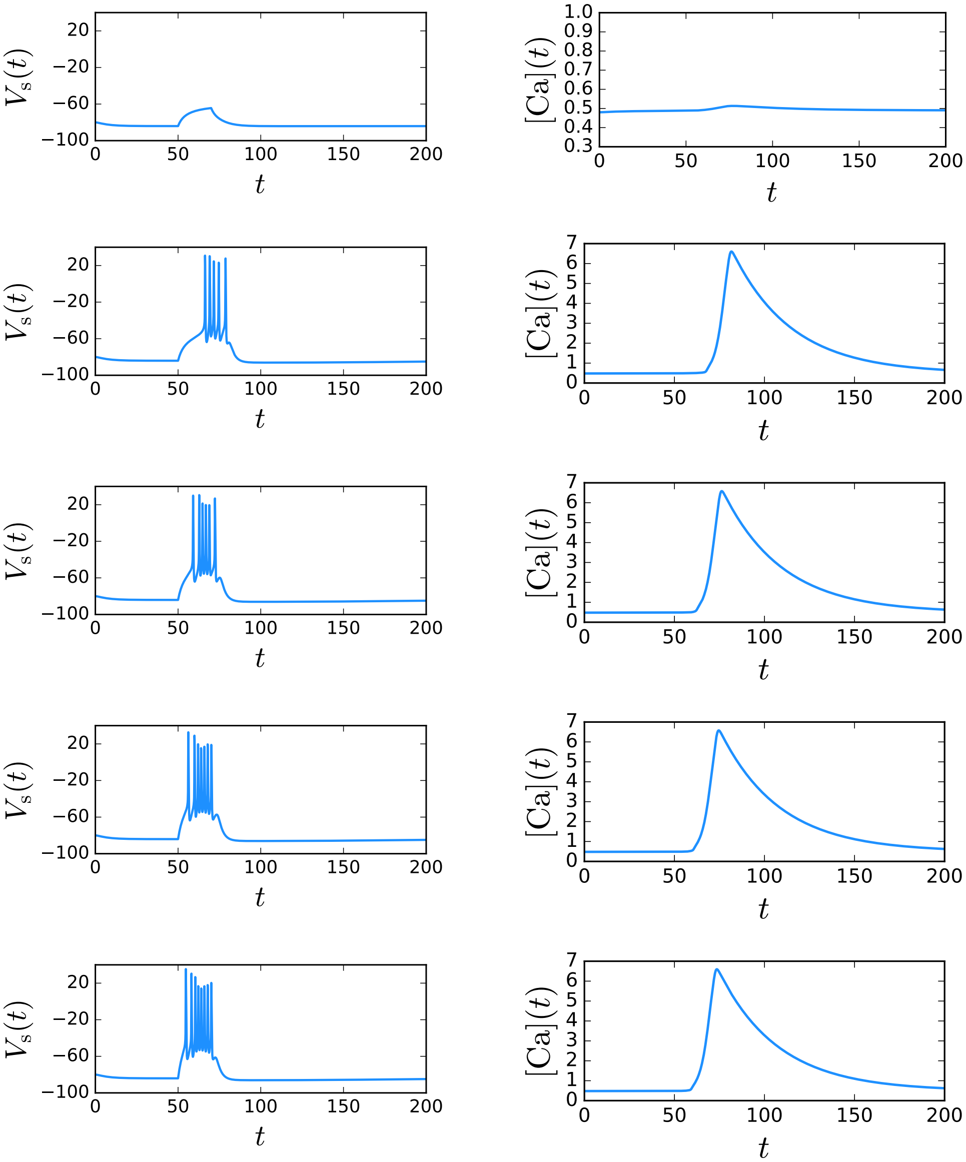

Figure 4.11: The same plots as in Figure 4.10, now with the calcium conductance increased, mimicking the presence of a calcium agonist. With sufficient agonists, the all-or-nothing behavior can again be evoked, again in line with experimental evidence [Jin et al., 2007].

4.4 Inferring the $\mathrm{HVC_{RA}}$ model parameters

4.4.1 Information from transients

For single neurons, patch clamp electrophysiology can be used to accurately measure membrane voltage response to injected current waveforms [Sakmann and Neher, 1984, Covey and Carter, 2015]. What's more, synaptic blockers may be used as well to eliminate synaptic currents to a large degree, functionally isolating the cell from its environment. Since the current is delivered externally, it is known, but among the variables in the model description, only the soma membrane voltage is observable.

The currents delivered in a current clamp experiment are user-defined. Step currents are informative, but only weakly. A weak step current (below spiking threshold) would probe only a single point in the neuron phase space – the fixed point at hyperpolarization – but even a higher step current would probe just a single, periodic closed curve in phase space, the spiking limit cycle. The information transferred from the resulting voltage traces would be far too little to successfully determine the many missing parameters. The same would be true for other common current waveforms such as saws, pulses, and sinusoids.

An alternative viewpoint is to deliver an injected current that contains underlying structure, but more densely fills phase space. This can be achieved by subjecting the neuron to a current that spans a broad power spectrum and many characteristic frequencies. Since the neuron is stable (either to a limit cycle or fixed point) in the dynamic range in question, static currents alone will quickly relax to their limiting stable manifolds. Thus, a useful current is one that produces information by continually pushing the system to *transients* between these limiting regions.

With this in mind, we will carry out the variational annealing protocol using V_s time traces from an injected current consisting of a normalized version of one component of the chaotic Lorenz oscillator, an idea that has proven useful in past studies [Toth et al., 2011, Kostuk et al., 2012, Meliza et al., 2014, Kadakia et al., 2016]. The normalization is chosen to span a realistic range of injected currents, plus or minus a few hundred pA; the time axis is also scaled so the highest frequency is not larger than the fastest timescale of the neural model, about 0.01 ms. In addition to this "chaotic" current, we add a pseudo-noisy current generated as described in Section 4.3.5.

We create a synthetic observed voltage trace by injecting this current waveform (Figure 4.12) into the somatic compartment and integrating the neuron dynamics, Eqs. 4.24, forward for 600 ms. The timestep

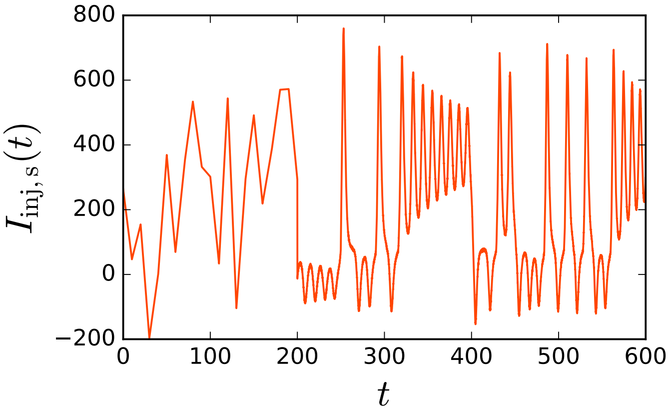

Figure 4.12: Injected somatic current used to generate the data for the numerical experiments on the HVC$_{\mathrm{RA}}$ model. The current consists of a pseudo-noisy waveform and the output of the Lorenz attractor, intended to push the system along its transients to more fully probe the dynamical phase space.

used in the integration is 0.02 ms, in line with the commonly used 50 kHz sample rate of current injection and data acquisition in typical electrophysiological setups; in addition the 600 ms recordings are completely within acceptable *in vitro* recording times, and will allow for multiple epochs to record before the neuron dies. The synthetic observation vector, $\mathbf{Y} = \{y_{V_s}(t_n)\}$ is then produced by adding idd Gaussian noise of null mean and deviation $\sigma_m = 2$ mV to the soma voltage time trace.

4.4.2 Timescale separation and an experimental protocol

Carrying out a variational annealing procedure with this set of synthetic observations does not succeed – many of the parameters are estimated quite poorly, and in many cases the optimization does not converge, likely due to a poorly conditioned cost surface $A(\mathbf{X}|\mathbf{Y})$, arising from the small observation dimension (1 of 7 variables), considerable nonlinearity in the model equations, and large unknown parameter space. The failure of the procedure is also attributed partly to the interplay of two disparate timescales in the model, one governing the slow calcium dynamics, and another the fast spiking dynamics in the neuron cell body. This disparity can introduce extreme sensitivities to parameterization, which in an optimization context manifest as additional roughness or nonconvexity in the cost surface.

To address these issues within an experimentally realizable framework let's consider the following. We have seen that the compartmental model contains fast spiking currents, K and Na, periodically modulated on a slower scale by the calcium-gated potassium current. Intracellular calcium moves five to ten times slower

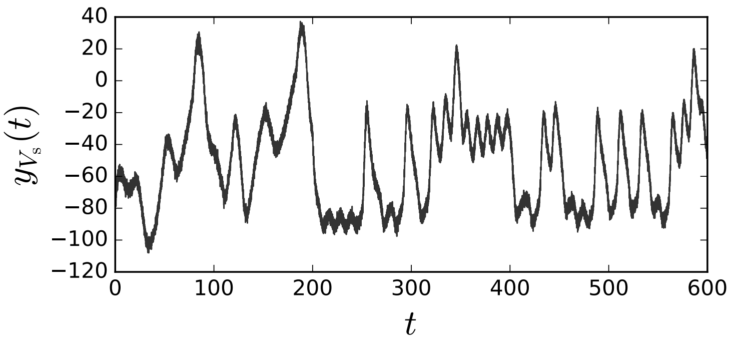

Figure 4.13: Synthetic somatic voltage observations for the reduced two-compartment HVC_{RA} model with fast currents switched off.

than spike-mediating Na and K currents. If we decouple the neuron into fast and slow terms by setting $\bar{g}_{\text{Na}}$ and $\bar{g}_{\text{K}}$ to zero, effectively blocking the fast currents, it may then be possible to estimate the parameters governing the slow variation – Ca and Ca-mediated channels – alone. Then, setting these estimated parameters fixed in the full model, the remaining parameters governing the fast variables can be estimated through a second application of the inference routine.

This idea can be implemented in practice by carrying out two recordings of V_s with the same patch clamp, the first as usual and the second with the application of a sodium channel blocker such as TTX, which can be done without repatching (carrying this out in the opposite order would likely subject the neuron to lingering effects of the channel blocker in the unadulterated recording trace). The observed traces are then used in reverse order in the annealing routines, as suggested above. The idea is that decoupling the fast and slow processes would smooth the cost function surface and more easily reveal the lowest minima.

4.4.3 Estimation of slow subsystem: computational details

To this end, we will first create synthetic observations of the reduced model neuron by integrating Eq. 4.24 in steps of $\Delta t = 0.02$ with $\bar{g}_{\text{Na}} = 0$ and $\bar{g}_{\text{K}} = 0$ and adding iid Gaussian noise with deviation 2 mV. The spiking mechanism is eliminated, so no single spikes are seen, only broader, slower, depolarizations (Figure 4.13).

Each step of variational annealing carries out an optimization over all collocated state variables $x_d(t_n)$, as well as parameters $\mathbf{p} = \{p_i\}$ [4]. Since the sodium and potassium currents have been suppressed,

---

[4] A notational note: In later chapters, **p** will refer to momenta variables; the distinction should be obvious from the context, and when it is not, will be explicitly stated

their associated kinetic parameters do not enter the dynamical equations, and are thereby removed from the search space. We also add the caveat that the background calcium concentration parameters, $[\mathrm{Ca}]_0$ and $[\mathrm{Ca}]_{\mathrm{out}}$, were required to be held fixed for the inference to produce accurate estimations. Interestingly, this does not imply that these parameters be known *a priori*. In fact, to highlight this, let us fix them to *incorrect* values $[\mathrm{Ca}]_0 = 0.285$ and $[\mathrm{Ca}]_{\mathrm{out}} = 1000$ in the model equations. As will be explained in detail later, the accurate estimations arising from incorrect prior information will be attributed to a partial degeneracy of the system.

$A(\mathbf{X}|\mathbf{Y})$ is very high-dimensional. There are four collocated dynamical variables to be estimated at every $\Delta t = 0.02$ ms in a 600 ms window (since $\bar{g}_{\mathrm{K}} = \bar{g}_{\mathrm{Na}} = 0$, then $m(t)$, $n(t)$ and $h(t)$ are eliminated). Adding the 17 unknown parameters relevant for the slow dynamics with $[\mathrm{Ca}]_0$ and $[\mathrm{Ca}]_{\mathrm{out}}$ fixed, produces a 120,017-dimensional search space $\mathbf{X} \equiv \{\mathbf{x}(t_n), \mathbf{p}\}$. As in the variational annealing examples in Section 3.2.2, the interior point method IPOPT will be used to carry out the optimization at each stage of the annealing.

In actual implementation, $A(\mathbf{X}|\mathbf{Y})$ requires a discretization of the continuous neuron dynamics, Eqs. 4.24. One benefit of this collocated cost function is that implicit and explicit integrators are equally simple to implement, since integration is not actually carried forward in time. We will use a 4th-order Simpson rule, which interpolates a function quadratically using third-order Lagrange polynomials between every other timestep. In this scheme, an interpolated function $f(x(t))$ and its integral in the region of interpolation $[t_n, t_n + 2\Delta t]$ are [Abramowitz and Stegun, 1964]:

$$\begin{aligned} f(x(t)) &\approx \sum_{j=1}^{3} f(x(t_j)) \prod_{k\neq j}^{3} \frac{t - t_k}{t_j - t_k} \\ &= \frac{(t-(t_n+\Delta t))(t-(t_n+2\Delta t))}{2\Delta t^2} f(x(t_n)) \\ &\quad + \frac{(t-t_n)(t-(t_n+2\Delta t))}{-\Delta t^2} f(x(t_n+\Delta t)) \\ &\quad\quad + \frac{(t-t_n)(t-(t_n+\Delta t))}{2\Delta t^2} f(x(t_n+2\Delta t)) \end{aligned} \tag{4.27}$$

and

$$\begin{aligned} \int_{t_n}^{t_n+2\Delta t} f(x(t))dt &\approx \frac{\Delta t}{3}\left[f(x(t_n)) + 4f(x(t_n+\Delta t)) + f(x(t_n+2\Delta t))\right] \\ &\equiv \frac{\Delta t}{3}\left[f(x(t_n)) + 4f(x(t_{n+1})) + f(x(t_{n+2}))\right]. \end{aligned} \tag{4.28}$$

Table 4.5: Upper and lower bounds on the parameters for the constrained optimization in the reduced HVC$_{\text{RA}}$ model. $[\text{Ca}]_0$ and $[\text{Ca}]_{\text{out}}$ are fixed at the erroneous values 0.285 and 1000, respectively, so they do not enter the optimization routine. Also listed are the upper and lower bounds and initializing R_{f0} values for the four active state variables.

parameter	lower	upper	parameter	lower	upper
$\bar{g}_{\text{L}}$	0.1	10	C_{m}	1	100
$\bar{g}_{\text{Ca-L}}$	0	10	k_{s}	1	100
$\bar{g}_{\text{K/Ca}}$	0	5000	ϕ	1e-5	1e-2
$\bar{g}_{\text{SD}}$	1	50	τ_{r0}	0	1
E_{L}	-110	-70	τ_{r1}	0	1
E_{K}	-100	-75	τ_{r2}	0	1
θ_r	-50	-10	τ_{Ca}	20	50
σ_r	5	25	$[\text{Ca}]_0$	-	-
$\theta_{\tau,r}$	-50	-10	$[\text{Ca}]_{\text{out}}$	-	-
$\sigma_{\tau,r}$	5	25			
state	lower	upper	R_{f0}		
$V_{\text{s}}(t_n)$	-120	50	1e-3		
$V_{\text{d}}(t_n)$	-120	50	1e-3		
$r(t_n)$	0	1	1e1		
$[\text{Ca}](t_n)$	0	1	1e1		

Applying this discretization of the dynamical term in the Gaussian error approximation $A(\mathbf{X}|\mathbf{Y})$ [5]:

$$\begin{aligned} \frac{dx_d}{dt} &= f_d(\mathbf{x},\mathbf{p}) \qquad \text{with} \qquad f_d(\mathbf{x},\mathbf{p}) = \text{Eq. 4.24} \\ \rightarrow x_d(t_{n+2}) &= x_d(t_n) + \int_{t_n}^{t_{n+2}} f_d(\mathbf{x},\mathbf{p})dt \\ &\approx x_d(t_n) + \frac{\Delta t}{3}\left[f_d(\mathbf{x}(t_n),\mathbf{p}) + 4f_d(\mathbf{x}(t_{n+1}),\mathbf{p}) + f_d(\mathbf{x}(t_{n+2}),\mathbf{p})\right]. \end{aligned} \tag{4.29}$$

Note that this interpolation, being 3rd order, occurs at every *other* step, in time intervals of $2\Delta t$. The full Gaussian error cost function for the HVC$_{\text{RA}}$ neuron is therefore (cross reference Eq. 2.44):

$$\begin{aligned} A(\mathbf{X}|\mathbf{Y}) = \sum_{n=1}^{N} \frac{R_{\text{m}}}{2}\left(V_{\text{s}}(t_n) - y_{V_{\text{s}}}(t_n)\right)^2 \\ + \sum_{n\text{ even}}^{N-2} \sum_{d=1}^{D} \frac{R_{\text{f}}(d)}{2}\Big(x_d(t_{n+2}) - x_d(t_n) \\ - \frac{\Delta t}{3}\Big[f_d(\mathbf{x}(t_n),\mathbf{p}) + 4f_d(\mathbf{x}(t_{n+1}),\mathbf{p}) + f_d(\mathbf{x}(t_{n+2}),\mathbf{p})\Big]\Big)^2, \end{aligned} \tag{4.30}$$

---

[5] Though not shown for purposes of brevity, the vector field $\mathbf{f}$ also depends on time here through the time-dependence of the injected current: $\mathbf{f} = \mathbf{f}(\mathbf{x},\mathbf{p},t)$. The discretization scheme simply evaluates the current at the same time as the state variable; i.e $f_d(\mathbf{x}(t_n),\mathbf{p}) = f_d(\mathbf{x}(t_n),\mathbf{p},t_n)$, etc.

where $\mathbf{X} = \{\mathbf{x}(t_n), \mathbf{p}\}$, with

$$\begin{aligned}
&\{\mathbf{x}(t_n)\} = \{V_\mathrm{s}(t_n), V_\mathrm{d}(t_n), r(t_n), [\mathrm{Ca}](t_n)\}\\
&\mathbf{p} = \text{ listed in Table 4.5,}\\
&\mathbf{Y} = \{y_{V_\mathrm{s}}(t_n)\}
\end{aligned} \tag{4.31}$$

and $N = 600$ timesteps and $D = 4$ dimensions. The observation inverse covariance of the measured somatic voltage is $R_\mathrm{m} = 1/2^2 = 0.25\ \mathrm{mV}^{-2}$. The optimizations will also require inequality constraints on the state variables and parameters: these are listed in Table 4.5, and reflect relatively agnostic priors on conductances, which can vary by orders of magnitudes among distinct neurons. The kinetic parameters and timescales are also constrained to physically plausible ranges. The reversal potentials are quite uniform and not difficult to estimate by eye anyway, so their ranges are tighter. Finally, the state variables $\{x_d(t_n)\}$ are constrained to their known physical ranges.

As described at length in Chapter 3, variational annealing strings together Q isolated optimizations within an iterative scheme. For the first step, $R_\mathrm{f}(d)$ is set to $R_\mathrm{f0}(d)$, which are listed in Table 4.5. The values are chosen to normalize the model errors to similar ranges: since V_s and V_d have larger dynamical ranges, their associated R_f0 are a few orders smaller. The Q parallel optimizations at this first annealing step are initialized by choosing the 120,000 collocated state variables and 17 free parameters from their bounds listed in Table 4.5. The annealing progresses for $\beta_\mathrm{max} = 30$ steps with $\alpha = 2$, as described in detail Section 3.1. For this problem, we will run $Q = 100$ parallel optimizations.

4.4.4 Estimation of slow subsystem: results

At each value of β, the optimization returns $Q = 100$ paths trajectories, in addition to 17 parameter estimates:

$$\hat{\mathbf{X}}^{q,\beta} = \{\hat{x}_d^{q,\beta}(t_n), \hat{p}_i^{q,\beta}\} \tag{4.32}$$

The values $A(\hat{\mathbf{X}}^{q,\beta}|\mathbf{Y})$ are plotted in the anneal plot shown in Figure 4.14. All Q anneals appear to converge to the same value of A =1.59e4, well in the vicinity of the expected limiting error value for the perfect model (Eq. 3.10; here it is E = 1.50e4), suggesting that all estimates have converged to the true trajectory.

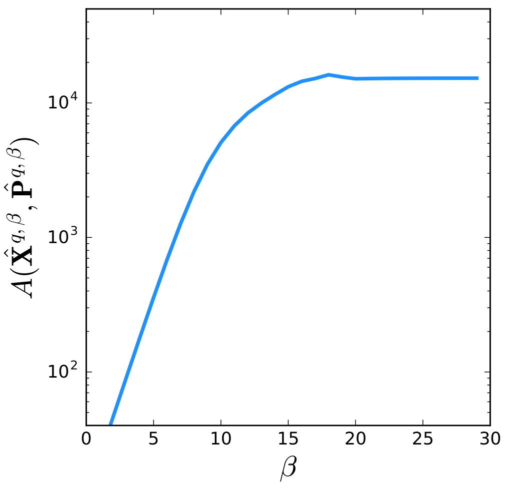

Figure 4.14: Annealing plot for the variational annealing estimation of the reduced two-compartment $\mathrm{HVC_{RA}}$ model. A single minimum is found for all estimations, with limiting value within the expected range of the global minimum (Eq. 3.10).

The associated paths are plotted in Figure 4.15, and the estimated parameters are listed in Table 4.6. As a comparison, when the optimization is carried out directly without annealing, 100 distinct estimates are returned, none of which even mildly mirror the accuracy found with variational annealing.

There are a few things of note about the estimated trajectories and parameters. The voltages and gating variable $r(t_n)$ are estimated to excellent accuracy, while the calcium estimate appears to be a shifted replica of the true trace. Recall that we fixed the background and extracellular calcium concentrations to erroneous – thought not wholly unphysical – values. The shifted calcium estimate is a reflection of this inaccurate prior, yet the precision of the other state variables, $V_\mathrm{s}(t)$, $V_\mathrm{d}(t)$ and $r(t)$ suggests that the model itself exhibits a degeneracy that cannot be probed by the measurements of somatic voltage alone. Inspecting further, we see that the parameter estimates, most of which are estimated to excellent accuracy, exhibit a few inconsistencies involving the calcium ions: ϕ, k_s, and $\bar{g}_\text{Ca-L}$ are all off by factor of 2 or so.

The picture we should adopt is the following: there is a degenerate manifold in the ϕ-k_s-$\bar{g}_\text{Ca-L} - [\mathrm{Ca}]_\mathrm{out} - [\mathrm{Ca}]_0$ parameter subspace, whereby differing combinations of these parameters can produce identical traces of $V_\mathrm{s}(t)$, $V_\mathrm{d}(t)$, and $r(t)$, but $[\mathrm{Ca}](t)$ traces shifted to varying degrees. From the point of view of the experimentalist, the degeneracy is *invisible*, and from the perspective of inference, it is problematic – the degenerate subspace transforms point minima into valleys that may hinder the optimization from reliably terminating. This is why Ca_0 and Ca_out needed to be fixed in the first place (even at incorrect values), and it comforting to see in this case that a biophysically plausible estimate emerges nonetheless. Momentarily,

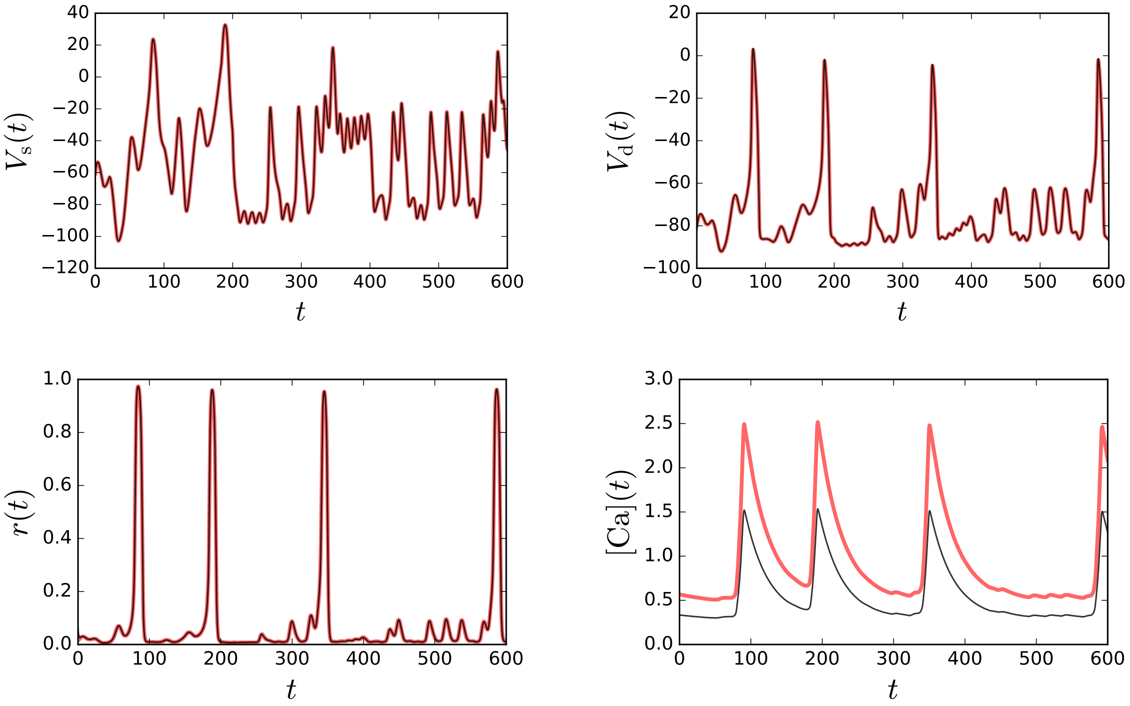

Figure 4.15: Estimates of a few select state variables of the reduced two-compartment $\mathrm{HVC_{RA}}$ model. The estimated trajectories are excellent; the calcium trace is slightly shifted due to a partial degeneracy of the system.

we will look in a bit more detail at this degeneracy from the point of view of the bifurcations of the two-compartment model.

4.4.5 Full model with slow parameters fixed

Having estimated the 17 slow parameters, we may now turn back on the sodium and potassium conductances and estimate the remaining parameters of the full $\mathrm{HVC_{RA}}$ model from a measured soma membrane voltage trace. Let us use the same injected current (Fig. 4.12) to create a set of synthetic observations, again found by infusing the true integrated trajectory with $\sigma = 2$ mV noise. The observed voltage trace is plotted in Figure 4.16.

The parameters in Table 4.6 are fixed to their *estimated* values in Table 4.6, while the remaining parameters are optimized subject to the bounds listed in Table 4.7. Carrying out this variational anneal (with the same algorithmic hyperparameters α and $\beta_{\max}$ as in the reduced model), the terminal estimate corresponding to the lowest value of $A(\hat{\mathbf{X}}^{q,\beta_{max}}|\mathbf{Y})$ among the $Q = 100$ anneals is shown in Table 4.8 and Figure 4.17. These estimated states and parameters producing this minimal A value are denoted by $\{\hat{\mathbf{x}}^{q_0,\beta_{\max}}(t_n)\}$ and $\{\hat{\mathbf{p}}^{q_0,\beta_{\max}}\}$,

Table 4.6: Estimated parameters of reduced $\mathrm{HVC_{RA}}$ model, with only somatic voltage measured. Most parameters are estimated correctly to high accuracy; discrepancies in the other parameters are explainable via system degeneracies. $\theta_{\tau,r}$ and $\sigma_{\tau,r}$ are not relevant to the estimation since the corresponding coefficients $\tau_{r,0}$ and $\tau_{r,1}$ are zero.

parameter	actual	estimated	parameter	actual	estimated
$\bar{g}_{\mathrm{L}}$	3.00	3.03	C_{m}	21.00	21.01
$\bar{g}_{\mathrm{Ca\text{-}L}}$	0.060	0.155	k_{s}	3.5	2.1
$\bar{g}_{\mathrm{K/Ca}}$	240	243	ϕ	8.67e-5	5.22e-5
$\bar{g}_{\mathrm{SD}}$	5.00	4.96	τ_{r0}	1.00	0.99
E_{L}	-80.0	-80.1	τ_{r1}	0	0.00
E_{K}	-90.0	-89.9	τ_{r2}	0	0.08
θ_r	-40.0	-39.9	τ_{Ca}	33.0	33.1
σ_r	10.0	10.0	$\theta_{\tau,r}$	-	-50.0
$\sigma_{\tau,r}$	-	5.00			

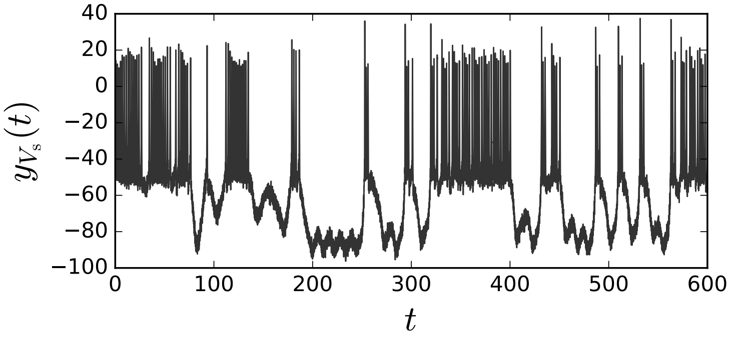

Figure 4.16: Trace of the synthetic voltage observations for the estimation of the full $\mathrm{HVC_{RA}}$ model, with slow parameters held fixed to their estimated values in Table 4.6.

respectively.

The parameter and state estimates are, for the most part, excellent, though there are some discrepancies. The calcium concentrations $\mathrm{Ca_0}$ and $\mathrm{Ca_{out}}$ are still fixed at erroneous values (0.285 and 1000 instead of 0.48 and 2500); further, some the slow parameters fixed from the estimation of the reduced model are now held fixed at their incorrectly-estimated values. For this reason, the estimate of $[\mathrm{Ca}](t)$ is again shifted from its true trace. Since the estimation of the fast components is still quite accurate, this is further indication that the inaccuracy in these parameter estimates is not a deficiency in the variational anneal, but a true degeneracy of the model description.

There are also a few parameters relevant to fast spiking that are inaccurate, $\bar{g}_{\mathrm{Na}}$ for example. The juxtaposition of this inaccuracy with the extreme precision in the state estimates in Fig. 4.17 suggests that

Table 4.7: Upper and lower bounds on the parameters for the constrained optimization in the full HVC$_{\mathrm{RA}}$ model. The remaining parameters of the full model are held fixed at those listed in Table 4.6. Also listed are the upper and lower bounds and initializing R_{f0} values for the seven state variables.

parameter	lower	upper	parameter	lower	upper
$\bar{g}_{\mathrm{K}}$	0	5000	θ_m	-50	-10
$\bar{g}_{\mathrm{Na}}$	0	5000	σ_m	6.25	16.67
E_{Na}	50	60	$\theta_{\tau,m}$	-40	-20
τ_{m0}	0	1	$\sigma_{\tau,m}$	-50	-5
τ_{m1}	0	1	θ_n	-50	-10
τ_{m2}	0	1	σ_n	6.25	16.67
τ_{n0}	0	1	$\theta_{\tau,n}$	-40	-20
τ_{n1}	0	1	$\sigma_{\tau,n}$	-50	-5
τ_{n2}	0	1	θ_h	-50	-10
τ_{h0}	0	1	σ_h	-16.67	-6.25
τ_{h1}	0	1	$\theta_{\tau,h}$	-50	-20
τ_{h2}	0	1	$\sigma_{\tau,h}$	-50	-5

state	lower	upper	R_{f0}
$V_{\mathrm{s}}(t_n)$	-120	50	1e-3
$V_{\mathrm{d}}(t_n)$	-120	50	1e-3
$m(t_n)$	0	1	1e1
$h(t_n)$	0	1	1e1
$n(t_n)$	0	1	1e1
$r(t_n)$	0	1	1e1
$[\mathrm{Ca}](t_n)$	0	1	1e1

the slow subsystem degeneracy may extend to the fast subsystem as well – a lower [Ca] may require a lower sodium conductance to produce an identical spike excitability. On the other hand, it may point to a fundamental insensitivity of the model behavior to these parameters.

Yet how do we verify the veracity of the parameters in real experiments, where we have no access to their "true" values or to the time traces of the hidden variables? We have already seen that variational annealing may produce very poor estimates of the hidden subspace even when estimates of the measured variables closely mimic observations – recall the estimate of the 10-D Lorenz 1996 system with a single measured variable in Figure 3.5a. What we require is some measure of cross-validation. Actually, this can be done in a natural way for dynamical systems: since we have an estimate of the system's state at the end of the estimation window, we can then generate a future time trace prediction of the system for $t > T$ by integrating Eqs. 4.24 from T forward, incorporating the estimated parameters in the dynamical equations. Since information is transferred from parameters to the measured variables via the dynamics, errors either in

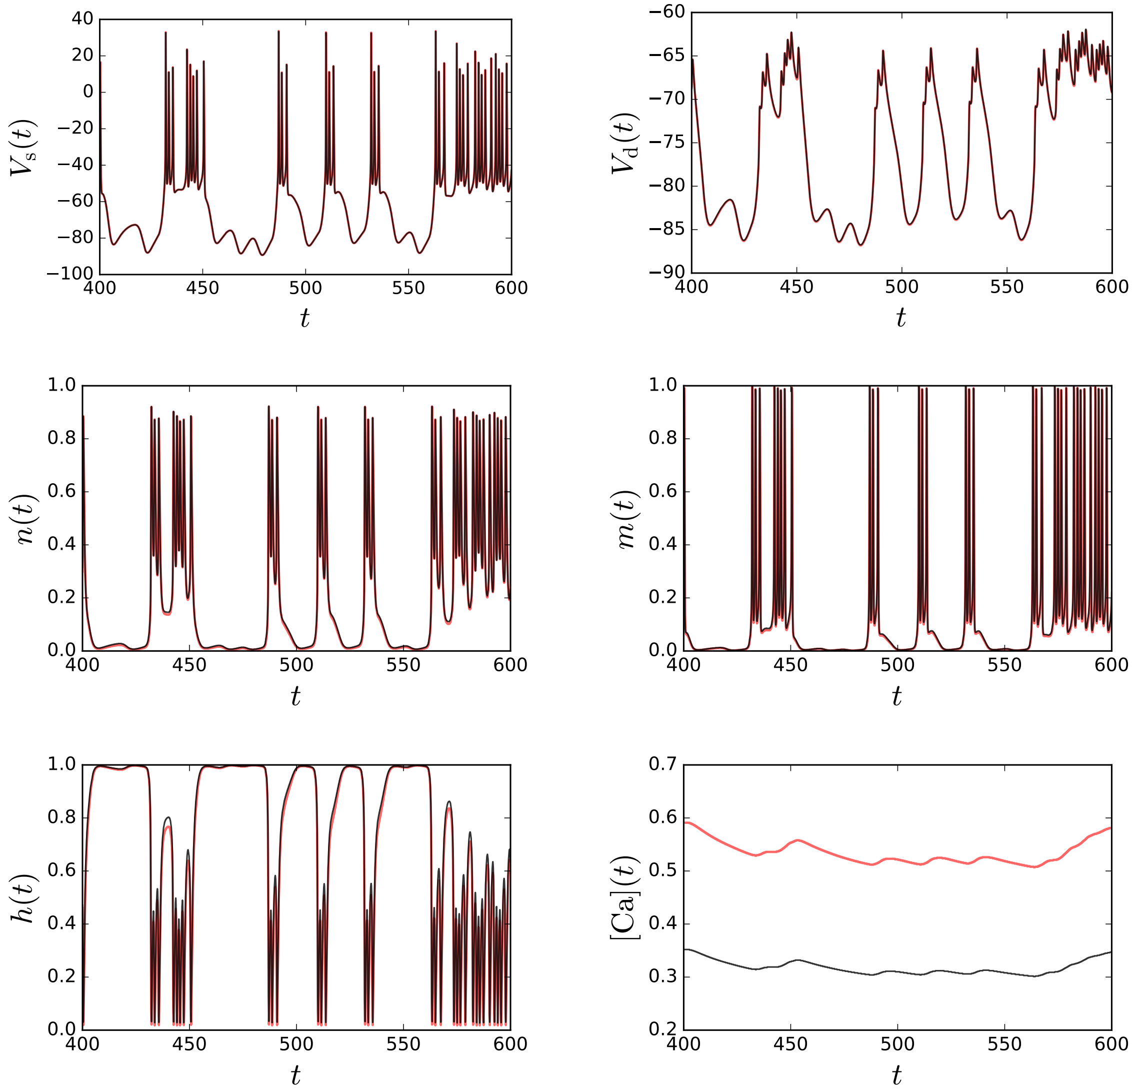

Figure 4.17: Time traces of estimated variables in the full $\rm HVC_{RA}$ model, using the estimated parameters of the reduced model estimation listed in Table 4.6. These parameters include those that were erroneously set to incorrect values: $[\rm Ca]_0$ and $[\rm Ca]_{out}$. The estimates are excellent, with $[\rm Ca](t)$ once again shifted due to the partial degeneracy of the model.

Table 4.8: Estimations of remaining parameters of full HVC$_{\mathrm{RA}}$ model, using only observations of the somatic voltage. The parameters not listed are held fixed at their estimated values from Table 4.6.

parameter	actual	estimated	parameter	actual	estimated
$\bar{g}_{\mathrm{K}}$	120	120.36	θ_m	-30	-30.44
$\bar{g}_{\mathrm{Na}}$	1050	757.81	σ_m	9.5	9.68
E_{Na}	55.0	60.0	$\theta_{\tau,m}$	-	-29.92
τ_{m0}	0.01	7.23e-4	$\sigma_{\tau,m}$	-	-8.24
τ_{m1}	0	5.46e-3	θ_n	-35	-34.28
τ_{m2}	0	0.00	σ_n	10.0	10.8
τ_{n0}	0.10	0.10	$\theta_{\tau,n}$	-27	-32.41
τ_{n1}	0	0.10	$\sigma_{\tau,n}$	-15	-12.41
τ_{n2}	0.50	0.42	θ_h	-45	-43.43
τ_{h0}	0.10	0.12	σ_h	-7.00	-7.05
τ_{h1}	0.00	0.04	$\theta_{\tau,h}$	-40.5	-41.3
τ_{h2}	0.75	0.69	$\sigma_{\tau,h}$	-6.0	-5.0

the final state estimation *or* in the parameters would be reflected in a predicted voltage trace $V_{\mathrm{s}}(t), t > T$ that conflicts with observed voltage recordings under the same stimulus.

To mimic such a validation for the HVA$_{\mathrm{RA}}$ neuron model, let us generate a forward prediction by integrating Eqs. 4.24 from $\hat{\mathbf{x}}^{q_0,\beta_{\max}}(T)$ using the estimated parameters listed in Tables 4.6 and 4.8, under a new injected current. We will generate "future observations" by integrating the true model from its final state $\mathbf{x}(T)$, using the original presumed variables from Table 4.4. Similar noise as above is added to this voltage trace.

A comparison of these two traces is shown in Figure 4.18. The estimated voltage (blue) tracks the synthetic observed voltage trace to excellent accuracy under the influence of this new somatic injected current. This is further indication that the estimated parameters, however inaccurate, are an adequate representation of this model behavior, and that the degeneracies only manifest in absolute [Ca] differences that pose no restrictions on the observed voltage traces.

We should note however that since the stimuli is injected into the cell body, calcium bursts are more difficult to evoke. The possibility remains that the particular stimulus used in these estimations has failed to adequately probe the phase space of the full system, and that the discrepancies in the parameter estimates seen in Tables 4.6 and 4.8 will manifest themselves in response to other stimuli, such as dendritically-injected currents. We test this by repeating the forward prediction comparison illustrated in Figure 4.18 but now using a dendritically-injected current stimulus consisting of pseudo-noisy waveform like that in Figure 4.8a. The

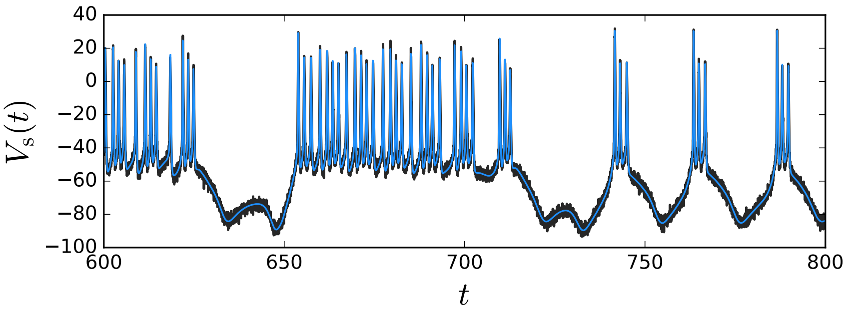

Figure 4.18: Comparison of forward prediction of the somatic voltage of the full two-compartment $\mathrm{HVC_{RA}}$ model using parameter estimates listed in Table 4.6 and 4.8 against the synthetic observations. The prediction is excellent, despite some “incorrectly” estimated parameters. This precision is due to the partial degeneracy of the model, suggesting that a reduced model description may suffice in practice.

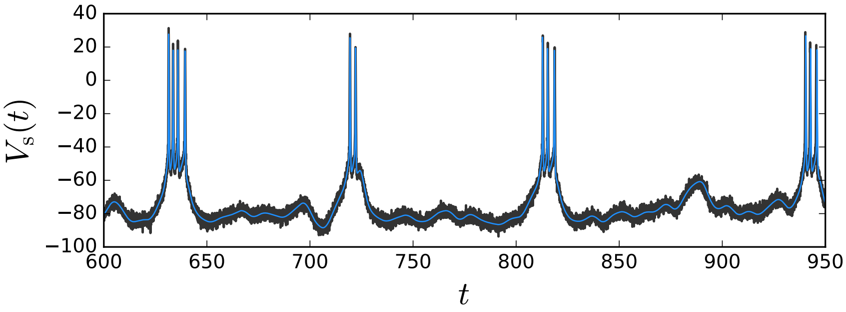

Figure 4.19: As in Figure 4.18, a comparison of the forward prediction of the full $\mathrm{HVC_{RA}}$ model using estimated parameters against the synthetic observations, now in response to noisy dendritic currents (mimicking synaptic stimuli). As before, the traces are estimated to excellent accuracy.

resulting trace is shown in Figure 4.19, and puts our worries to rest.

4.5 Degeneracies in the fast-slow $\mathrm{HVC_{RA}}$ model: A bifurcation study

We have seen that the two-compartment $\mathrm{HVC_{RA}}$ neuron model produces short, stereotyped bursts by modulating fast, spiking dynamics with the slow movement of calcium ions. It is of interest to understand how the apparent degeneracies in the parameter space of the $\mathrm{HVC_{RA}}$ model manifest in its dynamical bifurcations, since the latter are determined by the parameterization in the first place.

To get a handle on this, let us first plot the bifurcation diagram for the model with the original “true” parameters listed in Table 4.4, used to generate the synthetic observations. As in Figure 4.5, we turn off the calcium dynamics ($d[\mathrm{Ca}]/dt = 0$) and treat [Ca] as a bifurcation parameter, assuming a constant injected

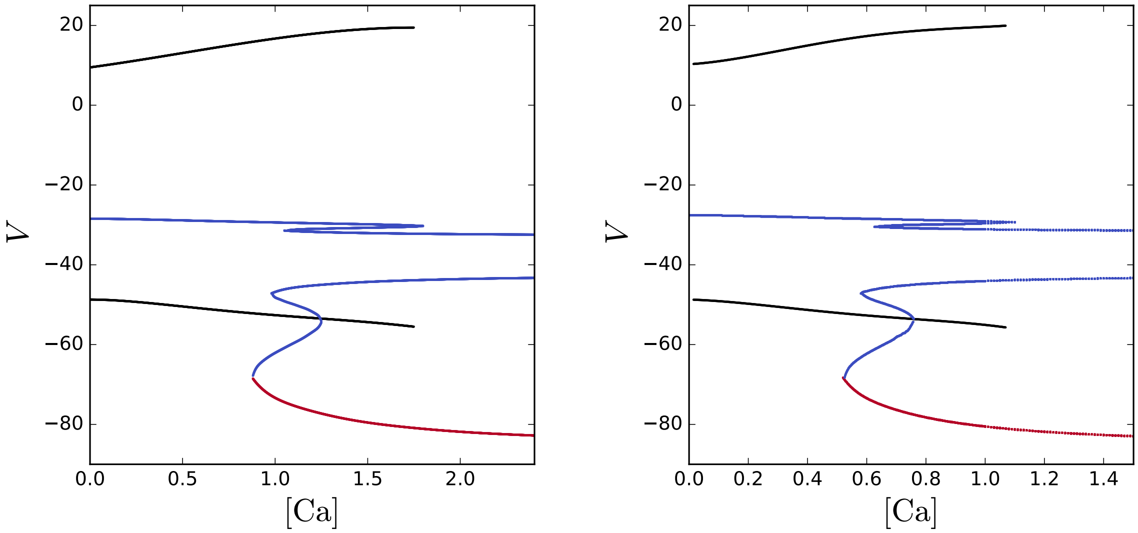

(a) blah

(b) blah

Figure 4.20: Bifurcation diagram showing the somatic voltages of the fixed points of the HVC_{RA} model with [Ca] dynamics shut off, using the parameters listed in Table 4.4 that were used to generate the data. The bifurcation parameter is the fixed intracellular [Ca]. (b) Analogous bifurcation diagram, now generated using the estimated parameters (Table 4.8). The diagram is identical to that in (a), except shifted along the [Ca] axis, in accordance with the shifted calcium traces shown in Figure 4.17.

current of 400 pA into the dendrite. The roots of the nonlinear system of equations, Eqs. 4.24, are found by using a nonlinear solver with 100 different initial guesses. The nonlinear solver will generally converge to the closest root. The stability is tested again by determining the eigenvalues of the Jacobian at these roots; purely negative eigenvalues indicates stability. In addition, the limits of stable limit cycles for a particular [Ca] are determined by integrating the system forward for a few seconds beginning at one of 100 random initial states within the dynamical range of the model.

The resulting diagram shown in Figure 4.20a is a bit more involved than the simpler one-compartment model, but still contains the salient features: a singly-stable spiking state at low [Ca], an intermediate bistable region containing both a spiking limit cycle and a low-voltage stable hyperpolarized rest state, and a singly-stable equilibrium at high [Ca].

How does the model parameterized by the estimated parameters compare? Repeating this procedure with the values from Tables 4.6 and 4.8, we find a bifurcation diagram that appears identical. But note the *x*-axis! The limit cycle bounds are essentially identical, the location of the hyperpolarized equilibria are identical – the only distinction is that comparable features in the two diagrams occur in distinct [Ca] ranges.

This is of course what was anticipated by the estimated trajectories, and it is comforting that global dynamical picture has not fundamentally changed for this parameterization. Rather, it has merely shifted.

The message of this model, and its accompanying bifurcation structure, is that partial degeneracies are not uncommon in neural models. To be sure, the degeneracy is not real – the calcium trace of the estimated and true trajectories are distinct in Figures 4.15 and 4.17. The degeneracy is "partial" in the sense that it occurs only in the measured subspace, and is therefore invisible to any inference procedure, no matter how effective. The value of numerical experiments with synthetic data is that this can be uncovered initially, and can be used to tailor experimental protocols accordingly.

4.6 Inference in neural networks and the challenges ahead

Intracellular voltage traces can be gathered from a handful of neurons at most in real brains, yet these only represent one of millions to trillions of neurons. We expect the partial degeneracy issue from the previous section to arise in a most deflating way: a massive number of combinations of synaptic weights and inputs from the unmeasured neurons would yield the same measured voltage trace, giving us no real information about how the neurons are connected, how quickly the synaptic neurotransmitters act and rectify their inputs, etc. The natural solution is to reduce the complexity of the model description, sometimes down to only a few variables that describe an entire population of neurons. This is, after all, the fundamental decision in neural modeling – what is the correct balance of tractability and complexity?

The theoretical inference tools described in this and the previous chapter rely on a problem formulation that is, while quite general, still not always realized. First, we assume a dynamical model described by either a discrete forward mapping or a set of ordinary differential equations. Dynamics described via partial differential equations in time and space would not readily fit in the framework demanded by the iterative Bayesian integrals; or at least, would require some more thought. This also rules out neural models of the integrate-and-fire (IF) type, which contain continuous dynamics in the subthreshold regime, but require external resetting and latent periods which cannot be described in the framework of differential calculus [Dayan and Abbott, 2005,Izhikevich, 2007]. Yet these are ubiquitous in modeling due to their computational tractability.

One step down the hierarchy of complexity are models that track only the rate of spike generation, $r(t)$, reducing the dimensionality of the model to one variable per neuron. The spike rate modulates via external and synaptic input and decays otherwise, typically described as

$$\frac{dr_i}{dt} = -r_i + \sum_{j=1}^{N} J_{ij}\phi(gr_j), \tag{4.33}$$

where i, j index distinct neurons, g is a "gain" constant for the synaptic input, $\phi(x)$ is some bounded function to prevent runaway excitation. Typically $\phi(x)$ is chosen as the hyperbolic tangent. This simple framework has allowed interesting theoretical results in the large N limit when the synaptic couplings J_{ij} are, rather than precisely defined, Gaussian variables of mean 0 and variance J^2/N. In this limit, it can be shown that stable cycles give way to chaos at a critical value of g, a result that is independent of ϕ or the exact instantiation of J_{ij} [Sompolinksy et al., 1988, van Vreeswijk and Sompolinsky, 1996].

Understanding the richness of animal behavior and consciousness will obviously require models beyond this mean-field limit. Even the multi-layered deep, recurrent neural networks now being tested as models of artificial intelligence will eventually exhaust their predictive power, leaving us no choice but to contend with the true biophysical mechanisms of actual neurons. The information that we can gather from the dynamical transients in this inverse problem will be of fundamental importance.

Chapter 5

Continuous time and the path integral of nonlinear inference

We will spend the remainder of this thesis investigating a number of theoretical aspects of the general Bayesian inference formulation in continuous time. In what follows, it will be necessary to adopt the following perspective: that the estimate state vector is a dynamical system in its own right. This is not entirely unlike the perspective taken by the sequential filter, where the mean of the conditional distribution evolves in discrete time:

$$\mathbf{m}_0 \to \mathbf{m}_1 \to \mathbf{m}_2 \to \ldots \tag{5.1}$$

Yet our approach will be fundamentally different, instead deriving dynamical equations from the stationary points of a functional integral that is *global* in time. In other words, while the evolution of $\mathbf{m}_n$ in the the filtering context relies on information from neighboring times alone, the stationary paths of a global functional integral rely on information from the entire trajectory at once. It is the richness and subtlety of this resulting dynamical system, describing the propagation of an estimate relying on both an underlying dynamical model and on observed data, that gives this thesis its name.

5.1 Stochastic path integrals and the Onsager-Machlup function

To proceed, we first cast the conditional expectation,

$$p(\mathbf{x}_N|\mathbf{Y}_N) \propto p(\mathbf{y}_N|\mathbf{x}_N)\left[\int \prod_{n=0}^{N-1} d\mathbf{x}_n p(\mathbf{x}_{n+1}|\mathbf{x}_n)p(\mathbf{y}_n|\mathbf{x}_n)\right] p(\mathbf{x}_0), \qquad \text{(Eq. 2.5 revisited)}$$

in continuous time. What results is a form presaged more than a half-century ago in the context of stochastic systems; it was then that Onsager and Machlup showed that continuous stochastic processes can be written equivalently as functional path integrals analogous to those in quantum field theory and statistical mechanics [Onsager and Machlup, 1953]. In light of the high-dimensional Bayesian integrals that appear in Eq. 2.5 this seems reasonable – the continuous limit readily gives a path integral measure, and exponential forms for the transition probabilities may produce a quantity akin to a classical action [Feynman and Hibbs, 1965, Abarbanel, 2013].

Such a continuous stochastic process in and of itself is equivalent to the inference problem without data – it is the forward evolution of $\mathbf{x}(t)$ via dynamical drift and stochastic fluctuations. So as a guide to the path integral formulation of $A(\mathbf{X}|\mathbf{Y})$ in the estimation problem, let us first treat this no measurement case. Specifically, we will derive the path integral for $\mathbf{x}(t)$ dynamics described by the Langevin equation

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x},t) + \boldsymbol{\eta}(t), \tag{5.2}$$

which is equivalent to the stochastic differential equation,

$$d\mathbf{x} = \mathbf{f}(\mathbf{x})dt + d\mathbf{w}(t), \tag{5.3}$$

where $\mathbf{w}(t)$ is a Wiener process whose statistics are governed by those of $\boldsymbol{\eta}(t)$ and are homogeneous, namely [Gardiner, 1985]

$$\boldsymbol{\eta}(t) \ \sim \ \mathcal{N}(0,\mathbf{Q}). \tag{5.4}$$

In writing the discrete Bayesian integral, Eq. 2.5, factors of the transition probabilities $p(\mathbf{x}_{n+1}|\mathbf{x}_n)$ appear when the distribution $p(\mathbf{x}_0)$ is propagated forward. To move toward the continuous case, we must exploit the properties of the Wiener process directly. The probability of transition between $\mathbf{x}_n$ and $\mathbf{x}_{n+1}$ at times t_n and t_{n+1}, respectively, is proportional to the Wiener measure of paths $\mathbf{w}_i$ connecting these points [Gardiner, 1985, Chaichian and Demichev, 2001, Chow and Buice, 2015]:

$$p(\mathbf{x}_{n+1}|\mathbf{x}_n) \propto \int d\mathbf{w}_n \exp\left[-(\mathbf{w}_{n+1}-\mathbf{w}_n)^2|_{(\Delta t\mathbf{Q})^{-1}}\right]. \tag{5.5}$$

Extending this from t_0 to t_N, exploiting the Chapman-Kolmogorov equation, and formally taking the continuous limit gives

$$p(\mathbf{x}_N|\mathbf{x}_0) \propto \int \lim_{\substack{N\to\infty \\ \Delta t\to 0}} \left[\prod_{n=0}^{N-1} d\mathbf{w}_n\right] \exp\left[-\sum_{n=0}^{N-1}(\mathbf{w}_{n+1}-\mathbf{w}_n)^2|_{(\Delta t\mathbf{Q})^{-1}}\right] \tag{5.6}$$

where n indicates the discrete times t_n, equally partitioning the interval such that $\Delta t = t_{n+1} - t_n$. For notational ease, we will write $N \to \infty$ only in the limit, the simultaneous limit $\Delta t \to 0$ being implicit. Changing variables from $\mathbf{w}$ to $\mathbf{x}$ demands a choice for discretization scheme which we parameterize by α:

$$\mathbf{x}_{n+1} - \mathbf{x}_n - [\alpha\mathbf{f}(\mathbf{x}_{n+1}) + (1-\alpha)\mathbf{f}(\mathbf{x}_n)]\Delta t = \mathbf{w}_{n+1} - \mathbf{w}_n, \tag{5.7}$$

where $\alpha = 0, 1/2$ correspond to Ito and Stratanovich calculus, respectively. The Jacobian involved in this change of variables is

$$\frac{\partial w_n^a}{\partial x_m^b} = \delta_{mn}\delta_{ab} - \Delta t \sum_{i=1}^{n}\left[\alpha\frac{\partial f_a(\mathbf{x}_i)}{\partial x_m^b}\delta_{m,i} + (1-\alpha)\frac{\partial f_a(\mathbf{x}_{i-1})}{\partial x_m^b}\delta_{m,i-1}\right], \tag{5.8}$$

which is block lower triangular. Its determinant is therefore the product of the determinants of the diagonal blocks:

$$\begin{aligned}
\det\left|\frac{\partial w_n^a}{\partial x_m^b}\right| = \prod_n \det\left|\frac{\partial w_n^a}{\partial x_n^b}\right| &= \prod_n \det\left|\delta_{ab} - \alpha\frac{\partial f_a(\mathbf{x}_n)}{\partial x_n^b}\Delta t\right| \\
&= \prod_n \det\left|\exp[-\alpha\mathbf{J}(\mathbf{x}_n)\Delta t]\right| + \mathscr{O}(\Delta t^2) \\
&= \prod_n \exp\left[-\mathrm{Tr}(\alpha\mathbf{J}(\mathbf{x}_n)\Delta t)\right] + \mathscr{O}(\Delta t^2) \\
&\to \exp\left[-\int dt\, \mathrm{Tr}(\alpha\mathbf{J}(\mathbf{x}))\right] \quad \text{as } N\to\infty
\end{aligned} \tag{5.9}$$

where $J_{ab}(\mathbf{x}') = \frac{\partial f_a(\mathbf{x})}{\partial x_n^b}|_{\mathbf{x}'}$ is the Jacobian of the vector field evaluated at $\mathbf{x}'$. Under the change of variables from the fluctuating variable $\mathbf{w}(t)$ to $\mathbf{x}(t)$, this term will contribute to the exponential in Eq. 5.6, and it appears to impose a dependence of the continuous integral upon the discretization through the presence of α. Interestingly, a careful derivation shows that this dependence actually cancels, so the path integral over $\mathscr{D}\mathbf{x}(t)$ is independent of the discretization scheme [Dekker, 1978, Hunt and Ross, 1981]. Nevertheless, the correct limit of Eq. 5.6 is not given by the naive limit of Eq. 5.3, $\Delta\mathbf{w}/\Delta t \to \dot{\mathbf{x}} - \mathbf{f}(\mathbf{x})$. Let us derive the correct expression specifically for the Ito and Stratanovich cases, the exact details of more general cases can be found elsewhere [Hunt and Ross, 1981].

For Ito discretization ($\alpha = 0$), the change of variable determinant, Eq. 5.9, is unity, while the argument of the exponential in Eq. 5.6 is expressed in $\mathbf{x}$ as:

$$\sum_{n=0}^{N-1} \left(\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\Delta t} - \mathbf{f}(\mathbf{x}_n) \right)^2 \Big|_{\mathbf{Q}^{-1}\Delta t}. \tag{5.10}$$

Expanding this quadratic, two of the terms have straightforward continuous limits, assuming additive noise:

$$\lim_{N\to\infty} \sum_{n=0}^{N-1} \left(\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\Delta t} \right)^T \frac{\Delta t \mathbf{Q}^{-1}}{2} \left(\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\Delta t} \right) = \int dt \dot{\mathbf{x}}^2 |_{\mathbf{Q}^{-1}} \tag{5.11}$$

$$\lim_{N\to\infty} \sum_{n=0}^{N-1} \mathbf{f}^T(\mathbf{x}_n) \frac{\Delta t \mathbf{Q}^{-1}}{2} \mathbf{f}(\mathbf{x}_n) = \int dt \mathbf{f}(\mathbf{x})^2 |_{\mathbf{Q}^{-1}} \tag{5.12}$$

The cross term, on the other hand, must be evaluated using Ito's Lemma [Gardiner, 1985]:

$$\begin{aligned}
\lim_{N\to\infty} \sum_{n=0}^{N-1} (\mathbf{x}_{n+1} - \mathbf{x}_n)^T \mathbf{Q}^{-1} \mathbf{f}(\mathbf{x}_n) &= \lim_{N\to\infty} \sum_{n=0}^{N-1} \Delta\mathbf{x}_n^T \mathbf{Q}^{-1} \mathbf{f}(\mathbf{x}_n) \\
&= \int d\mathbf{x}^T \mathbf{Q}^{-1} \mathbf{f}(\mathbf{x}) - \frac{1}{2} \int dt \, \mathrm{Tr} \left(\mathbf{Q}^{1/2} \mathbf{Q}^{-1} \frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{x}} \mathbf{Q}^{1/2} \right) \\
&= \int dt \dot{\mathbf{x}}^T \mathbf{Q}^{-1} \mathbf{f}(\mathbf{x}) - \frac{1}{2} \int dt \, \mathrm{Tr}(\mathbf{J}(\mathbf{x}))
\end{aligned} \tag{5.13}$$

Combining Eqs. 5.11, 5.12, and 5.13, we see that the quadratic term in the continuous limit of Eq. 5.6 is given by the naive expression *plus* an added divergence:

$$\exp\left[-\int (\mathbf{w}_{n+1} - \mathbf{w}_n)^2 |_{(dt\mathbf{Q})^{-1}} \right] \to \exp\left[-\int dt (\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}))^2 |_{\mathbf{Q}^{-1}} - \frac{1}{2} \int dt \boldsymbol{\nabla} \cdot \mathbf{f}(\mathbf{x}) \right], \tag{5.14}$$

using that the trace of the Jacobian is the divergence of the vector field.

The situation is reversed in the Stratanovich case: the fundamental theorem of Newtonian calculus now holds, so the divergence terms disappears in the quadratic expansion; on the other hand, it appears in the change of variables determinant Eq. 5.9 by virtue of the fact that $\alpha = 1/2$, producing the exact same expression as in the Ito case. Thus, irrespective of the discretization scheme, the stochastic system Eq. 5.3 can be expressed as the functional path integral

$$p(\mathbf{x}(t)|\mathbf{x}(0)) \propto \int \mathscr{D}\mathbf{x} e^{-\int dt \mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t))} = \int \mathscr{D}\mathbf{x} e^{-A[\mathbf{x}(t)]} \tag{5.15}$$

$$\mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t) = (\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}))^2|_{\mathbf{Q}^{-1}} + \frac{1}{2}\nabla \cdot \mathbf{f}(\mathbf{x}). \tag{5.16}$$

The notation $\mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t)$ is not incidental. Eq. 5.15 is analogous to a quantum propagation amplitude between $\mathbf{x}(0)$ and $\mathbf{x}(t)$ in imaginary time, where $A[\mathbf{x}]$ is the action functional and $\mathscr{L}$ is the corresponding Lagrangian known as the Onsager-Machlup function [1] [Onsager and Machlup, 1953,Feynman and Hibbs, 1965]. Before turning to a similar action functional for the situation with measured data, let us first examine in detail the Onsager-Machlup function $\mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t)$ via the quantum analogy.

5.2 The Onsager-Machlup Euler-Lagrange equations

Recall from the discrete time Bayesian formulation that the overriding goal of the estimation problem is the calculation of expected values (and possibly higher moments) of functions defined within the observation window. In the discrete case, we estimated the conditional expectation by working in the Laplace approximation, producing solutions that contributed most heavily to the high-dimensional integral. The message was that when the model precision $\mathbf{Q}^{-1}$ is appreciable, the exponential argument is fractal but there exists a single, extremely dominant minimum. The continuous time analogue, Eq. 5.15, suggests the application of analogous methodologies from classical mechanics and quantum theory, which contain the same mathematical structures. For one, dominant contributions to the path integral arise from stationary points of the action, just as in the discrete case. But here we are not minimizing functions $A(\mathbf{X}|\mathbf{Y})$ over a discrete set of variables $\mathbf{X}$, rather functionals $A[\mathbf{x}(t)]$ over functions $\mathbf{x}(t)$. The solutions to this variational problem are described by the Euler-Lagrange equations, which for the Onsager-Machlup function in Eq. 5.16 produce a

---

[1] In their original paper, Onsager and Machlup did not include the divergence term

set of D second-order ordinary differential equations:

$$\begin{aligned}\frac{d}{dt}\frac{\partial \mathscr{L}}{\partial \dot{\mathbf{x}}} &= \frac{\partial \mathscr{L}}{\partial \mathbf{x}} \\ \mathbf{Q}^{-1}(\ddot{\mathbf{x}} - \mathbf{J}^T\dot{\mathbf{x}}) &= -\mathbf{Q}^{-1}(\mathbf{J}\dot{\mathbf{x}} - \mathbf{J}\mathbf{f}) + \frac{1}{2}\boldsymbol{\nabla}(\boldsymbol{\nabla}\cdot\mathbf{f}(\mathbf{x})) \\ \ddot{\mathbf{x}} &= \mathbf{J}^T\dot{\mathbf{x}} - \mathbf{J}\dot{\mathbf{x}} + \mathbf{J}\mathbf{f} + \frac{1}{2}\mathbf{Q}\boldsymbol{\nabla}(\boldsymbol{\nabla}\cdot\mathbf{f}(\mathbf{x})).\end{aligned} \tag{5.17}$$

More transparently, written in component notation, and as a system of first order ODEs:

$$\begin{aligned}\dot{x}_a &= v_a \\ \dot{v}_a &= \partial_b f_a(\mathbf{x}) v_b - \partial_a f_b(\mathbf{x}) v_b + \partial_a f_b(\mathbf{x}) f_b(\mathbf{x}) + \frac{1}{2} Q_{ab}\partial_b\partial_c f_c(\mathbf{x}),\end{aligned} \tag{5.18}$$

These are Euler-Lagrange equations of the Onsager-Machlup function, which in theory be integrated forward from an initial condition to produce what is effectively the "classical" motion of this stochastic system. There is a fine point in all of this: the Euler-Lagrange equations depend on the vanishing of boundary terms that arise from integration by parts (no sum over a):

$$\delta x_a \frac{\partial \mathscr{L}}{\partial \dot{x}_a}\bigg|_{t=0,T} = 0 \tag{5.19}$$

In the mechanical system, it is assumed that all possible paths pass through known locations at the bounds of the observation window, or $\delta x_a(T) = \delta x_a(0) = 0$. But these boundary conditions cannot be enforced in the estimation problem since the system is only partially observed. Instead, we must enforce that the partial derivative factor vanish, which for the Gaussian error approximation produces $\dot{\mathbf{x}} = \mathbf{f}$; the system obeys the deterministic drift at its boundaries.

Note that since the stochastic system reduces to a deterministic system as the noise variance vanishes ($\mathbf{Q} \to 0$), we should expect that the deterministic limit of the stochastic dynamics would simply reproduce the original dynamics $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ suggested by Eq. 5.3. Before turning to the case with both process noise $\mathbf{Q} \neq 0$ and observed data, let us first acquaint ourselves with some unexpected features of the Euler-Lagrange equations for this case.

5.2.1 Deterministic dynamics: The traceless Jacobian

On the one hand, the solutions of $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ are indeed integral curves of Eqs. 5.18 in the $\mathbf{Q} = 0$ limit:

$$\begin{aligned}\dot{v}_a = \frac{d}{dt} f_a(\mathbf{x}) &= \partial_b f_a(\mathbf{x}) \dot{x}_b \\ &= \partial_b f_a(\mathbf{x}) v_b - \partial_a f_b(\mathbf{x}) v_b + \partial_a f_b(\mathbf{x}) f_b(\mathbf{x}) \Big|_{v_a = f_a}\end{aligned} \tag{5.20}$$

On the other hand, the Euler-Lagrange equations reside in an augmented space $\{x_a, v_a\}$ of dimension $2D$. The path integral devoid of stochasticity has artificially embedded the original nonlinear system of dimension D in a larger manifold. It is important to realize that this embedding is not a mere mathematical redundancy – the enlarged space fundamentally affects the local stability of the original dynamics, $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$. This can be seen immediately via Jacobian of the Euler-Lagrange equations:

$$\delta \dot{x}_a = \delta v_a \tag{5.21}$$

$$\begin{aligned}\delta \dot{v}_a &= v_b \partial_b \partial_c f_a(\mathbf{x}) \delta x_c + \partial_b f_a(\mathbf{x}) \delta v_b \\ &\quad - v_b \partial_a \partial_c f_b(\mathbf{x}) \delta x_c - \partial_a F_b(\mathbf{x}) \delta v_b \\ &\quad + \partial_a f_b(\mathbf{x}) \partial_c f_b(\mathbf{x}) \delta x_c + \partial_a \partial_c f_b(\mathbf{x}) f_b(\mathbf{x}) \delta x_c \\ &= (\delta x_c \text{ terms}) + (\partial_b f_a(\mathbf{x}) - \partial_a f_b(\mathbf{x})) \delta v_b \\ &= \mathbf{K}_{ac} \delta x_c + \mathbf{A}_{ab} v_b.\end{aligned} \tag{5.22}$$

Here, $\mathbf{K}_{ac}$ and $\mathbf{A}_{ab}$ are matrices, with the latter being antisymmetric. The structure of the Jacobian is therefore

$$\mathbf{J}_{\text{EL},\mathbf{Q}=0} = \left[\begin{array}{c|c} \mathbb{O}_{DxD} & \mathbb{1}_{DxD} \\ \hline \mathbf{K}_{ac} & \mathbf{A}_{ab} \end{array}\right] \tag{5.23}$$

which is traceless. Since the trace of the Jacobian is the sum of its eigenvalues, the fact that it vanishes implies that every negative eigenvalue must be paired with a positive one – stable and unstable directions coexist. For purely oscillatory systems, this has no negative ramifications, since the eigenvalues may well all be imaginary. On the other hand, what was a stable limit cycle in the configuration space $\{\mathbf{x}\}$ could now be rendered locally unstable in the Lagrangian coordinates $\{\mathbf{x}, \mathbf{y}\}$. Unless the eigenvalues recross imaginary axis, this instability will persist.

5.2.2 Instability in the classical motion

What is the consequence of the traceless Jacobian? Recall that our general formulation of the nonlinear estimation problem assumes that the system is only sparsely observable – we cannot integrate $\dot{x}_a = f_a(\mathbf{x})$ forward since the initial condition is incomplete. Instead, we may impose a boundary conditions, $\dot{x}_a(T) = f_a(\mathbf{x}(T))$ and attempt to perturb the initial conditions until the boundary conditions are met. In theory, this could be done systematically – it the basis of the shooting technique which will arise again later in this thesis – but in practice, the instabilities exposed by the local Jacobian in Euler-Lagrange coordinates may produce huge changes in $\mathbf{x}(T)$ from tiny nudges in $\mathbf{x}(0)$. In other words, the final state is hypersensitive to changes in the initial state, and successful perturbation would ultimately require precision beyond the limitations of reasonable computation.

This instability exists across a broad range of dynamical systems $\mathbf{f}(\mathbf{x})$ – even along stable limit cycles and equilibria! As an illustration of this, consider the Van der Pol oscillator, a nonlinear system containing a stable limit cycle and an unstable fixed point at the origin [Jordan and Smith, 1999]:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) = \begin{cases} \mu(x_1 - x_1^3/3 - x_2) \\ \frac{1}{\mu}x_1 \end{cases} \tag{5.24}$$

with $\mu = 0.3$. Integrating this system forward from the random initial condition $\mathbf{x}(0) = \mathbf{x}_{\text{init}} = \{2.30, -0.81\}$ generates the phase plot and time trace shown in Fig. 5.1. The stability and periodicity of the limit cycle are evident.

Casting the Van der Pol oscillator into the $4D$ Euler-Lagrange coordinates via Eq. 5.20 produces the following nonlinear differential system:

$$\begin{aligned} \dot{\mathbf{x}} &= \begin{cases} v_1 \\ v_2 \end{cases} \\ \dot{\mathbf{v}} &= \begin{cases} -v_2(\mu + 1/\mu) + \mu^2(1 - x_1^2)(x_1 - x_1^3/3 - x_2) + x_1/\mu^2 \\ v_1(\mu + 1/\mu) - \mu^2(x_1 - x_1^3/3 - x_2) \end{cases} \end{aligned} \tag{5.25}$$

To integrate the system forward in this 4-D space, we initialize the system such that the initial velocities $\mathbf{v}(0)$

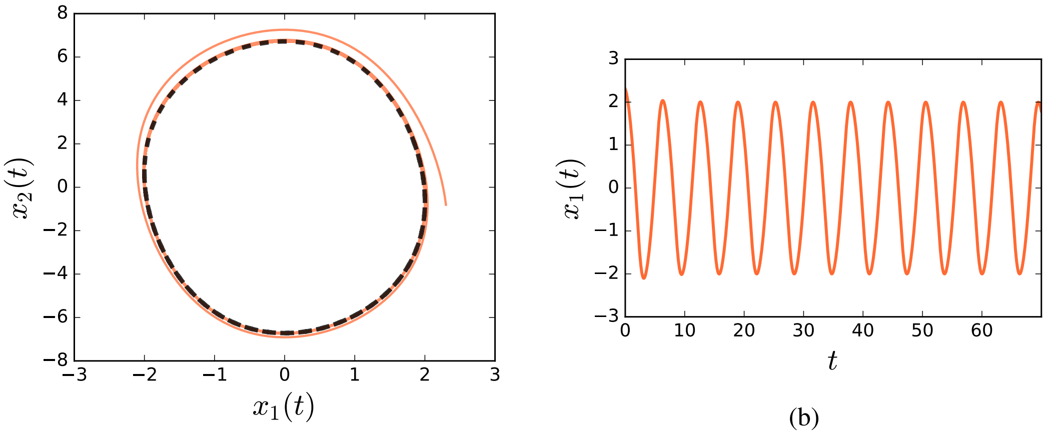

(a)

(b)

Figure 5.1: (a) Phase plot of the Van der Pol oscillator; the stable limit cycle is indicated by the black dotted line. (b) Associated time trace of x_1.

are chosen to place the system *precisely* on the integral curve of the original system:

$$\mathbf{x}(0) = \mathbf{x}_{\text{init}} \tag{5.26}$$

$$\mathbf{v}(0) = \mathbf{f}(\mathbf{x}_{\text{init}}) \tag{5.27}$$

Since, as argued above, $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ solves the enlarged dynamical system, then projecting the integral curves of Eq. 5.25 to the original space $\mathbf{x}(t)$ should in principle reproduce Fig. 5.1a identically.

The issue is that due to discretization errors in numerical integration, the integral curve $\mathbf{x} = \mathbf{f}(\mathbf{x})$ may not be followed exactly, *even from the same initial state*. Figures 5.2a and 5.2b show such a numerical integration using a fourth-order integrator with timestep $\Delta t = 0.01$. The system tracks the limit cycle for some time, but eventually leaves the attractor via motion in the augmented directions v_1 and v_2. This motion is shown explicitly in Fig. 5.3, where the phase portrait now includes a third axis showing the discrepancy between the original system's vector field x-component $f_1(\mathbf{x}(t))$ and the Lagrangian system's independent velocity variable $v_1(t)$. We see $v_1(t)$ slowly deviate from the true dynamical vector field; as the $4D$ system contracts in the stable x_1-x_2 plane towards its limiting attractor, it must expand in the v_1-v_2 directions as a consequence of the traceless Jacobian we found in Eq. 5.23. We cannot precisely satisfy the continuous dynamics, and the system succumbs to the instability in the $\mathbf{v}$ manifold.

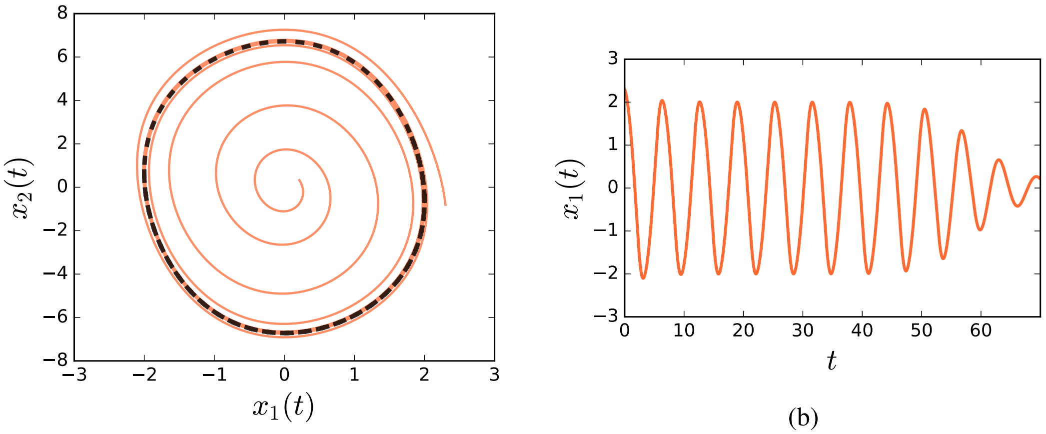

(b)

(a) bs

Figure 5.2: (a) Phase plot of the Van der Pol oscillator as in Figure 5.1, but now in Lagrangian phase space. The limit cycle is indicated by the black dotted line; however, by casting the system into the Lagrangian coordinates and deriving the dynamics from Hamilton's principle, the stability of the limit cycle only exists in the reduced $\mathbf{x}$ coordinates, with associated instability along the $\mathbf{v}$ directions. When the system is initialized at the same point as in Figure 5.1, it attracts to the limit cycle, but eventually succumbs to these instabilities and diverges. (b) Associated time trace of x_1.

5.3 The dynamical formulation of nonlinear inference

Hypersensitivity and persistent instability are expected in chaotic systems. Yet we have seen that this behavior can be artificially concocted in stable, dissipative systems by casting the original dynamics within a variational principle. In the noiseless deterministic system without measurements, this procedure is a formality – the dynamics are already known ($\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$), so couching the dynamics in a functional path integral is an artificial construct. But what is the implication of the Euler-Lagrange equations for the actual itestimation problem, in which there is both stochasticity and data?

Proceeding as before, let us cast the estimation dynamics into a path integral via the same discretization procedure. We first formally promote the observations to a continuous vector, $\mathbf{y}_n \to \mathbf{y}(t)$ and express the conditional integrals $p(\mathbf{y}_n|\mathbf{x}_n)$ in the appropriate continuous limit, as was done for the noiseless case. The conditional expected value of functionals $G(\mathbf{x}(t))$ defined in the observation window from $[0,T]$ is expressed, in analogy to Eq. 2.5, as:

$$\langle G(\mathbf{x}(t))|\mathbf{y}(t)\rangle \propto \lim_{\substack{N\to\infty \\ dt\to 0}} \int \prod_{n=0}^{N-1} d\mathbf{w}_n G(\mathbf{x}_n) p(\mathbf{x}_{n+1}|\mathbf{x}_n) p(\mathbf{y}_n|\mathbf{x}_n) \tag{5.28}$$

Incorporating i) the Wiener transition probability Eq. 5.5, ii) the change of variables Jacobian Eq. 5.7, and

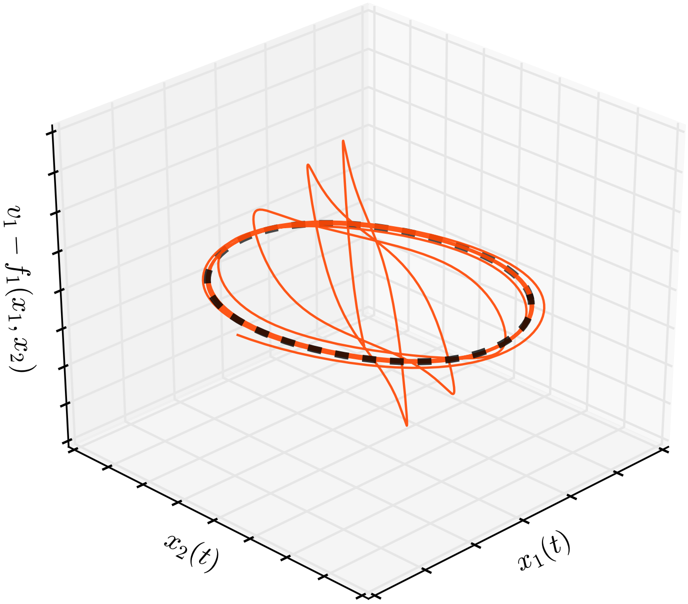

Figure 5.3: Three-dimensional plot showing the unstable manifold of the Van der Pol oscillator when integrated in Lagrangian coordinates. The z-axis is the deviation of the velocity variable v_1 from the value of $f_1(x_1, x_2)$; in $\mathbf{x}$-space, these quantities are identified, but are independent in Lagrangian coordinates. While the deviation begins small as the system converges to the limit cycle, numerical errors compound and eventually succumb the system to unstable directions along v_1 and v_2, producing a deviation of these variables from the vector field of the original system.

iii) the assumption of iid normally-distribution observations, $\mathbf{y}(t) \sim \mathscr{N}(\mathbf{x}(t), \mathbf{R}(t))$, the conditional expectation in the Gaussian error approximation can be written as the path integral

$$\langle G(\mathbf{x}) | \mathbf{y} \rangle = \frac{\int \mathscr{D}\mathbf{x} G(\mathbf{x}) e^{-\int dt \mathscr{L}(\mathbf{x}, \dot{\mathbf{x}}, t)}}{\int \mathscr{D}\mathbf{x} e^{-\int dt \mathscr{L}(\mathbf{x}, \dot{\mathbf{x}}, t)}} \tag{5.29}$$

where

$$\mathscr{L}_{\mathrm{I}}(\mathbf{x}, \dot{\mathbf{x}}, t) = (\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}))^2 \big|_{\mathbf{Q}^{-1}} + (\mathbf{x} - \mathbf{y})^2 \big|_{\mathbf{R}^{-1}} + \frac{1}{2} \nabla \cdot \mathbf{f}(\mathbf{x}), \tag{5.30}$$

and $\mathscr{L}_{\mathrm{I}}$ is notation for the Lagrangian of nonlinear dynamical **I**nference. For completeness, the corresponding Lagrangian in the case of direct observations and uncorrelated process noise, in analogy to Eq. 2.44, is

$$\mathscr{L}_{\mathrm{I}}(\mathbf{x}, \dot{\mathbf{x}}, t) = \sum_{d=1}^{D} \frac{R_{\mathrm{f}}(t,d)}{2} (\dot{x}_d - f_d(\mathbf{x}))^2 + \sum_{d=1}^{D} \frac{R_{\mathrm{m}}(t,d)}{2} (x_d - y_d)^2 + \frac{1}{2} \frac{\partial f_d(\mathbf{x})}{\partial x_d}, \tag{5.31}$$

This expression is the continuous path integral analogue of Eqs. 2.35 and 2.44; it is the naive continuous limit plus a divergence. The quantities $\mathbf{x} = \mathbf{x}(t)$ and $\mathbf{y} = \mathbf{y}(t)$, in contrast to Eq. 2.34, are not ND-dimensional vectors but D-dimensional functions, and the associated Lagrangian $\mathscr{L}(\mathbf{x}, \dot{\mathbf{x}}, t)$ now contains explicit time-dependence through the continuous stream of measurements, $\mathbf{y}(t)$ [2].

As before, we can derive the corresponding Euler-Lagrange equations and investigate the stability of the linearized system. The calculation is straightforward, adding only linear observation terms to the Onsager-Machlup dynamics of Eq. 5.18:

$$\begin{aligned}
\dot{x}_a &= f_a(\mathbf{x}, \mathbf{v}) = v_a \\
\dot{v}_a &= \partial_b f_a(\mathbf{x}) v_b - \partial_a f_b(\mathbf{x}) v_b + \partial_a f_b(\mathbf{x}) f_b(\mathbf{x}) \\
&\quad + Q_{ab} R^{-1}_{bc}(x_c - y_c) + \frac{1}{2} Q_{ab} \partial_b \partial_c f_c(\mathbf{x}),
\end{aligned} \tag{5.32}$$

subject to the boundary conditions:

$$\begin{aligned}
v_a(0) &= f_a(\mathbf{x}(0)) \\
v_a(T) &= f_a(\mathbf{x}(T)).
\end{aligned} \tag{5.33}$$

Since $\dot{v}_a$ is augmented only with x_a terms, the corresponding Jacobian for the linearized dynamics retains the same structure of the matrix Eq. 5.23 – in particular, it is still traceless:

$$\mathbf{J}_{\text{EL}} = \left[\begin{array}{c|c} \mathbb{O}_{DxD} & \mathbb{1}_{DxD} \\ \hline \tilde{\mathbf{K}}_{ac} & \mathbf{A}_{ab} \end{array} \right], \tag{5.34}$$

---

[2] There is a further, subtle discrepancy between the discrete and continuous time Lagrangians, Eqs. 2.44 and 5.31; the R_{f} factor is rescaled by $\sqrt{\Delta t}$, and so R_{f} is now an inverse covariance density.

with

$$\begin{aligned}\mathbf{A}_{ab} &= \partial_b f_a(\mathbf{x}) - \partial_a f_b(\mathbf{x})\\ \tilde{\mathbf{K}}_{ac} &= v_b \partial_b \partial_c f_a(\mathbf{x}) - v_b \partial_a \partial_c f_b(\mathbf{x})\\ &\quad + \partial_a f_b(\mathbf{x}) \partial_c f_b(\mathbf{x}) + \partial_a \partial_c f_b(\mathbf{x}) f_b(\mathbf{x})\\ &\qquad + Q_{ab} R_{bc}^{-1} + \frac{1}{2} Q_{ab} \partial_c \partial_b \partial_d f_d(\mathbf{x}) \end{aligned} \tag{5.35}$$

Thus by combining a stochastic Markov model with a continuous stream of observed data, we have generated a *dynamical model of nonlinear inference*. Put another way, the Euler-Lagrange equations Eq. 5.32 produce an infinite set of integral curves that describe the propagation of the *estimate* through the observation window.

This perspective – an estimate propagating forward in time – is of course familiar: as we have seen, the Kalman filter also provides a prescription for determining optimal estimates iteratively from $t = 0$. We found that the downside of the filter in nonlinear models was its susceptibility to multimodality in the conditional distributions, compounded even further by the presence of a weakly informative prior. The estimation dynamics arising from the variational solutions exhibit instead hypersensitivity, manifested as persistent instability of integral curves.

There is a finer point in the classical analogy: classical equations of motion derived within a variational principle are almost always formulated as initial value problems in practice. Uniqueness of solutions then follow under certain assumptions on smoothness and continuity [Tenenbaum and Pollard, 1985]. On the other hand, the nonlinear inference equations constitute a boundary value problem, which contain no such requirements. Not only are the dynamics hypersensitive to errors in the initial state, but may possess many distinct solutions $\{\hat{\mathbf{x}}(t)\}$, which must then be compared on the basis of their associated action functional values, $\{A[\hat{\mathbf{x}}(t)]\}$.

Together, these subtleties illustrate that while nonlinear inference can be cast formally as a mechanical system with well-defined equations of motion, the usual avenues of approach are far from useful. Chapter 6 is devoted to exploring the stationary solutions of the nonlinear inference action functional in a Hamiltonian formulation, where we will uncover some unexpected features of underlying symmetries and invariants of motion. We will propose an annealing method in canonical coordinates similar to that of Chapter 3, which may offer a partial solution to the issue of persistent stability by leveraging the freedom of motion in a momentum submanifold. Before doing so, we will end this chapter with a brief (and as yet unresolved) analogy with classical electromagnetism.

5.3.1 The electromagnetic analogy

When the state space $\mathbf{x}$ is 3-dimensional, the Lagrangian of nonlinear inference, Eq. 5.30, is familiar from another context in classical physics: a charged particle under the influence of electromagnetic forces. This Lagrangian, for a particle of charge q and mass m moving within electric and magnetic fields is [Goldstein et al., 2002]

$$\mathscr{L}_{\mathrm{EM}} = \frac{\dot{\mathbf{x}}^2}{2m} + q\dot{\mathbf{x}}^T \mathbf{A}(\mathbf{x}) - q\phi(\mathbf{x}) \tag{5.36}$$

where the fields are derived from the potentials $\mathbf{A}$ and ϕ:

$$\begin{aligned} \mathbf{E}(\mathbf{x},t) &= -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t} \\ \mathbf{B}(\mathbf{x},t) &= \nabla \times \mathbf{A} \end{aligned} \tag{5.37}$$

Identifying $\mathscr{L}_{\mathrm{EM}}$ and $\mathscr{L}_{\mathrm{I}}$ creates the following correspondence [3]:

$$\begin{aligned} \mathbf{A} &= \mathbf{A}_{\mathrm{I}}(\mathbf{x},t) = -\mathbf{Q}^{-1}\mathbf{f}(\mathbf{x},t) \\ \phi_{\mathrm{I}} &= \phi_{\mathrm{I}}(\mathbf{x},t) = -\mathbf{f}(\mathbf{x},t)^2\big|_{\mathbf{Q}^{-1}} - (\mathbf{x}-\mathbf{y}(t))^2\big|_{\mathbf{R}^{-1}} - \frac{1}{2}\boldsymbol{\nabla}\cdot\mathbf{f}(\mathbf{x},t), \end{aligned} \tag{5.38}$$

with the charge equal to unity and the mass matrix given by $\mathbf{Q}^{-1}$.

The classical solutions of the path integral of nonlinear inference are equivalent to a charge under the influence of electric and magnetic fields. Interestingly, these fields are not arbitrary: $\mathbf{A}$ and ϕ – and therefore $\mathbf{E}(\mathbf{x},t)$ and $\mathbf{B}(\mathbf{x},t)$ – are interdependent, both functions of the drift dynamics $\mathbf{f}(\mathbf{x},t)$. The electromagnetic analogy will not be explored further in this thesis, being introduced primarily to suggest future work in applying methods of particle tracking and stabilization utilized in plasma physics. In particular, one may exploit the notion of guiding centers, which separate the slow drift in non-uniform $\mathbf{B}$-fields from the fast

---

[3] The dependence of $\mathbf{f}$ upon time is noted explicitly here to indicate how it may lead to explicit time dependence in the potentials.

surrounding oscillations [Northrop, 1961]. Perhaps there is a natural way to separate these fast and time scales in the inference dynamics – to incorporate the faster oscillations in a systematic way, the idea being that these may contribute mostly strongly to estimate destabilization. The beauty of the electromagnetic-inference correspondence is that it reveals the origin of this instability from the action of quantities familiar from classical electrodynamics. The physical picture we should keep in our mind is a delicate balance between magnetic and electric forces, either one of which could rapidly destabilize the motion.

5.3.2 E-M dynamics of the two-dimensional linear system

In a moment, we will revisit all the aspects of the inference problem noted in this chapter in a parallel formalism – Hamiltonian dynamics. Before doing so, let us end this chapter with an illustration of the electromagnetic-inference correspondence in the simplest nontrivial case – linear drift dynamics in two dimensions:

$$\begin{aligned} f_1 &= x_2 \\ f_2 &= -x_1 \\ f_3 &= 0 \end{aligned} \tag{5.39}$$

Assuming uniform model precision $\mathbf{Q} = R_\mathrm{f}\mathbf{I}$ and both components directly observed to a precision of R_m, the associated potentials from Eq. 5.38 are:

$$\begin{aligned} \mathbf{A} &= -R_\mathrm{f} x_2 \mathbf{i} + R_\mathrm{f} x_1 \mathbf{j} \\ \phi &= -\frac{R_\mathrm{f}}{2}(x_1^2 + x_2^2) - \frac{R_\mathrm{m}}{2}(x_1 - y_1(t))^2 - \frac{R_\mathrm{m}}{2}(x_2 - y_2(t))^2 \end{aligned} \tag{5.40}$$

corresponding to an upward uniform **B**-field with planar, outwardly electric fields perturbed *in time* by observed data – the associated 'Inference Lorenz force' $\mathbf{f}_\mathrm{I}$ is:

$$\mathbf{f}_\mathrm{I} = \mathbf{f}_\mathbf{B} + \mathbf{f}_\mathbf{E} = [R_\mathrm{f}(2v_y + x_1) + R_\mathrm{m}(x_1 - y_1(t))]\mathbf{i} + [-R_\mathrm{f}(2v_x + x_2) + R_\mathrm{m}(x_2 - y_2(t))]\mathbf{i} \tag{5.41}$$

This system is only oscillatory in the unobserved limit $R_\mathrm{m} = 0$, and only along curves initiated at the exact dynamics $f_i = v_i$. Otherwise, it contains unstable spirals; that is, motion in the radial direction via the **E**-field causes oscillatory motion directed by the upward **B**-field. This is in stark contrast to the estimate propagation in the sequential filter, which is under appropriate assumptions continuously corrected by incoming data in a stable and convergent manner.

Chapter 6

A Hamiltonian formulation of nonlinear dynamical inference

6.1 Boundary value problems in optimal control

The previous chapter uncovered connections between estimate propagation and mechanical systems. Let us pursue the analogy a bit further. In classical mechanics, Lagrangian and Hamiltonian formulations equivalently solve a mechanical problem. *Equivalently* in the sense that the integral curves of the N second-order E-L equations corresponds to those of $2N$ first-order Hamiltonian equations. In the Chapter 5, we developed a path integral and Lagrangian theory for the nonlinear dynamical inference problem, but found that while the integral curves may be well-defined, they may not be solvable in practice due to persistent instabilities – even if instability is absent in the dynamical model itself. The central distinction between mechanics and dynamical inference is that the former is naturally framed as an initial value problem, for which solutions are unique and well-posed, while the latter can only be framed as a boundary value problem since the initial conditions are incomplete.

Such boundary value problems are well-known in the theory of optimal control [Liberzon, 2012, Betts, 2010]. Optimal control problems seek a control vector $\mathbf{u}(t)$ that optimally drives a dynamical state vector $\mathbf{x}(t)$ from an initial point to a final point subject to boundary constraints, while also minimizing a functional $J[\mathbf{x}(t),\mathbf{u}(t)]$ along the full path trajectory. Consider, for example the problem of attempting to drive between San Diego and Los Angeles using the least amount of gas, given that mpg ratings and viscous air drag are strongly speed-dependent, and given that by some miracle the I-5 is clear. In this scenario, i) the control variable $\mathbf{u}(t)$ to be determined is our weight on the accelerator pedal, ii) the dynamical equations connects the speed $\mathbf{x}(t)$ to this force and others, and iii) the cost J depends on the instantaneous fuel usage rate, which is in turn a function of the speed $\mathbf{x}(t)$. This problem is expressed mathematically as

$$\text{minimize } J[\mathbf{x}(t),\mathbf{u}(t)] = \Phi[\mathbf{x}(t_\mathrm{i}),\mathbf{x}(t_\mathrm{f}),t_\mathrm{i},t_\mathrm{f}] + \int_{t_\mathrm{i}}^{t_\mathrm{f}} \mathscr{L}(\mathbf{x}(t),\mathbf{u}(t))dt \tag{6.1}$$

subject to

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},\mathbf{u},t)$$

$$g_\mathrm{i} = \phi_\mathrm{i}(\mathbf{x}(t_\mathrm{i}),\mathbf{u}(t_\mathrm{i})) \quad \text{and} \quad g_\mathrm{f} = \phi_\mathrm{f}(\mathbf{x}(t_\mathrm{f}),\mathbf{u}(t_\mathrm{f})). \tag{6.2}$$

The extremization of the functional J is carried out by defining an augmented performance functional $H[\mathbf{x},\mathbf{u},\boldsymbol{\lambda}]$, which incorporates the dynamical equations $\dot{\mathbf{x}} = \mathbf{f}$ via Lagrange multipliers $\boldsymbol{\lambda}$, and then solve for stationary variations $\delta H = 0$. The notation H is no accident, as the variational principle produces a set of Hamilton's equations, one of which is a restatement of the state dynamics, the other of which is dynamics of the Lagrange multipliers (or adjoints, or momenta). There is a final equation expressing the optimality of the control in terms of H. In some cases, this expression can be used to express, *explicitly*, $\mathbf{u}$ as a function of $\mathbf{x}$ and $\boldsymbol{\lambda}$. When this is true, the result is a fully Hamiltonian problem in $\mathbf{x}$ and $\boldsymbol{\lambda}$.

6.1.1 Instabilities in Hamiltonian dynamics and the shooting method

Finer details and complications of the optimal control formulation, particularly for non-smooth controls, will not be discussed in this thesis. What is critical for now is that the mathematical structure of this problem resembles that of dynamical inference; even if the $\mathbf{x}(t_\mathrm{i})$ are specified, the set of coupled ODEs still form a boundary value problem since the initial values of $\boldsymbol{\lambda}$ are unknown [Betts, 2010].

A common and intuitive method to solve this boundary-value problem is the so-called shooting technique and its variants [Betts, 2010]. In this scheme, the unknown initial conditions are guessed ($\lambda(t_\mathrm{i})$) and Hamilton's equations are solved forward in time. The error in the presumed boundary conditions at the final point, $\varepsilon \equiv g_\mathrm{f} - \phi_\mathrm{f}$, is evaluated. When ε is above some tolerance, the guessed initial condition is adjusted, and the integration is repeated. This procedure is continued accordingly until the error falls below a specified threshold.

This reasonably logical process is complicated by the presence of instabilities entirely analogous to those encountered in the previous chapter. These instabilities are a direct consequence of the fact that the problem is Hamiltonian. Explicitly, for Hamiltonians systems,

$$\begin{aligned}\dot{\mathbf{q}} &= \frac{\partial H}{\partial \mathbf{p}} \\ \dot{\mathbf{p}} &= -\frac{\partial H}{\partial \mathbf{q}},\end{aligned} \tag{6.3}$$

whereby the Jacobian matrix is

$$\begin{aligned}\mathbf{J} &= \begin{bmatrix} \frac{\partial \dot{\mathbf{q}}}{\partial \mathbf{q}} & \frac{\partial \dot{\mathbf{q}}}{\partial \mathbf{p}} \\ \frac{\partial \dot{\mathbf{p}}}{\partial \mathbf{q}} & \frac{\partial \dot{\mathbf{p}}}{\partial \mathbf{p}} \end{bmatrix} \\ &= \begin{bmatrix} \frac{\partial^2 H}{\partial \mathbf{p} \partial \mathbf{q}} & \frac{\partial^2 H}{\partial \mathbf{p} \partial \mathbf{p}} \\ -\frac{\partial^2 H}{\partial \mathbf{q} \partial \mathbf{q}} & -\frac{\partial^2 H}{\partial \mathbf{q} \partial \mathbf{p}} \end{bmatrix},\end{aligned} \tag{6.4}$$

whose trace vanishes. That is, since differential area projections $|d\mathbf{p}_a \times d\mathbf{q}_a|$ are conserved along Hamiltonian flows, any initial small ball of initial conditions that contracts over time in one direction must expand in another – exactly what we found in the dynamics of nonlinear inference [Betts, 2010].

6.1.2 Multiple shooting

Obviously the problems associated with hypersensitivity to initial conditions compound as the time window lengthens. One way to minimize the effect of initial-state sensitivity is by partitioning t_f interval into N smaller segments, called phases, each individually subject to a shooting requirement. The resulting trajectories are then strung together to a continuous trajectory by identifying the resulting state of the $k-1$ segment and the initializing state of the k segment. In practice, these N conditions ($N-1$ for the the inner phases, plus the terminal boundary condition ϕ_f) are incorporated into a nonlinear program as constraints [Betts, 2010]. This procedure often eliminates or reduces many of the issues with hypersensitivity found in shooting, albeit at the introduction of extra variables to be optimized – each phase now needs its own free state vector $\mathbf{x}_k$.

In multiple shooting, the optimization cost function consists of the original cost function A plus Lagrange multiplier terms enforcing the constraints and boundary conditions. Assuming that a Newton-type nonlinear programming algorithm is used to minimize this function iteratively, the associated Hessian will include second derivatives of A, as well as the Jacobian of the constraint equations. However, multiple shooting requires distinct state vectors in *each* phase, so the associated Hessian has ballooned from D^2 elements, in the case of indirect single-pass shooting, to $(ND)^2$ elements.

Fortunately, the constraint equations $\mathbf{c}_k$ at the phase boundaries are only functions of the state variables on either side of the boundary, $\mathbf{x}_{k-1}$ and $\mathbf{x}_k$. The Jacobian $J_{jk} = \partial_j c_k$ is therefore block tridiagonal at most, with blocks of size DxD. Thus, the number of nonzero elements in the Hessian scales not with N^2 but rather linearly with N. This reduces the runtime of the optimization protocols considerably, an aspect exploited by modern methods in nonlinear optimization for efficient calculation, permitting applications with $10^5 - 10^6$ variables and constraints. For this reason, multiple shooting is a preferred method in optimal control applications where sensitivity and nonlinearity are common [Betts, 2010].

Yet, for heavily nonlinear or chaotic dynamics, there is the matter of spurious minima: if the optimization procedure begins with a poor initial guess, then the extremum satisfying the constraint conditions may only be local in the space of candidate paths. This is essentially a consequence of the lack of general uniqueness theorems for two-point boundary value problems of nonlinear ODEs. It is something we have also seen play out several times through the course of this thesis.

As a simple illustration, consider the following Hamiltonian system $\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}) = \mathbf{f}(x_1, x_2, p_1, p_2)$:

$$
\begin{aligned}
\dot{x}_1 &= p_1 \\
\dot{x}_2 &= p_2 \\
\dot{p}_1 &= -x_1 - 2x_1x_2^2 \\
\dot{p}_2 &= -x_2 - 2x_2x_1^2,
\end{aligned}
\tag{6.5}
$$

which is a 2D harmonic oscillator with nonlinear coupling, known to be chaotic for a wide array of initial conditions [Sprott, 2010]. As a numerical experiment, let us integrate the system forward for 2001 timesteps of $\Delta t = 0.01$, beginning from a given $\{\mathbf{x}(0), \mathbf{p}(0)\}$ to generate the "known" partial boundary condition $\mathbf{x}(T)$. Then we attempt to re-estimate this system given only the initial and final states $\mathbf{x}(0)$ and $\mathbf{x}(T)$, using multiple shooting with $N = 200$ partitions of size $10\Delta t$. The variables $\mathbf{Z} \equiv \{\mathbf{z}_1, ...\mathbf{z}_N\}$ are the arguments of the cost function $A(\mathbf{Z})$ which quantifies the square error between $\mathbf{z}_{n+1}$ and the forward shooting of $\mathbf{z}_n$, in addition to

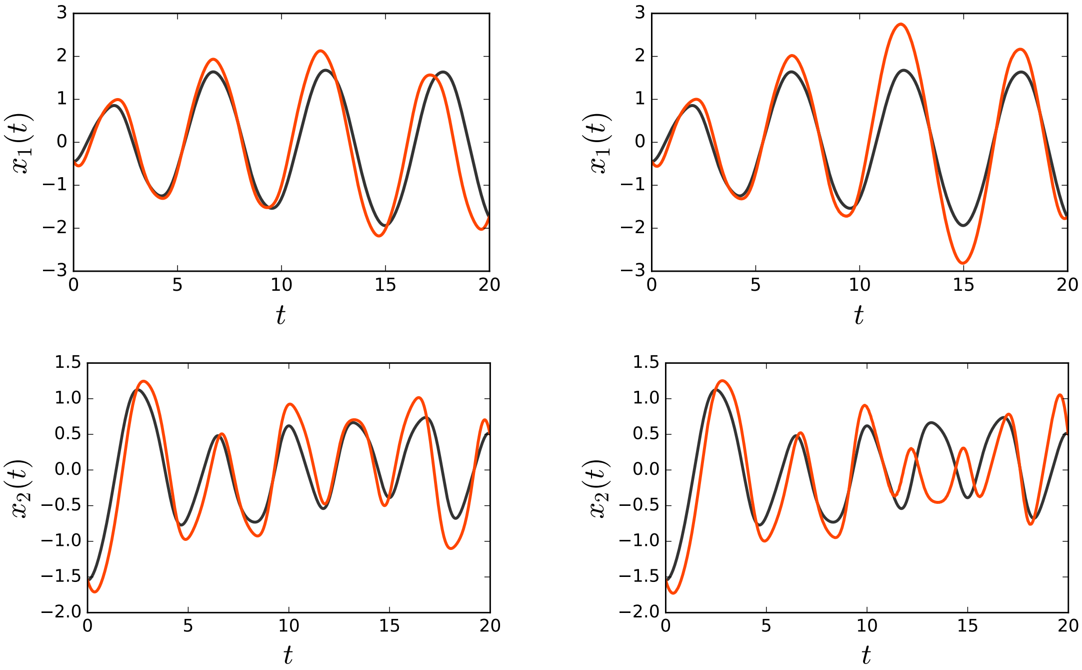

(a)

(b)

Figure 6.1: (a) Estimated trajectory (orange) versus true trajectory (black) of two variables of the the four-dimensional chaotic system, Eqs. 6.5, using multiple shooting. Despite satisfying the boundary conditions and returning a low value of the cost function $A(\mathbf{Z})$, near machine precision, the estimate is poor. (b) Another estimated trajectory satisfying the boundary conditions with low error cost. The estimate is worse than (a).

satisfying the boundary conditions:

$$A(\mathbf{Z}) = \sum_{n=1}^{N-1} (\mathbf{z}_{n+1} - \mathbf{F}_{\mathrm{E}}^{10}(\mathbf{z}_n))^2 + (\mathbf{x}_1 - \mathbf{x}(0))^2 + (\mathbf{x}(T) - \mathbf{P}_x \mathbf{F}_{\mathrm{E}}^{10}(\mathbf{z}_N))^2, \tag{6.6}$$

where $\mathbf{P}_x$ is the projection matrix onto the $\mathbf{x}$ components, $\mathbf{F}_{\mathrm{E}}$ is the Euler step approximation of the dynamics $\mathbf{f}$, and $\mathbf{F}_{\mathbf{E}}^2(\mathbf{x}_n) = (\mathbf{F}_{\mathrm{E}} \circ \mathbf{F}_{\mathrm{E}})(\mathbf{x}_n)$, etc. Running this nonlinear program with a randomly chosen initialization of the optimization often produces the "true" trajectory first generating the boundary conditions. Yet, as seen in Figure 6.1, along with this trajectory are found several other "spurious" trajectories nevertheless satisfying the boundary conditions, minimizing A to the same level of error near machine precision. Though this model is chaotic, it is small – only four dimensions – and the shooting windows are brief; even for this relatively benign case, the lack of a strong prior holds demonstrable consequences.

6.1.3 Collocation

Multiple shooting is often used within a formulation of optimal control known as the *indirect method*. Here, the variational equations solving the optimal control problem – typically a set of Hamilton's equations – are solved. The general idea in indirect methods, as noted in [Betts, 2010] is "optimize then discretize": Hamilton's equations are explicitly derived from a variational principle, the control is inverted into a function of **x** and **p**, and these equations are discretized and solved in a nonlinear program. An alternate formulation, more amenable in the presence of state and control constraints are *direct* methods, which optimize a discretization of the cost function itself. In direct methods, neither Hamilton's equations nor an explicit equation for optimum control is found. As such, these methods are typically much more robust when the state or control constraints are active and/or the optimal control condition cannot be inverted.

One such method is *direct collocation*, or *direct transcription*. The cost function is discretized at collocated points with interpolating polynomials, these interpolating variables then used as the free variables in the optimization search. Clearly, variational annealing fits in this class. Due to the local manner in which the dynamical equations are discretized, the resulting Jacobian and Hessian matrices in collocation schemes are again sparse and banded, making them rather attractive for higher-dimensional problems.

Direct collocation is mentioned here only for sake of completeness, and further discussion of the compromises inherent in various methods in the context of optimal control can be found elsewhere [Kirk, 1970,Betts, 1998,Betts, 2010,Rao, 2010,Liberzon, 2012]. Let us now turn to the formulation of the nonlinear inference problem in the Hamiltonian picture, returning to appropriate methods of solution following a careful discussion of the issues at hand.

6.2 Toward a Hamiltonian formulation of nonlinear inference

The optimal control theory and methods of solution have thus far not been described mathematically in any heavy detail. Nonetheless, the primary takeaway is the following: boundary value problems in optimal control can often be cast explicitly as Hamiltonian systems, which are inherently unstable. To some extent hypersensitivity can be addressed with multiple shooting techniques, though a good initial guess is still required. Collocation offers a further avenue of approach, though it is not without its limitations.

In Chapter 4, we found that the nonlinear dynamical inference problem admits a natural Lagrangian function, thereby casting it as a dynamical system in its own right. But the resulting differential system was not amenable to practical computation, even with known initial conditions. It is of interest to see if the Hamiltonian formulation, though equivalent in the continuum limit, is perhaps more benign in a practical sense.

To this end, let us begin with the continuous Lagrangian of the inference formulation that we derived in Chapter 4,

$$\mathscr{L}_{\mathrm{I}}(\mathbf{x},\dot{\mathbf{x}},t) = (\dot{\mathbf{x}}-\mathbf{f}(\mathbf{x}))^2\big|_{\mathbf{Q}^{-1}} + (\mathbf{x}-\mathbf{y})^2\big|_{\mathbf{R}^{-1}} + \frac{1}{2}\nabla\cdot\mathbf{f}(\mathbf{x}). \tag{5.30 revisited}$$

The Hamiltonian is derived from $\mathscr{L}_{\mathrm{I}}$ through a Legendre transformation,

$$H(\mathbf{x},\mathbf{p}) = \mathbf{p}\cdot\dot{\mathbf{x}} - \mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t)\Big|_{\dot{\mathbf{x}}(\mathbf{p},\mathbf{x})}, \tag{6.7}$$

where $\mathbf{p}$ is defined by $\mathbf{p} \equiv \frac{\partial\mathscr{L}}{\partial\dot{\mathbf{x}}}$. An explicit representation of H requires that this expression can be inverted to express $\dot{\mathbf{x}}$ *explicitly* as a function of $\mathbf{p}$ and $\mathbf{x}$. In practice, this may not always be possible if $\mathscr{L}$ is sufficiently nonlinear in its velocity variables.

In the Gaussian error approximation, this inversion poses no issues. Applying the Legendre transform to $\mathscr{L}_{\mathrm{I}}$ gives the Hamiltonian of nonlinear **I**nference:

$$H_{\mathrm{I}}(\mathbf{x},\mathbf{p}) = \mathbf{p}^2\big|_{\mathbf{Q}} - (\mathbf{x}-\mathbf{y})^2\big|_{\mathbf{R}^{-1}} + \mathbf{p}\cdot\mathbf{f}(\mathbf{x}) - \frac{1}{2}\nabla\cdot\mathbf{f}(\mathbf{x}), \tag{6.8}$$

or in component form, assuming direct observations and uncorrelated, constant process noise (in analogy to Eq. 5.31),

$$H_{\mathrm{I}}(\mathbf{x},\mathbf{p}) = \frac{p_a^2}{2R_{\mathrm{f}}(a)} - \frac{R_{\mathrm{m}}(t,a)}{2}(x_a-y_a)^2 + p_a f_a(\mathbf{x}) - \frac{1}{2}\frac{\partial f_a(\mathbf{x})}{\partial x_a}. \tag{6.9}$$

We will use the component notation exclusively from here forward, as the vector notation is somewhat opaque when writing the equations of motion. The dynamics are then governed by the associated Hamilton's equations of motion:

$$
\begin{aligned}
\frac{dx_a}{dt} &= \frac{\partial H_{\mathrm{I}}}{\partial p_a} = f_a(\mathbf{x}) + \frac{p_a}{R_{\mathrm{f}}(a)} \\
-\frac{dp_a}{dt} &= \frac{\partial H_{\mathrm{I}}}{\partial x_a} = p_b \frac{\partial f_b(\mathbf{x})}{\partial x_a} - R_{\mathrm{m}}(a,t) \cdot (x_a - y_a) - \frac{1}{2} \frac{\partial^2 f_b}{\partial x_a \partial x_b}.
\end{aligned}
\tag{6.10}
$$

In this problem, the indeterminacy of the initial state, partially on the measured subspace and certainly on the unmeasured one, implies that boundary conditions cannot be enforced on $\mathbf{x}$. Conversely, having enforced $\dot{\mathbf{x}}(t_{\mathrm{i}}) = \mathbf{f}(\mathbf{x}(t_{\mathrm{i}}))$ and $\dot{\mathbf{x}}(t_{\mathrm{f}}) = \mathbf{f}(\mathbf{x}(t_{\mathrm{f}}))$ in the Lagrangian formulation, Eqs. 6.10 suggest the vanishing of $\mathbf{p}$ as the natural analogue in the Hamiltonian formulation:

$$
p_a(0) = p_a(T) = 0 \tag{6.11}
$$

As discussed above, while multiple shooting may solve some of the sensitivity issues of single-pass shooting, spurious trajectories may still arise. We have seen fractal cost surfaces arise in the context of variational annealing with chaotic dynamics, so the appearance of several paths satisfying this two-point Hamiltonian boundary value problem would not be wholly unexpected.

In time, we will present a method to partially address this problem in the variational context. Before doing so, let us first investigate a fundamental invariant of Hamiltonian systems, and the subtleties therein. This invariant is the symplectic bilinear form, whose preservation is responsible for Liouville's Theorem. Somewhat less well known in the physics community are the failings of this symmetry in numerical computation, which will be central in the dynamical problem of nonlinear inference.

6.3 Symplectic invariants in Hamiltonian systems

Phase space volume $\prod_d dq_d dp_d$ is conserved along Hamiltonian flows as a consequence of Liouville's Theorem. In fact there is a more fundamental conserved quantity, the symplectic bilinear form, which for a $2D$ Hamiltonian space is

$$
\omega = \sum_{d=1}^{D} dq_d \wedge dp_d, \tag{6.12}
$$

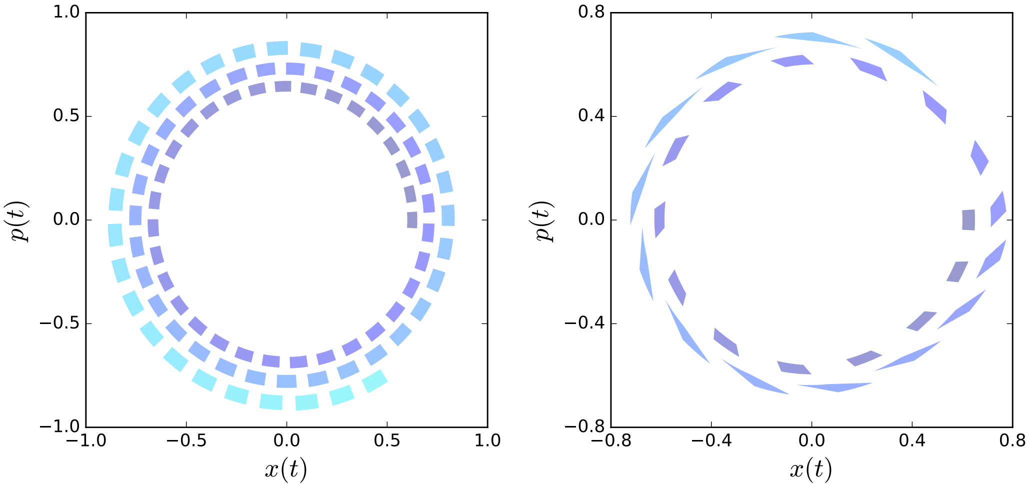

(a) (b)

Figure 6.2: (a) Evolution in phase space of a small region of initial conditions of a $2D$ harmonic oscillator, solved using the forward Euler step numerical scheme. The phase space areas grow in time, contradicting Liouville's Theorem for Hamiltonian systems. (b) Similar forward integration using forward Euler step, now for a nonlinear oscillator system. While the deformation of phase space areas is consistent with Liouville's Theorem, the increase in area magnitude is not.

where the wedge product ω between two differential one forms is the signed area vector of the parallelogram they span. Phase space conservation follows from Eq. 6.12, simply by taking D wedge products of ω with itself.

Manifolds which contain such a conserved quantity are known as symplectic manifolds; equipping these manifolds with an energy function, i.e. the Hamiltonian $H(q_a, p_a)$, then naturally produces a vector field $\mathbf{v} = (\dot{q}_a, \dot{p}_a)$ via Hamilton's equations [Wendlandt and Marsden, 1997, Marsden and West, 2001, Goldstein et al., 2002]. The symplectic form is then preserved along the integral curves of this vector field. Curiously, this symmetry is quite generally violated outside of the continuum limit. If the Hamiltonian system is integrated forward *numerically*, bilinear forms and differential phase space area accrue errors, *whatever* the integration timestep. This is not a side effect of nonlinearity, occurring even in linear dynamics. To illustrate this, consider the simple harmonic oscillator in Hamiltonian coordinates with unit frequency,

$$\dot{x} = p \qquad , \qquad \dot{p} = -x, \tag{6.13}$$

integrated forward for $N = 450$ timesteps using an explicit Euler scheme with timestep $\Delta t = 0.04$:

$$\begin{aligned} x_{n+1} &= x_n + p_n \Delta t \\ p_{n+1} &= p_n + x_n \Delta t. \end{aligned} \tag{6.14}$$

Figure 6.2a shows the evolution of a small region of initial conditions centered at $[x(0), p(0)] = [0.62, 0]$ with this numerical integrator, plotted every 5th step for ease of visualization. Symplecticity dictates that while the region may deform, its area should remain invariant; instead the expected circular orbit gradually gives way to an unbounded spiral, the region steadily growing in size. Figure 6.2b shows the same phenomenon for a nonlinear system with quartic potential, integrated for $N = 270$ steps of size $\Delta t = 0.05$:

$$\begin{aligned} H(x,p) &= \frac{p^2}{2} + \frac{x^2}{10} - \frac{x^3}{15} + \frac{x^4}{4} \\ \dot{x} &= p \\ \dot{p} &= -\frac{x}{5} + \frac{x^2}{5} - x^3. \end{aligned} \tag{6.15}$$

These plots appear to contradict the fact that the time evolution of Hamiltonian systems is a symplectomorphism, the loss of this fundamental symmetry arising entirely from the discrete numerical integration scheme. Thus, numerical integration by an explicit Euler scheme does not preserve the symplectic invariant ω. Is this an unavoidable byproduct of numerical integration? Are there numerical integrators that can preserve these invariants?

6.3.1 Symplectic integrators

Many, in fact. Such *symplectic integrators* preserve ω exactly at every step, despite the unavoidable presence of errors in the trajectories themselves [Hairer et al., 2006]. One example, applicable to Newtonian systems, is the *leapfrog* or Verlet method, of which one variant is:

$$\begin{aligned} \mathbf{p}_{n+1/2} &= \mathbf{p}_n + \frac{dt}{2}\mathbf{f}(\mathbf{x}_n) \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + dt\,\mathbf{p}_{n+1/2} \\ \mathbf{p}_{n+1} &= \mathbf{p}_{n+1/2} + \frac{dt}{2}\mathbf{f}(\mathbf{x}_{n+1}), \end{aligned} \tag{6.16}$$

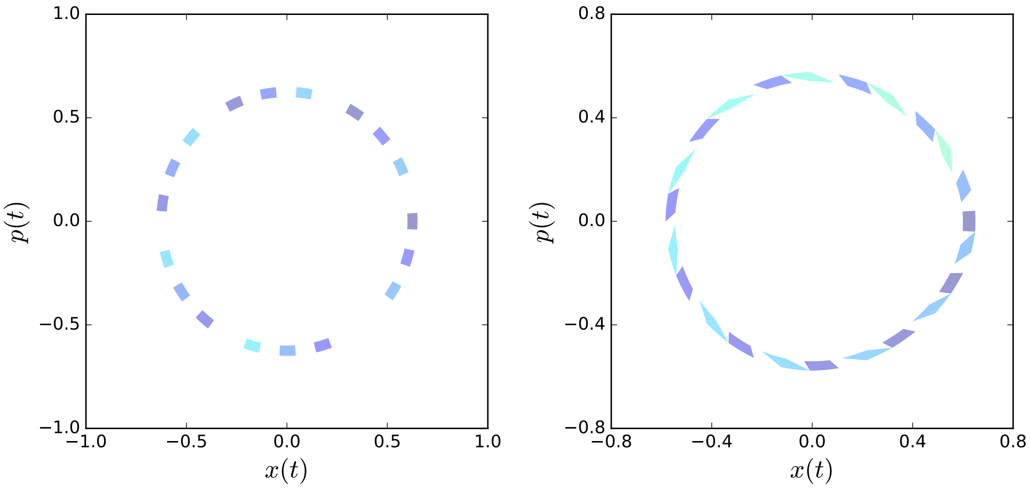

(a) (b)

Figure 6.3: Evolution of the same systems as in Figure 6.2, now using the symplectic leapfrog integration scheme. While the phase space areas deform in the nonlinear case, their magnitudes are nonetheless preserved in accordance with Liouville's Theorem.

assuming unity mass. Evolving the linear and nonlinear systems Eqs. 6.13 and 6.15 with the leapfrog integrator produces the trajectories shown in Figure 6.3 (The figures are plotted at larger increments than in Figure 6.2 for better visualization). In both cases, the differential areas may contort, but are preserved step-by-step.

Proving that an integrator such as the leapfrog method is symplectic is straightforward. To begin, note that symplectomorphisms (*canonical transforms* in physics) are defined as transformations that leave ω invariant, and within the language of differential forms, this invariance can be shown to follow naturally from the particular form of Hamilton's equations. In other words, if a coordinate transformation is shown to preserve Hamilton's equations, it is indeed canonical.

A more compact condition on canonical transformations can be derived using the invariance of Hamilton's equations. Consider the Jacobian matrix of transformation $\mathbf{M}$ from $\boldsymbol{\eta} = \{\mathbf{x}, \mathbf{p}\}$ to $\boldsymbol{\xi} = \{\mathbf{x}', \mathbf{p}'\}$. This matrix is:

$$\mathbf{M} \equiv \frac{\partial \xi}{\partial \eta} = \frac{\partial(\mathbf{x}', \mathbf{p}')}{\partial(\mathbf{x}, \mathbf{p})}. \tag{6.17}$$

Expressing the time derivative of the new coordinates in terms of the old and using Hamilton's equations in the latter:

$$\begin{aligned}\dot{\boldsymbol{\xi}} &= \mathbf{M}\dot{\boldsymbol{\eta}} \\ &= \mathbf{M}\mathbf{J}\frac{\partial H}{\partial \boldsymbol{\eta}},\end{aligned} \tag{6.18}$$

where Hamilton's equations have been written compactly via **J**, which is the skew block matrix with **I** in the upper right and $-\mathbf{I}$ in the lower left. Now, H could also be considered a function of $\boldsymbol{\xi}$ through the coordinate transformation, whereby

$$\dot{\boldsymbol{\xi}} = \mathbf{M}\mathbf{J}\mathbf{M}^T\frac{\partial H}{\partial \boldsymbol{\xi}} \tag{6.19}$$

under some smoothness and invertibility assumptions on the coordinate transformation. Note that the transpose of the transformation appears in the coordinate change; in the language of tensors, this is a consequence of gradients transforming as a dual vector [Goldstein et al., 2002].

On the other hand, if **M** represents a canonical transformation, then Hamilton's hold in the transformed coordinates $\boldsymbol{\xi}$:

$$\dot{\boldsymbol{\xi}} = \mathbf{J}\frac{\partial H}{\partial \boldsymbol{\xi}}. \tag{6.20}$$

Identifying Eq. 6.19 and Eq. 6.20, the symplectic condition on **M** is therefore:

$$\mathbf{J} = \mathbf{M}\mathbf{J}\mathbf{M}^T. \tag{6.21}$$

Eq. 6.21 is the symplectic condition on a coordinate transformation **M**. In particular, flows along Hamilton's equations of motion are themselves canonical, and satisfy Eq. 6.21 in the infinitesimal limit [Goldstein et al., 2002]. Conversely, Eq. 6.21 provides a simple check on whether evolution in discrete time – i.e. a given numerical integrater – is canonical or not.

To this end, let use the symplectic condition to explicitly illustrate the breaking of symplectic symmetry in the Euler scheme. In the Euler scheme, the evolution of a Hamiltonian system in finite time Δt is given by:

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \Delta t \left.\frac{\partial H}{\partial \mathbf{p}}\right|_{\mathbf{x}_n,\mathbf{p}_n} \equiv \mathbf{x}_n + H_{\mathbf{p}}\Delta t \tag{6.22}$$

$$\mathbf{p}_{n+1} = \mathbf{p}_n - \Delta t \left.\frac{\partial H}{\partial \mathbf{x}}\right|_{\mathbf{x}_n,\mathbf{p}_n} \equiv \mathbf{p}_n - H_{\mathbf{x}}\Delta t \tag{6.23}$$

The transformation matrix is:

$$\mathbf{M} = \frac{\partial(\mathbf{x}_{n+1},\mathbf{p}_{n+1})}{\partial(\mathbf{x}_n,\mathbf{p}_n)} = \begin{bmatrix} \mathbf{I}+H_{\mathbf{xp}}\Delta t & H_{\mathbf{pp}}\Delta t \\ -H_{\mathbf{xx}}\Delta t & \mathbf{I}-H_{\mathbf{xp}}\Delta t \end{bmatrix}, \tag{6.24}$$

whereby

$$\mathbf{MJM}^T = \begin{bmatrix} \mathbf{I}+H_{\mathbf{xp}}\Delta t & H_{\mathbf{pp}}\Delta t \\ -H_{\mathbf{xx}}\Delta t & \mathbf{I}-H_{\mathbf{xp}}\Delta t \end{bmatrix} \begin{bmatrix} 0 & \mathbf{I} \\ -\mathbf{I} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{I}+H_{\mathbf{xp}}\Delta t & -H_{\mathbf{xx}}\Delta t \\ H_{\mathbf{pp}}\Delta t & \mathbf{I}-H_{\mathbf{xp}}\Delta t \end{bmatrix} \tag{6.25}$$

$$= \begin{bmatrix} 0 & \mathbf{I}+\mathscr{O}\left(\Delta t^2\right) \\ -\mathbf{I}+\mathscr{O}\left(\Delta t^2\right) & 0 \end{bmatrix}, \tag{6.26}$$

where the second partials $H_{\mathbf{xp}}$, $H_{\mathbf{xx}}$, and $H_{\mathbf{pp}}$ are evaluated at $(\mathbf{x}_n,\mathbf{p}_n)$ and $\mathscr{O}\left(\Delta t^2\right)$ consists only of terms proportional to Δt^2. Two things are apparent here: i) since Eq. 6.26 is not exactly equal to **J**, the Euler step is not symplectic, and ii) the scheme violates symplecticity only to second order. The vanishing of terms first order in Δt indicates that symplecticity is preserved in the continuous limit (as shown above), however this finite integrator will unavoidably accumulate errors, as we found in Figures 6.2.

We can use the same idea to illustrate that other integrators are exactly symplectic, such as leapfrog; here, we show symplecticity in the midpoint method, as we will use this particular integrator later in this chapter. The midpoint rule discretizes the time evolution via:

$$\begin{aligned} \mathbf{x}_{n+1} &= \mathbf{x}_n + \Delta t \left.\frac{\partial H}{\partial \mathbf{p}}\right|_{\mathbf{x}_{n+1/2},\mathbf{p}_{n+1/2}} \equiv \mathbf{x}_n + H_{\mathbf{p}}\Delta t \\ \mathbf{p}_{n+1} &= \mathbf{p}_n - \Delta t \left.\frac{\partial H}{\partial \mathbf{x}}\right|_{\mathbf{x}_{n+1/2},\mathbf{p}_{n+1/2}} \equiv \mathbf{p}_n - H_{\mathbf{x}}\Delta t \end{aligned} \tag{6.27}$$

where the partials of H are evaluated at the midpoints

$$\begin{aligned}\mathbf{x}_{n+1/2} &= \frac{\mathbf{x}_n + \mathbf{x}_{n+1}}{2} \\ \mathbf{p}_{n+1/2} &= \frac{\mathbf{p}_n + \mathbf{p}_{n+1}}{2}.\end{aligned} \tag{6.28}$$

Due to the symmetry of this integrator, it will be easiest to work with the combined coordinates $\mathbf{z} = \{\mathbf{x}, \mathbf{p}\}$, whereby equation Eq. 6.27 can be written

$$\mathbf{z}_{n+1} = \mathbf{z}_n + \Delta t \mathbf{J} \boldsymbol{\nabla} H\big|_{\mathbf{z}_{n+1/2}}. \tag{6.29}$$

The transformation matrix is found by taking the gradient of this vector equation with respect to $\mathbf{z}_n$, and using the fact that H depends on *both* $\mathbf{z}_n$ and $\mathbf{z}_{n+1}$ through its argument $\mathbf{z}_{n+1/2}$. That is,

$$\begin{aligned}\frac{\partial}{\partial \mathbf{z}_n}\Big\{\mathbf{z}_{n+1} &= \mathbf{z}_n + \Delta t \mathbf{J} \nabla H\big|_{\mathbf{z}_{n+1/2}}\Big\} \\ \mathbf{M} &= \mathbf{I} + \Delta t \mathbf{J}(\nabla^2 H)\frac{\partial}{\partial \mathbf{z}_n}(\mathbf{z}_{n+1/2}) \\ \mathbf{M} &= \mathbf{I} + \frac{\Delta t}{2}\mathbf{J}(\nabla^2 H)(\mathbf{I} + \mathbf{M}) \\ \mathbf{M} &= (\mathbf{I} - \frac{\Delta t}{2}\mathbf{J}\nabla^2 H)^{-1}(\mathbf{I} + \frac{\Delta t}{2}\mathbf{J}\nabla^2 H)\end{aligned} \tag{6.30}$$

To prove the symplectic condition, consider $\mathbf{J}\mathbf{M}^{-\mathrm{T}}$, where $\mathbf{M}^{-\mathrm{T}}$ is the inverse transpose:

$$\mathbf{J}\mathbf{M}^{-\mathrm{T}} = \mathbf{J}(\mathbf{I} - \frac{\Delta t}{2}\nabla^2 H \mathbf{J})^{-1}(\mathbf{I} + \frac{\Delta t}{2}\nabla^2 H \mathbf{J}), \tag{6.31}$$

using that $\mathbf{J}^{\mathrm{T}} = \mathbf{J}^{-1} = -\mathbf{J}$. The first two factors on the right hand side of this equation can be written:

$$\begin{aligned}
\mathbf{J}(\mathbf{I}-\frac{\Delta t}{2}\nabla^2 H\mathbf{J})^{-1} &= -\mathbf{J}^{-1}(\mathbf{I}-\frac{\Delta t}{2}\nabla^2 H\mathbf{J})^{-1} \\
&= -((\mathbf{I}-\frac{\Delta t}{2}\nabla^2 H\mathbf{J})\mathbf{J})^{-1} \\
&= -(\mathbf{J}+\frac{\Delta t}{2}\nabla^2 H)^{-1} \\
&= -(\mathbf{J}(\mathbf{I}+\frac{\Delta t}{2}\mathbf{J}^{-1}\nabla^2 H))^{-1} \\
&= -(\mathbf{I}+\frac{\Delta t}{2}\mathbf{J}^{-1}\nabla^2 H)^{-1}\mathbf{J}^{-1} \\
&= (\mathbf{I}-\frac{\Delta t}{2}\mathbf{J}\nabla^2 H)^{-1}\mathbf{J}. \qquad (6.32)
\end{aligned}$$

Inserting this back into Eq. 6.31 and multiplying the middle $\mathbf{J}$ through the last factor gives:

$$\begin{aligned}
\mathbf{J}\mathbf{M}^{-\mathrm{T}} &= (\mathbf{I}-\frac{\Delta t}{2}\mathbf{J}\nabla^2 H)^{-1}(\mathbf{I}+\frac{\Delta t}{2}\mathbf{J}\nabla^2 H)\mathbf{J} \\
&= \mathbf{M}\mathbf{J}, \qquad (6.33)
\end{aligned}$$

from which the symplectic condition Eq. 6.21 immediately follows. Similar procedures can be used to show that classical RK4 and the trapezoidal rule are not symplectic, while leapfrog, symplectic Euler, and other classes of Runge-Kutta methods are [Hairer et al., 2006].

6.3.2 Exploiting symplecticity in nonlinear dynamical inference

The pitfalls of integrating Hamiltonian systems with arbitrary numerical schemes first became clear through early computer simulations of celestial motion [Channell and Scovel, 1990, Yoshida, 1990, Wisdom and Holman, 1991]. The divergent trajectories of these lengthly integrations, which should have produced bounded orbits, indicated that fundamental conservation laws were somehow being violated. In light of the instability and divergence we have seen in the Lagrangian formulation of nonlinear inference, it is of interest to see if the Hamiltonian formulation, in which these symmetries can be explicitly enforced with symplectic numerical integration schemes, may solve this problem – at least in part.

Since the Hamiltonian formulation underlies much of optimal control theory and dynamic programming, it would seem curious that the idea of symplectic integration had not yet already been explored in these contexts. There has been some work in applying symplectic integrators to optimal control, but they are careful to note the critical difference versus mechanics: this is a two-point boundary value problem which, if using multiple shooting or collocation for example, does not demand long time integrations with free endpoints that symplectic integration intends to address [Chyba et al., 2009]. Not surprisingly, it is found that the benefits are fairly problem-dependent.

A further twist in the story comes in a recent exploration of symplectic invariants arising directly from variations of the discrete action, which can be developed both in Hamiltonian and Lagrangian formulations[1] [Wendlandt and Marsden, 1997, Marsden and West, 2001]. In fact, it turns out that symplectic structures are preserved quite readily in this way for a large class of classical actions. In this sense, preserving symplecticity appears more natural in the setting of variational principles derived from these discrete actions than in the Hamiltonian context, where integrators must be explicitly devised. Having already investigated variational annealing in the framework of the least action principle, let us therefore proceed by investigating the symplectic structure of discrete actions. Our hope is that the benefits of symplecticity can somehow be melded with the idea of incremental annealing in the estimation of strongly nonlinear and chaotic models, where the conditional distributions are far from Gaussian.

6.4 Symplecticity from discretized classical actions

6.4.1 Discrete Euler-Lagrange equations and symplectic flows

The conservation of phase space in Hamiltonian spaces is well-known in the physics literature; less commonly noted is that Lagrangian spaces also admit symplectic bilinear forms that are conserved along their integral curves. This bilinear form is [Marsden and West, 2001]:

$$\omega_{\mathscr{L}} = \sum_{a,b=1}^{D} \frac{\partial^2 \mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t)}{\partial x_a \partial \dot{x}_b} dx_a \wedge dx_b + \sum_{a,b=1}^{D} \frac{\partial^2 \mathscr{L}(\mathbf{x},\dot{\mathbf{x}},t)}{\partial \dot{x}_a \partial \dot{x}_b} d\dot{x}_a \wedge dx_b. \tag{6.34}$$

Perhaps the reason that this invariant is omitted in the classical physics literature is that unlike the differential area $dp_i \wedge dq_i$, it lacks an obvious physical interpretation. The correspondence to Hamiltonian mechanics is simple to note, however, by using the definition of the canonical momenta $p_a = \partial/\partial \dot{x}_a$.

It was recently found that in the Lagrangian formulation, integrators that preserve symplectic in-

---

[1] This dichotomy is somewhat reminiscent of that between direct and indirect methods in optimal control – the former of which follows a notion of *discretize then optimize*, the latter *optimize then discretize*.

variants arise almost automatically [Wendlandt and Marsden, 1997, Marsden and West, 2001, Hairer et al., 2006]. In fact, invoking Hamilton's principle directly on the discrete produces at once discrete analogues of the integrator, the bilinear form, and the Euler-Lagrange equations. Being derived from Hamilton's principle, such schemes are naturally called *variational integrators*. Various works pointing to the connection between variational principles and the conservations of symplectic invariants had appeared throughout the late 80s and 90s; the theory was formalized first in [Wendlandt and Marsden, 1997] and proceeded by many extensions since [Marsden and West, 2001, Hairer et al., 2006, Lall and West, 2006, Leok and Zhang, 2011].

The formulation begins with the discrete analogue of Hamilton's principle, called the *Discrete Variational Principle* (DVP). The DVP applies to a discrete Lagrangian L_d, which is a map $L_\mathrm{d} : Q \times Q \to \mathbb{R}$, where Q for our purposes is the state space of the dynamical system: $\mathbf{x}_n \in Q = \mathbb{R}^n$. The discrete action $S_\mathrm{d} : Q^{N+1} \to \mathbb{R}$ is then defined as:

$$S_\mathrm{d} = \int_0^T \mathscr{L}(\mathbf{x}, \dot{\mathbf{x}}, t)dt \to \sum_{n=0}^{N-1} L_\mathrm{d}(\mathbf{x}_n, \mathbf{x}_{n+1}). \tag{6.35}$$

Defining a discrete Lagrangian based on a continuous functional form can be done in a variety of ways concomitant with this definition, the only requirement being that S_d can be represented as a sum of discrete Lagrangians L_d with homogeneous dependence on "adjacent" times and state vectors $\dot{\mathrm{W}}$ith this discrete action in hand, the DVP state that the evolution of the state in discrete time extremizes S_d subject to fixed endpoints q_0 and q_N. The derivatives are straightforward, resulting in the discrete Euler-Lagrange equations:

$$0 = D_2 L_\mathrm{d}(\mathbf{x}_{n-1}, \mathbf{x}_n) + D_1 L_\mathrm{d}(\mathbf{x}_n, \mathbf{x}_{n+1}) \qquad ; \qquad n = 1, ..., N-1 \tag{6.36}$$

where the notation D_1 in the second line indicates derivative with respect to the first argument.

However, we see that the DVP under with endpoints fixed is useless from the point of view of the inference problem: q_0 and q_N cannot be held fixed because they are unknown. Employing instead the DVP without assuming fixed endpoints gives:

$$0 = \sum_{n=1}^{N-1} [D_2 L_\mathrm{d}(\mathbf{x}_{n-1}, \mathbf{x}_n) + D_1 L_\mathrm{d}(\mathbf{x}_n, \mathbf{x}_{n+1})] \cdot \delta \mathbf{x}_n + D_1 L_\mathrm{d}(\mathbf{x}_0, \mathbf{x}_1) \cdot \delta \mathbf{x}_0 + D_2 L_\mathrm{d}(\mathbf{x}_{N-1}, \mathbf{x}_N) \cdot \delta \mathbf{x}_N. \tag{6.37}$$

On integral curves of the Euler-Lagrange equations, the summed term vanishes. We can remove the boundary terms by demanding that, rather than fixed endpoints, the derivative terms vanish,

$$D_1 L_\mathrm{d}(\mathbf{x}_0, \mathbf{x}_1) = D_2 L_\mathrm{d}(\mathbf{x}_{N-1}, \mathbf{x}_N) = 0. \tag{6.38}$$

These are the "natural boundary conditions" for the discrete Lagrangian of the inference problem.

It is apparent that the Euler-Lagrange dynamics, together with appropriate boundary conditions, follow directly from the stationarity of the action when considered a function of $\mathbf{x}_0, ..., \mathbf{x}_N$. What's more, the flow of the discrete Euler-Lagrange equations preserves a discrete symplectic form. This bilinear form essentially arises as the exterior derivative of the one-form comprising the boundary terms at $n = 0$ and $n = N$. Proper treatment in terms of the tangent bundles on which the Lagrangian is defined and the associate tangent spaces of one forms is found in [Marsden and West, 2001], but briefly, these boundary one-forms sum to:

$$\Theta^+(\mathbf{x}_{N-1}, \mathbf{x}_N) + \Theta^-(\mathbf{x}_0, \mathbf{x}_1) = \frac{\partial L_\mathrm{d}(\mathbf{x}_0, \mathbf{x}_1)}{\partial \mathbf{x}_0} d\mathbf{x}_0 + \frac{\partial L_\mathrm{d}(\mathbf{x}_{N-1}, \mathbf{x}_N)}{\partial \mathbf{x}_N} d\mathbf{x}_N. \tag{6.39}$$

These boundary terms arose from differentiating the discrete action S_d along Euler-Lagrange flows. Taking a second exterior derivative gives:

$$\begin{aligned}
d^2 S_\mathrm{d} &= d\Theta^+(\mathbf{x}_{N-1}, \mathbf{x}_N) + d\Theta^-(\mathbf{x}_0, \mathbf{x}_1) \\
&= d\left[\frac{\partial L_\mathrm{d}(\mathbf{x}_{N-1}, \mathbf{x}_N)}{\partial \mathbf{x}_N} d\mathbf{x}_N\right] + d\left[\frac{\partial L_\mathrm{d}(\mathbf{x}_0, \mathbf{x}_1)}{\partial \mathbf{x}_0} d\mathbf{x}_0\right] \\
&= \frac{\partial^2 L_\mathrm{d}(\mathbf{x}_{N-1}, \mathbf{x}_N)}{\partial \mathbf{x}_{N-1} \partial \mathbf{x}_N} d\mathbf{x}_N \wedge d\mathbf{x}_{N-1} + \frac{\partial^2 L_\mathrm{d}(\mathbf{x}_0, \mathbf{x}_1)}{\partial \mathbf{x}_0 \partial \mathbf{x}_1} d\mathbf{x}_0 \wedge d\mathbf{x}_1 \\
&= -\frac{\partial^2 L_\mathrm{d}(\mathbf{x}_{N-1}, \mathbf{x}_N)}{\partial \mathbf{x}_{N-1} \partial \mathbf{x}_N} d\mathbf{x}_{N-1} \wedge d\mathbf{x}_N + \frac{\partial^2 L_\mathrm{d}(\mathbf{x}_0, \mathbf{x}_1)}{\partial \mathbf{x}_0 \partial \mathbf{x}_1} d\mathbf{x}_0 \wedge d\mathbf{x}_1
\end{aligned} \tag{6.40}$$

using $d\mathbf{x}_0 \wedge d\mathbf{x}_0 = 0$. Since $d^2 S_\mathrm{d} = 0$, then the final line states that the following quantity is conserved between t_0 and t_N (or more generally for any subinterval of $[0, N]$):

$$\omega_{\mathrm{L_d}} = \frac{\partial^2 L_\mathrm{d}(\mathbf{x}_n, \mathbf{x}_{n+1})}{\partial \mathbf{x}_n \partial \mathbf{x}_{n+1}} d\mathbf{x}_n \wedge d\mathbf{x}_{n+1}, \tag{6.41}$$

which is naturally identified with the discrete symplectic bilinear form. We say then that, in analogy to the continuous case, the flow of the discrete Euler-Lagrange equations, Eq. 6.36, is *discretely symplectic* on $Q \times Q$.

6.4.2 From discrete actions to variational integrators

Having shown that the Euler-Lagrange equations, boundary conditions, and the conservation of symplectic invariants result automatically from the discrete variational principle applied to discrete actions of the form Eq. 6.35, let us now construct these discrete actions from their continuous counterparts. This is basically a matter of choosing discretization schemes for the position and velocity variables $\mathbf{x}$ and $\dot{\mathbf{x}}$. The corresponding discrete Euler-Lagrange equations then produce a variational integrator whose particular form depends on this discretization.

In light of the fact that two of the most rudimentary numerical integration procedures – the explicit Euler step and the trapezoidal rule – are not symplectic, while others such as midpoint rule are symplectic, it appears remarkable that arbitrarily discretizing the Lagrangian iteratively as in Eq. 6.35 *always* produces a symplectic map. At lower orders, the most natural methods of discretizing the position and velocity in the windows $[t_n, t_{n+1}]$ produce familiar schemes. Let us investigate two of these now.

First consider approximating the integral $\int \mathscr{L} dt$ by interpolating $\mathbf{x}(t)$ linearly between $\mathbf{x}_n$ and $\mathbf{x}_{n+1}$, whereby

$$S_{\mathrm{d}} \approx \sum_{n=1}^{N-1} \int_{t_n}^{t_{n+1}} dt \mathscr{L}(\mathbf{x}_n(1-\alpha(t)) + \mathbf{x}_{n+1}\alpha(t), (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t), \tag{6.42}$$

with $\alpha(t) = (t - t_n)/\Delta t$ and $\Delta t = t_{n+1} - t_n$. Approximating these integrals by the trapezoidal rule then gives:

$$L_{\mathrm{d}}(\mathbf{x}_n, \mathbf{x}_{n+1}) = \frac{\Delta t}{2} \mathscr{L}(\mathbf{x}_n, (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t) + \frac{\Delta t}{2} \mathscr{L}(\mathbf{x}_{n+1}, (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t) \tag{6.43}$$

Deriving the discrete Euler-Lagrange equations from L_{d} will produce an integration scheme from $\mathbf{x}_{n-1}$ to $\mathbf{x}_{n+1}$ whose symplecticity is guaranteed. Since symplectic integrators designed by directly discretizing Hamilton's equations are defined in canonical coordinates, the correspondence will be made easier with the following definition of discrete momentum:

$$\mathbf{p}_n = -D_1 L_{\mathrm{d}}(\mathbf{x}_n, \mathbf{x}_{n+1}) \tag{6.44}$$

Combining Eq. 6.44 with the discrete Euler-Lagrange equation gives

$$\mathbf{p}_{n+1} = D_2 L_{\mathrm{d}}(\mathbf{x}_n, \mathbf{x}_{n+1}). \tag{6.45}$$

This procedure gives an iterative scheme for generating momenta and positions from $\mathbf{x}_n$, though the expressions may of course be implicit. For Eq. 6.43 we get

$$\begin{aligned}
\mathbf{p}_n &= -\frac{\Delta t}{2}\frac{\partial \mathscr{L}(\mathbf{x}_n, \mathbf{v}_n)}{\partial \mathbf{x}} + \frac{1}{2}\frac{\partial \mathscr{L}(\mathbf{x}_n, \mathbf{v}_n)}{\partial \dot{\mathbf{x}}} + \frac{1}{2}\frac{\partial \mathscr{L}(\mathbf{x}_{n+1}, \mathbf{v}_n)}{\partial \dot{\mathbf{x}}} \\
\mathbf{p}_{n+1} &= \frac{1}{2}\frac{\partial \mathscr{L}(\mathbf{x}_n, \mathbf{v}_n)}{\partial \dot{\mathbf{x}}} + \frac{\Delta t}{2}\frac{\partial \mathscr{L}(\mathbf{x}_{n+1}, \mathbf{v}_n)}{\partial \mathbf{x}} + \frac{1}{2}\frac{\partial \mathscr{L}(\mathbf{x}_{n+1}, \mathbf{v}_n)}{\partial \dot{\mathbf{x}}}
\end{aligned} \tag{6.46}$$

where $\mathbf{v}_n \equiv (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t$. For a Lagrangian with canonical kinetic energy and potential $U(\mathbf{x})$, $\mathbf{f}(\mathbf{x}) = -\nabla U(\mathbf{x})$ this gives:

$$\begin{aligned}
\mathbf{p}_n &= -\frac{\Delta t}{2}\mathbf{f}(\mathbf{x}_n) + \mathbf{v}_n \\
\mathbf{p}_{n+1} &= \mathbf{v}_n - \frac{\Delta t}{2}\mathbf{f}(\mathbf{x}_{n+1})
\end{aligned} \tag{6.47}$$

Calling $\mathbf{v}_n$ as the intermediate momentum $\mathbf{p}_{n+1/2}$, and combining its definition with the two equations above gives

$$\begin{aligned}
\mathbf{p}_{n+1/2} &= \mathbf{p}_n + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x}_n) \\
\mathbf{x}_{n+1} &= \mathbf{x}_n + \Delta t\,\mathbf{p}_{n+1/2} \\
\mathbf{p}_{n+1} &= \mathbf{p}_{n+1/2} + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x}_{n+1}),
\end{aligned} \tag{6.48}$$

which is the symplectic leapfrog method of Eq. 6.16. Similarly, it can be shown that approximating Eq. 6.42 instead by the midpoint method will generate the symplectic midpoint method.

If the basic Euler scheme is not itself symplectic, then which variational integrator does a first order approximation of the discrete action produce? Consider approximating Eq. 6.42 by a left Riemann sum:

$$
\begin{aligned}
L_\mathrm{d}(\mathbf{x}_n, \mathbf{x}_{n+1}) &= \Delta t \mathscr{L}(\mathbf{x}_n, (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t) \\
\rightarrow \mathbf{p}_n &= -\Delta t \frac{\partial \mathscr{L}}{\partial \mathbf{x}} + \frac{\partial \mathscr{L}}{\partial \dot{\mathbf{x}}} \\
\mathbf{p}_{n+1} &= \frac{\partial \mathscr{L}}{\partial \dot{\mathbf{x}}},
\end{aligned}
\tag{6.49}
$$

and for the canonical Lagrangian with unity mass:

$$
\begin{aligned}
\mathbf{p}_n &= -\Delta t \mathbf{f}(\mathbf{x}) + (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t \\
\mathbf{p}_{n+1} &= (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t,
\end{aligned}
\tag{6.50}
$$

giving:

$$
\begin{aligned}
\mathbf{x}_{n+1} &= \mathbf{x}_n + \Delta t \mathbf{p}_{n+1} \\
\mathbf{p}_{n+1} &= \mathbf{p}_n + \Delta t \mathbf{f}(\mathbf{x}_n).
\end{aligned}
\tag{6.51}
$$

Comparing with Eq. 6.23, we see that to render the Euler scheme symplectic, the vector fields $\partial_p H$ and $\partial_x H$ must be evaluated not at t_n, but opposite *both* boundaries: $\mathbf{x}_n$ and $\mathbf{p}_{n+1}$.

6.4.3 Discrete variational Hamiltonian mechanics

The discrete variational principle can also be developed in a Hamiltonian space to produce analogous variational integrators in canonical coordinates [Wendlandt and Marsden, 1997, Lall and West, 2006]. These essentially arise as discrete Legendre transforms of the discrete Langrangian L_d, thought there are two choices corresponding to either the *right* or *left* Legendre transform. Either of these integrators could be derived for the Gaussian error action. Though this formulation is mentioned for completeness, it will not be pursued here; some details of the formulation can be found in [Kadakia et al., 2017].

6.5 A Hamiltonian algorithm for nonlinear estimation

In the past several sections, we have encountered the existence of symplectic symmetry in Hamiltonian systems, its violation in discrete time, and its subsequent reinstatement either from carefully constructed mappings of Hamilton's equations, or directly via a discrete variational principle.

Let us now return to the discussion of nonlinear inference as a mechanical system. The intent of the preceding sections was to suggest that geometric structures inherent in Hamiltonian flows may be leveraged to address the problem of persistent instability. This is unlikely for two reasons. First, instability is a continuous time issue that is a *consequence* of symplectic symmetry, not a byproduct of careless numerics. Second, the development of variational integrators from the discrete variational principle means that symplectic invariants are preserved in the Lagrangian formulation anyway, simply by satisfying the stationarity of $A(\mathbf{X}|\mathbf{Y})$ [2]. As mentioned, the variational principle produces not only the E-L equations, but the boundary conditions as well, implying that minimizing $A(\mathbf{X}|\mathbf{Y})$ at once both preserves symplecticity and enforces the boundary conditions. It is therefore unclear what more can be gained by moving to a Hamiltonian context.

On the other hand, the exposition of the inference problem is somewhat more natural in canonical coordinates than in the discrete action space Q. For example, applying the terminal boundary condition of the discrete action principle to the Gaussian error action Eq. 2.44 with $g_d(\cdot)$ discretized by the trapezoidal rule gives:

$$\begin{aligned}\frac{\partial}{\partial \mathbf{x}_N} L_{\mathrm{d}}(\mathbf{x}_{N-1},\mathbf{x}_N) = R_{\mathrm{f}}\Bigg[\mathbf{x}_N - \frac{\Delta t}{2}\left(\mathbf{f}(\mathbf{x}_{N-1}) + \mathbf{f}(\mathbf{x}_N) + x_a(t_N)\frac{\partial f_a(\mathbf{x}_N)}{\partial \mathbf{x}_N}\right) \\ + \frac{\Delta t^2}{2}\left(f_a(\mathbf{x}_{N-1}) + f_a(\mathbf{x}_N)\right)\frac{\partial f_a(\mathbf{x}_N)}{\partial \mathbf{x}_N}\Bigg]. \qquad (6.52)\end{aligned}$$

This boundary condition and its analogous expression at t_0 are i) dependent on R_{f} – which changes iteratively over several orders throughout the variational annealing procedure, ii) sensitive to the discretization format of the discrete action, thereby inheriting its error magnitudes and iii) not enforced explicitly, rather satisfied as a byproduct of the optimization itself. Conversely, boundary conditions in the Hamiltonian formulation of inference are trivial ($\mathbf{p}_N = \mathbf{p}_0 = 0$) and can be enforced more readily as constraints, so perhaps work-

---

[2] This is true for higher-order integrators as well. Though only first and second-order variational integrators were presented above, the discrete action formulation Eq. 6.35 can be extended straightforwardly to approximations with multiple stages in each time window $[t_n, t_{n+1}]$, allowing naturally for quadrature schemes that generate variational integrators of higher-order.

ing in canonical coordinates could remove some 'unnaturalness' of the description, better conditioning the optimization process.

Certainly, we might carry out the nonlinear estimation with Hamilton's equations, Eqs. 6.10, employing multiple shooting or collocation. But we would likely expect difficulties of the sort presaged by the $2D$ chaotic oscillator above: multiple solutions satisfying the boundary conditions Eqs. 6.11 may be uncovered by the nonlinear optimization.

After all, recall that the equations of motion arose initially as the saddle point approximation of a high-dimensional path integral, which technically sums an infinite number of contributing paths. Laplace's method solves this impossible calculation by retaining only those that contribute more than locally neighboring paths. If the action admits more than one stationary path, the individual weights must be compared, yet the equations of motion contain no prescription for this comparison – the action values themselves must be invoked. In this sense, the well-posed functional path integral could actually produce an ill-posed boundary value problem through the method of stationary phase. In the context of nonlinear estimation of chaotic systems, this is not just a possibility. It is ubiquitous.

In Ch. 3 we found that the issue of multiple minima could be often mitigated by incorporating nonlinearities in a systematic and gradual way. It would be desirable if a similar annealing protocol could be developed in canonical coordinates, the hope that the boundary value problem might be solved more reliably in this space, particularly if the annealing is initialized with a tractable, convex cost function.

To this end, consider the continuous Lagrangian of nonlinear inference, expressed in canonical coordinates:

$$\mathscr{L}_{\mathrm{I,H}}(\mathbf{x},\mathbf{p},t) = \mathbf{p}^2\big|_{\mathbf{Q}} + (\mathbf{y}-\mathbf{H}\mathbf{x}_n)^2\big|_{\mathbf{R}^{-1}} + \frac{1}{2}\boldsymbol{\nabla}\cdot\mathbf{f}(\mathbf{x}), \tag{6.53}$$

for which the "action" is, in discrete time [3]:

$$A_{\mathrm{H}}(\mathbf{Z}|\mathbf{Y}) = \sum_{n=0}^{N}\mathbf{p}_n^2\big|_{\mathbf{Q}_n} + \sum_{n=1}^{N}(\mathbf{y}_n-\mathbf{H}\mathbf{x}_n)^2\big|_{\mathbf{R}_n^{-1}} \tag{6.54}$$

where $\mathbf{Z} = \{\mathbf{x}_n, \mathbf{p}_n\}$. *Action* is quoted here since this expression is, in and of itself, meaningful in neither the Lagrangian nor Hamiltonian description: it represents essentially a time-integrated Lagrangian, but is ex-

---

[3] the divergence term arises only in the continuum limit of the Wiener integral

pressed in canonical coordinates, which are meaningful only in the Hamiltonian picture. However, along the integral curves of Hamilton's equations, it is equivalent to the original action. It therefore provides a quantification of the relative contribution to the expected value path integral, for distinct solutions of Hamilton's equations.

Of equal importance is the structure of Eq. 6.54: much like the discrete action in configuration space in the $R_{\mathrm{f}} \to 0$ limit, it is convex, though now for all R_{f}. The model nonlinearities, on the other hand, reside entirely in Hamilton's equations. We could exploit this decoupling much like in variational annealing if, somehow, the equations of motion themselves could themselves be "annealed."

In a nonlinear programming context, for example, we might optimize $A_{\mathrm{H}}(\mathbf{Z}|\mathbf{Y})$ with Hamilton's equations incorporated as *inequality* constraints. The iterative aspect of the algorithm would introduce the effect of Hamilton dynamics by gradually decreasing the constraint slack on the Hamilton's equations of motion, slowly deforming the feasible region and transforming the constraints from inactive to active.

Specifically, let us consider the following *Hamiltonian variational annealing* algorithm, in which R_{f} is fixed, and we perform the following optimization:

$$\begin{aligned} &\text{minimize} \quad A_{\mathrm{H}}(\mathbf{Z}) \\ &\text{subject to} \\ &\quad |\phi(\mathbf{z}_n, \mathbf{z}_{n+1}, \ldots, \mathbf{z}_{n+m-1})| < g_{\mathrm{H}} = \frac{g_{\mathrm{H}_0}}{\alpha^{\beta}}, \end{aligned} \tag{6.55}$$

where $\phi(\cdot)$ is an mth-order discretization of Hamilton's equations of motion. The constraint tolerance in the annealing protocol starts at a relatively large value g_{H_0} and is decreased by a factor $1/\alpha$ at each step $\beta = 1, 2, \ldots$ of the anneal, the idea being that as g_{H} gets smaller, the dynamical equations are imposed more strictly through Hamilton's equations (see Algorithm 2).

Optimizing Eq. 6.54 in this iteratively constrained manner allows us to work entirely in canonical coordinates and provides much the same functionality as the Lagrangian approach – by slowly introducing the dynamical nonlinearities of the model – in a more direct way. A distinct difference is that in canonical coordinates, natural boundary conditions may be enforced directly as a constraint on the initial and final momenta, $\mathbf{p}(t_0) = \mathbf{p}(t_N) = 0$, since they are now independent variables in the optimization itself.

Algorithm 2: Hamiltonian Variational Annealing (HA)

```
Input  : Observations $\mathbf{Y}$ and annealing parameters $\beta_{\max}$, $\alpha$, $g_{\mathrm{H0}}$
Output: Estimated trajectories, parameters, and momenta $\{\hat{\mathbf{Z}}^{q,\beta}\}$
for $q = 1,...,Q$, in parallel do
    Sample $\mathbf{Z}^q_{\mathrm{init}}$ uniformly from presumed dynamical range
end
for $\beta \leftarrow 0$ to $\beta_{\max}$ do
    $g_{\mathrm{H}} \leftarrow g_{\mathrm{H0}}/\alpha^{\beta}$
    for $q = 1,...,Q$, in parallel do
        $\hat{\mathbf{Z}}^{q,\beta} \leftarrow$ argmin  $[A_{\mathrm{H}}(\mathbf{Z}|\mathbf{Y})$ subject to Eq. 6.55], with initial guess $\mathbf{Z}^q_{\mathrm{init}}$
        $\mathbf{Z}^q_{\mathrm{init}} \leftarrow \hat{\mathbf{Z}}^{q,\beta}$
    end
end
```

6.5.1 Symplecticity, boundary conditions, and degeneracies

A few words on the differences between variational and Hamiltonian annealing are in order. Though both methods are similarly motivated, seeking stationary paths of the action through the gradual application of model dynamics, they achieve this goal in fundamentally different ways. In the Lagrangian description, iterative changes in the model precision R_{f} deform the action manifold from one that is quadratic in the measured components and flat in the unmeasured components, to one rendered highly non-convex by the nonlinearity of the vector field. The incremental way in which this is done tracks the local and global minima of the action systematically, even if for large R_{f} they occupy tiny, deep corners in the action manifold. At small R_{f}, due to the approximate flatness of the action in the unmeasured directions, the stationary paths are highly degenerate, and this degeneracy is slowly lifted with increasing model precision. Meanwhile, since the minimizing paths satisfy the Euler-Lagrange equations, by the arguments given in Section 6.4, the symplecticity of every minimum action path is guaranteed.

On the other hand, in the Hamiltonian formulation, the path is embedded in a higher-dimensional phase space of canonical coordinates. Part of the usefulness of this description is that the calculation can be carried out in a way that preserves symplectic structure. To enforce the symplectic map, however, Hamilton's equations i) would need to be discretized with a symplectic scheme, and ii) enforced as *equality* constraints. But since the solutions to Hamilton's equations coincide with the action minima in the first place, requirement (ii) would render the action redundant entirely. So it is clear that by relaxing Hamilton's equations, we are also relaxing the requirement that the stationary paths contain symplectic invariants – despite these paths residing in a Hamiltonian manifold.

Instead, by using the constraints as an annealing device to gradually impose the model dynamics, two things occur. First, in analogy to variational annealing, the degeneracies related to the partial observability of the system are slowly removed. In addition, a separate degeneracy in the momenta variables is also lifted: at low β, all $\mathbf{p}(n) \approx 0$. These degeneracies are both broken by gradually enforcing the model dynamics through the application of Hamilton's equations as constraints. Further, in contrast to variational annealing, they exist for arbitrary R_{f}, which has been demoted from its status as an algorithmic parameter altogether.

Second, symplecticity is imposed *gradually* through the enforcement of Hamilton's equations; this is in contrast to variational annealing, where the stationary paths are symplectic at *each* β. At the cost of a higher-dimensional space, we have gained more control over how the underlying symmetries of this inference dynamics are enforced.

It is of course still not certain that these aspects will actually increase the accuracy or reliability of the estimation. After all, we have magnified the complexity of the problem – not only has the search dimension doubled, but we have traded an unconstrained optimization procedure for a more computationally costly constrained one. Yet there may be some benefit to carrying out a search in the joint space of $\mathbf{x}$ and $\mathbf{p}$. We found in Section 5.2.2 that embedding the Van der Pol oscillator into a higher-dimensional canonical space inverted its attractive basins into saddles. Yet these same augmented momenta directions that compromised stability in forward integration might be exploited in the optimization context. These directions may furnish alternative paths out of what are otherwise stable local minima in configuration coordinates $x_a(t_n)$, increasing the likelihood that the global minimum of the action is found.

6.6 Numerical experiments of a chaotic attractor

Having motivated the Hamiltonian formulation of nonlinear inference, let us now compare it to variational annealing in estimating a chaotic dynamical system with partial observations. The implementation of each of these formulations will be as follows.

6.6.1 Implementation details

Variational annealing

Variational annealing will be carried out as described in Section 3.1, with $R_{\mathrm{f0}}(n,d) = R_{\mathrm{f0}}$ = 1e-4, $\alpha = 10^{1/4}$, and $\beta_{\mathrm{max}} = 40$, whereby the final R_{f} value is 10^6 (Table 6.1). Since the second-order symplectic integrators are much easier to implement than higher order ones, we will work to this order of approximation in both formulations. The dynamical model $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},t)$ is therefore discretized with a second-order Lagrange polynomial, giving

$$\begin{aligned} f(\mathbf{x}(t)) &\approx \sum_{j=1}^{2} f(\mathbf{x}(t_j)) \prod_{k\neq j}^{2} \frac{t-t_k}{t_j-t_k} \\ &= \frac{(t-t_{n+1})}{-\Delta t} f(\mathbf{x}_n) + \frac{(t-t_n)}{\Delta t} f(\mathbf{x}_{n+1}) \end{aligned} \tag{6.56}$$

whose integral yields the trapezoidal rule:

$$\int_{t_n}^{t_{n+1}} f_d(\mathbf{x}(t))dt \approx \frac{\Delta t}{2}[f_d(\mathbf{x}_n) + f_d(\mathbf{x}_{n+1})] \tag{6.57}$$

Thus the following discrete action will be minimized at each stage of the anneal:

$$\begin{aligned} A_{\mathrm{L}}(\mathbf{X} = \mathbf{x}_n|\mathbf{Y}) &= \sum_{n=0}^{N-1} L_{\mathrm{d}}(\mathbf{x}_n, \mathbf{x}_{n+1}) \\ L_{\mathrm{d}}(\mathbf{x}_n, \mathbf{x}_{n+1}) &= \sum_{d=1}^{L} \frac{R_{\mathrm{m}}(n,d)}{2} (y_d(t_n) - x_d(t_n))^2 \\ &\sum_{d=1}^{D} \frac{R_{\mathrm{f}}}{2} \left(x_d(t_{n+1}) - x_d(t_n) - \frac{\Delta t}{2}[f_d(\mathbf{x}_n) + f_d(\mathbf{x}_{n+1})] \right)^2 . \end{aligned} \tag{6.58}$$

Note that there is a subtlety in this particular discretization – it is *not* the variational integrator that corresponds to leapfrog integration derived in Section 6.4.2, despite ostensibly utilizing a trapezoidal rule. Here, we are instead approximating the “integral within” the action integral by the trapezoidal rule; in the previous case, we would approximate the full action integral directly from Eq. 5.31, giving:

$$\begin{aligned}
L_\mathrm{d}(\mathbf{x}_n, \mathbf{x}_{n+1}) &\sim \frac{1}{2}\mathscr{L}_1(\mathbf{x}_n, (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t) + \frac{1}{2}\mathscr{L}_1(\mathbf{x}_{n+1}, (\mathbf{x}_{n+1} - \mathbf{x}_n)/\Delta t) \\
&= \sum_{d=1}^{L} \frac{R_\mathrm{m}(n,d)}{2}(y_d(t_n) - x_d(t_n))^2 + \sum_{d=1}^{L} \frac{R_\mathrm{m}(n+1,d)}{2}(y_d(t_{n+1}) - x_d(t_{n+1}))^2 \\
&\quad + \sum_{d=1}^{D} \frac{R_\mathrm{f}}{2}\left(\frac{x_d(t_{n+1}) - x_d(t_n)}{\Delta t} - f_d(\mathbf{x}_n)\right)^2 + \sum_{d=1}^{D} \frac{R_\mathrm{f}}{2}\left(\frac{x_d(t_{n+1}) - x_d(t_n)}{\Delta t} - f_d(\mathbf{x}_{n+1})\right)^2 .
\end{aligned} \tag{6.59}$$

The expressions in Eqs. 6.58 and 6.59 differ by i) a couple of boundary terms involving the data, ii) a renormalization of R_f, and iii) cross terms coupling $\mathbf{f}(\mathbf{x}_n)$ and $\mathbf{f}(\mathbf{x}_{n+1})$. Items (i) and (ii) are minor issues not expected to affect the performance of the optimization (the stationary paths do not change at large R_f, so even a factor of $\sim 10^4$ given by $1/\Delta t^2$ would not affect the limiting result). Despite (iii), the beauty of variational integrators is that *any* discretization of the continuous action which can be written symmetrically as in Eq. 6.35 furnishes a symplectic map. Thus, despite the somewhat distinct way in which we implement the trapezoidal rule in this case, the stationary paths of Eq. 6.58 are indeed discretely symplectic.

At each β, the optimization will be performed again with the constrained interior point method IPOPT [Wächter and Biegler, 2006]. As in the example calculation in Section 3.2.2, there are inequality constraints on the state variables to ensure that they stay within the dynamical range of the model, $|x_d(n)| < 15$, but again this condition is not met in practice anyhow, so the constraints are effectively inactive. Analytical derivatives ∇A and $\nabla^2 A$, computed symbolically using the SymPy package in Python, are also provided.

Table 6.1: Algorithmic parameters for comparison of variational annealing with symplectic and non-symplectic Hamiltonian variational annealing.

Variational method	Annealing parameter	α	β	Initial Value	Other
Variational annealing	$R_\mathrm{f} = R_\mathrm{f0}\alpha^\beta$	$10^{1/4}$	0,1,...40	$R_\mathrm{f0} = 10^{-4}$	–
Symplectic Hamiltonian annealing	$g_\mathrm{H} = g_\mathrm{H0}/\alpha^\beta$	$10^{1/6}$	0,1,...36	$g_\mathrm{H0} = 10^2$	$R_\mathrm{f} = 10^6$, $g_p = 10^{-4}$
Non-symplectic Hamiltonian annealing	$g_\mathrm{H} = g_\mathrm{H0}/\alpha^\beta$	$10^{1/6}$	0,1,...36	$g_\mathrm{H0} = 10^2$	$R_\mathrm{f} = 10^6$, $g_p = 10^{-4}$

Symplectic Hamiltonian annealing

We will compare two variants of Hamiltonian annealing, one symplectic and one non-symplectic. In the symplectic variant, the following cost function will be minimized:

$$A_{\mathrm{H}}(\mathbf{z}_n) = \sum_{n=0}^{N}\sum_{d=1}^{D} \frac{p_d(t_n)^2}{2R_{\mathrm{f}}} + \sum_{n=0}^{N}\sum_{d=1}^{L} \frac{R_{\mathrm{m}}(n,d)}{2}(y_d(t_n) - x_d(t_n))^2, \tag{6.60}$$

where $\mathbf{z}_n = \{\mathbf{x}_n, \mathbf{p}_n\}$, subject to the following constraints,

$$\left| z_d(t_{n+1}) - z_d(t_n) - \Delta t J_{cd} \frac{\partial H_{\mathrm{I}}}{\partial z_d}(z_d^*, t_n^*) \right| < g_{\mathrm{H}} \tag{6.61}$$

$$\begin{aligned} &|p_d(t_0)| < g_p \\ &|p_d(t_N)| < g_p, \end{aligned} \tag{6.62}$$

with continuous Hamiltonian

$$H_{\mathrm{I}}(\mathbf{z} = \{\mathbf{x},\mathbf{p}\}, t) = \frac{\mathbf{p}^2}{2R_{\mathrm{f}}} - \frac{R_{\mathrm{m}}(t,d)}{2}(\mathbf{x} - \mathbf{y}(t))^2 + \mathbf{p}\cdot\mathbf{f}(\mathbf{x}) - \frac{1}{2}\nabla\cdot\mathbf{f}(\mathbf{x}), \tag{6.9 revisited}$$

whose partials in Eq. 6.61 are evaluated at the midpoints

$$\begin{aligned} t_n^* &= t_n + \Delta t/2 \\ z_d^* &= \frac{z_d(t_n) + z_d(t_{n+1})}{2}. \end{aligned} \tag{6.63}$$

Note that here, J_{cd} is the symplectic matrix $\mathbf{J}$ and the partials of H_{I} are given in Eqs. 6.10. The optimization will be carried out at each annealing step, the result of the prior being used for the initialization point in the subsequent step, as in variational annealing. The difference here is that R_{f} is held fixed at $R_{\mathrm{f}} = 10^6$, while g_{H} is multiplicatively decreased, from an initial value of $g_{\mathrm{H}0} = 10^2$ to a final value of 10^{-4} in steps of $\alpha = 10^{1/6}$ (Table 6.1).

Non-symplectic Hamiltonian annealing

This variant of the Hamiltonian algorithm is identical to the symplectic one, except the discretization of Hamilton's equations is chosen as the trapezoidal rule, which we have seen to be non-symplectic:

$$\left| z_d(t_{n+1}) - z_d(t_n) - \frac{\Delta t}{2} J_{cd} \left(\frac{\partial H_{\mathrm{I}}}{\partial z_d}(z_d(t_n), t_n) + \frac{\partial H_{\mathrm{I}}}{\partial z_d}(z_d(t_{n+1}), t_{n+1}) \right) \right| < g_{\mathrm{H}} \tag{6.64}$$

The annealing procedure is identical to the symplectic case, using the algorithmic parameters listed in Table 6.1.

Details of model and data generation

We will compare these three methods with a systematic study of the chaotic 10-dimensional Lorenz96 model:

$$f_d(\mathbf{x}) = -x_{d-2}x_{d-1} + x_{d-1}x_{d+1} - x_d + f \tag{2.29 revisited}$$

with forcing parameter $f = 8$. To create synthetic data, first $P = 100$ trajectories are generated by choosing random initial conditions from the dynamical range of the attractor and integrating each forward using an explicit fourth order Runge-Kutta scheme ($\Delta t = 0.01$). Initial transients of length $T_{\text{transient}} = 5$ are removed from each of these trajectories to ensure that the estimation is occurring on the attractor alone. The resulting 100 "true" trajectories $\{\mathbf{X}_p^*\}$ will then be used as a benchmark to compare the accuracy of the estimates. By choosing several paths in this way, we investigate the effect of differing degrees of local instability, thereby giving a comprehensive assessment of the full attractive manifold.

"Observed" datasets $\{\mathbf{Y}_p\}$ are then generated by corrupting each $\mathbf{X}_p^*$ with additive Gaussian noise of mean zero and unity variance. To consider the effect of nonlinearities, we consider separate cases with data sampled at either $\Delta t = 0.01$ and $\Delta t = 0.05$, through a total data collection window $[0, T]$, $T = 4$. To investigate the effect of sparse observations, we will consider the case when only $L = 1,\ 2,\ 3,\ 4,$ or 5 components of the full 10-dimensional space are observed.

The optimization at the first annealing stage ($\beta = 0$) must also be initialized. Distinct initializations are chosen by sampling $\mathbf{x}_d(t_n)$ uniformly across the dynamical range of the attractor, roughly $[-20, 20]$. For the Hamiltonian methods, $\mathbf{p}_d(t_n)$ is also sampled uniformly from a wide range, $\pm 10^3$. For each $\mathbf{Y}_p$, the optimization is performed in parallel for these $Q = 100$ initializations (the same initializations are used for each $\mathbf{Y}_p$), producing a set of estimated trajectories $\mathbf{X}^{\beta}_{q,p}$, and the accuracy of the annealing method is assessed by investigating the resulting estimates at the end of the annealing procedure, $\mathbf{X}^{\beta_{\max}}_{q,p}$. We denote for each p the path among for which A_{L} or A_{H} is maximal as the "optimal estimate" $\hat{\mathbf{x}}_p(t_n)$. This path gives the largest contribution to the expected value integral Eq. 2.5.

In sum, for each unique data trajectory $\mathbf{Y}_p$ we produce a progression of $\beta_{\max}$ path estimates $\mathbf{X}^{\beta}_{q,p}$ for each of the Q initial guesses of the annealing protocol. The Q initializations quantify the reliability of the algorithm, while the P distinct datasets give statistical information across the entire attractor. The former allows us address questions like: given no prior information about the system state, how often is an algorithm able to find the lowest action level? The latter provides confidence that our conclusions about a particular method are not sensitive to local properties such as model instability or system observability.

6.6.2 Results

Metrics

The accuracy and reliability of the inference methods will be quantified in two ways. The first is the pointwise accuracy of the inferred trajectories, during both the estimation window and prediction (forward integration). To generate these error estimates, only the lowest action paths $\hat{\mathbf{x}}_p(t_n)$ will be used, their mean-square-error then compared to the true trajectory:

$$E(t_n) = \frac{1}{DP}\sum_{p=1}^{P}\left(\mathbf{x}^{*}_{p}(t_n) - \hat{\mathbf{x}}_p(t_n)\right)^2 \tag{6.65}$$

This point-wise statistic, which is averaged over all dimensions and datasets, provides a global estimate for how the algorithm performs within the observation window, at its boundaries, and during prediction.

The second metric will assess the reliability of the algorithm, counting, for each data set, the number of initializing guesses that achieve an action value within a range of three standard deviations of the expected minimum of A (Eq. 3.10):

$$A \sim 1 \pm \frac{3}{\sqrt{(N+1)L/2}}. \tag{6.66}$$

(Note that this metric normalizes Eq. 3.10 to unity). This metric quantifies the sensitivity of the algorithms to distinct data sets, i.e. to particularly unstable sections of the attractor.

Lowest action levels and prediction errors

The lowest action levels for each of the P data sets are shown first for the case of $\Delta t = 0.05$ in the far left panel of Figure 6.4a. Each algorithm produces a plethora of different levels, suggesting that corresponding forward predictions are probably unreliable. This is indeed corroborated by the error metric in Figure 6.4b: the MSE at $t = T$ is ~ 10, much larger than the observational noise $\sigma^2 = 1$. Thus, while it appears the Lagrangian approach has the advantage $L = 1$, the estimation effectively fails for all three cases.

As expected, the reliability and errors improve with more measurements (middle and right panels of Figures 6.4a and 6.5b). Though the Lagrangian method gives slightly better estimates, the differences are small. For $L = 5$, only a few datasets are unable to produce the global minimum, and all annealing methods give similar quality estimates at $t = T$, thereby producing no discernible disparities in prediction errors.

Figure 6.5 repeats these same calculations with a smaller time step $\Delta t = 0.01$. While the low observability case (far left panels) presents no conclusions beyond that of the results for $\Delta t = 0.05$, the situation changes for more measurements, where a distinct advantage of the Hamiltonian method emerges. For sufficient observability, the Lagrangian method is unable to find the lowest action level for several paths along the chaotic attractor manifold. The Hamiltonian methods on the other hand consistently return lowest minima near unity for all but a few data sets. This is most pronounced for $L = 5$, where the Hamiltonian methods are essentially equivalent, with mean squared errors smaller by nearly two orders of magnitude over variational annealing.

Probability of finding global minimum

In the previous sections, the estimation and prediction errors were determined by using only paths corresponding to the lowest action level among the Q different initializations. We also compute the average proportion of the runs that produce these lowest values, providing a further measure of the reliability of our methods when given no prior information, Eq. 6.66.

These percentages averaged over the 100 datasets, are shown in Table 6.2 and Table 6.3 (for $L = 1$ these ratios are small for all three anneal types, on the order of 1 percent). By this metric, the Lagrangian method consistently outperforms both Hamiltonian methods. Between the two Hamiltonian methods, the

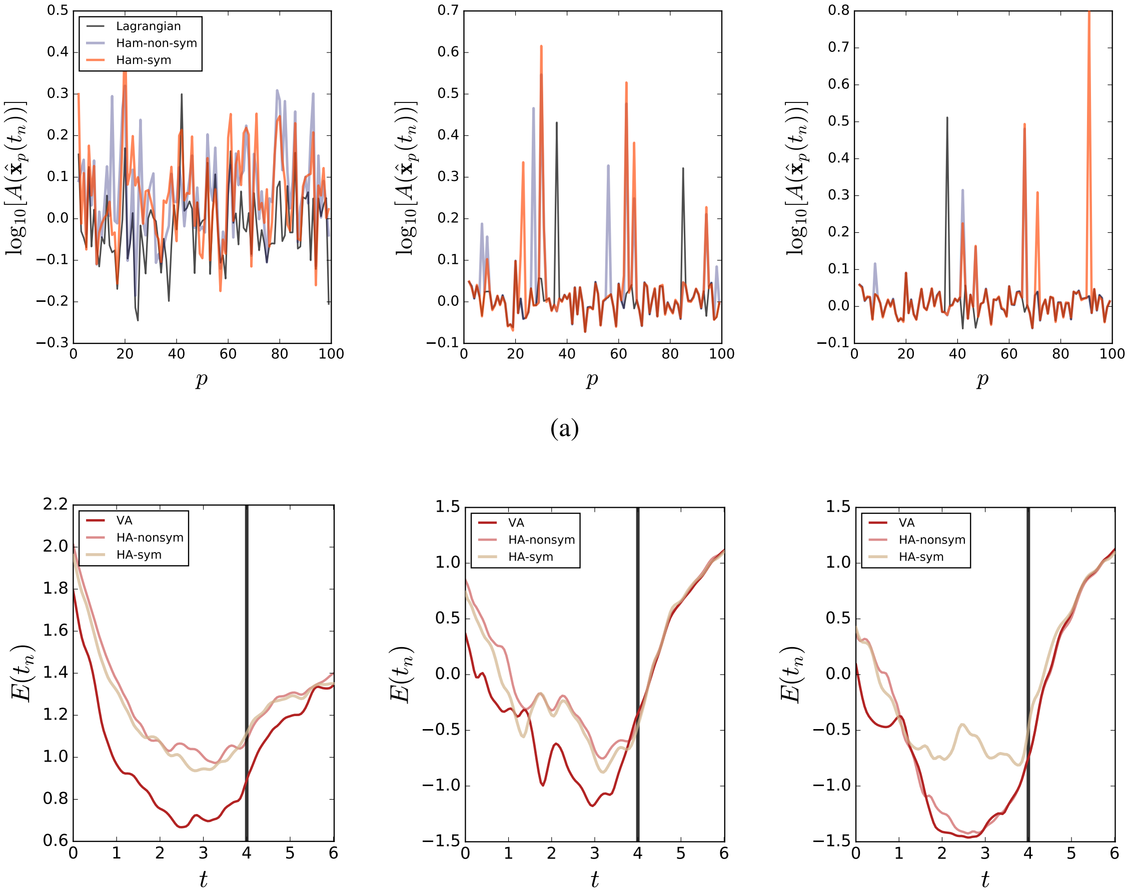

(a)

(b)

Figure 6.4: (a) Lowest action levels among the Q initial guesses, for each dataset p. The three estimation methods are distinguished by color. Here, the observation time step is $\Delta t = 0.05$, and the number of observed variables is $L = 1, 4$, and 5, for the three plots, respectively. While all methods perform better with more observations, there are a few datasets for which the global minima precludes detection. (b) Corresponding error magnitudes as a function of time, calculated by Eq. 6.65. The errors are close to the observation noise when only a single variable is measured (left panel), but reduce below the observation noise for $L = 4$ and $L = 5$. However, all three methods produce largely equivalent forward estimations when averaged over all datasets.

Table 6.2: Likelihood of finding the global minimum, averaged over all p, for the three annealing methods, for $L = 4$ measured variables.

Annealing method	$\Delta t = 0.01$	$\Delta t = 0.05$
Variational annealing	76	67
Non-symplectic Hamiltonian annealing	34	34
Symplectic Hamiltonian annealing	43	42

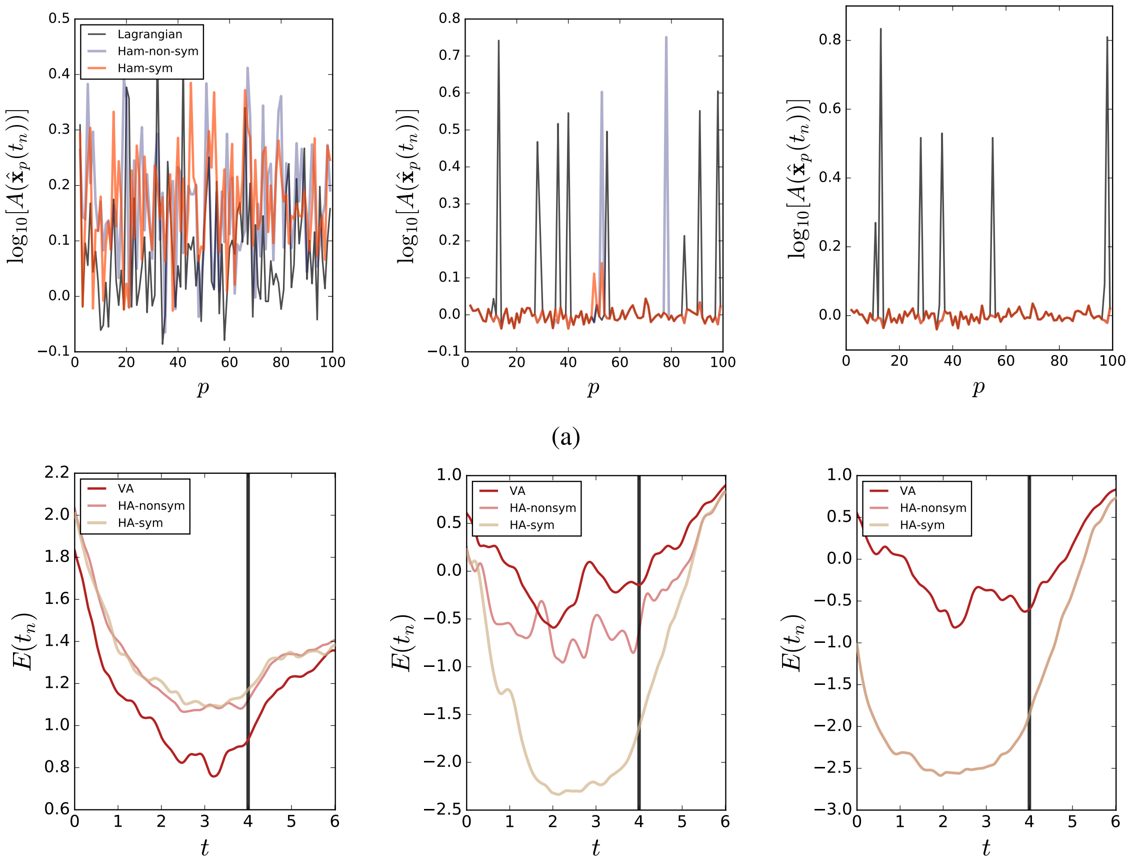

(a)

(b)

Figure 6.5: Same plots as in Figure 6.4, but now with a smaller observation timestep $\Delta t = 0.01$. While the averaged errors are again large for a single measured variable, the symplectic Hamiltonian variational annealing method substantially outperforms the other methods for $L = 4$, and both Hamiltonian methods outperform variational annealing by two orders of magnitude for $L = 5$. The large errors in the variational annealing methods arise from the overt failure to detect the global minimum for a number of datasets.

Table 6.3: Likelihood of finding the global minimum, averaged over all p, for the three annealing methods, for $L = 5$ measured variables.

Annealing method	$\Delta t = 0.01$	$\Delta t = 0.05$
Variational annealing	86	83
Non-symplectic Hamiltonian annealing	54	51
Symplectic Hamiltonian annealing	62	57

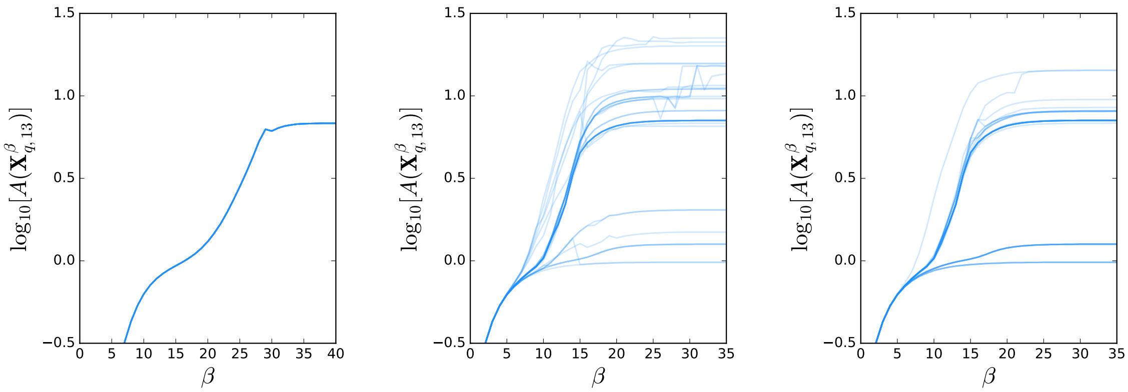

Figure 6.6: Annealing plots for variational annealing (left panel) and non-symplectic and symplectic Hamiltonian variational annealing (middle panel and right panels, respectively), for one of the outlier data sets in Figure 6.5a. While the Hamiltonian annealing methods consistently uncover more local minima than variational annealing, they nevertheless find the global minimum reliably. Variational annealing consistently converges to the same local basin, far above the global minimum, for all Q initializations.

symplectic is slightly favored.

The fact that variational annealing in $\mathbf{X}$ appears more reliable than Hamiltonian annealing in $\mathbf{X},\mathbf{P}$ by this metric, but less so by prediction errors, can be illuminated by the annealing plots corresponding to the extreme cases in the middle and right panels of Figures 6.5a. In Figure 6.6, we show the annealing plot for the 13^{th} dataset from the rightmost panel of Figure 6.5a ($\Delta t = 0.01$, $L = 5$), corresponding to one of the outliers in the Lagrangian calculation; for *all* of the $Q = 100$ initial paths, variational annealing only identifies a single level, significantly above the expected lowest normalized value near unity. Meanwhile, the Hamiltonian methods find this this and other local minima, in addition to the expected lowest value, suggesting that searches in $\{\mathbf{X},\mathbf{P}\}$, though more computationally demanding, may be beneficial nonetheless.

We may thus posit the following. Optimization in canonical coordinates opens up directions that were essentially “integrated out” by working in $\mathbf{X}$ alone. Though this produces a richer cost manifold A with more local minima, when couched in an annealing algorithm, the augmented momenta directions allow a connecting path between the minima of A that is essentially suppressed in the lower-dimensional projection down to configuration space.

As an aside, the idea of Hamiltonian annealing is similar in spirit to the method of Hamiltonian Monte Carlo (HMC) [Neal, 1996]. Distribution samples are gathered in HMC by allowing motion in a combined space consisting of both the original support and a fictitious momentum space. By alternating i) samples from the momentum variables and ii) motion along integral curves of Hamilton’s equations in this space, broader swaths of the the distribution are opened up more quickly, rapidly reducing sample correlation compared to the traditional Markov Chain method.

6.6.3 Discussion

The fact that variational annealing and Hamiltonian annealing return different solutions to the optimization is not surprising, given the distinct ways in which they lift the degeneracy associated with partial system observability. In the former, the gradual deformation from a near-quadratic to non-convex action surface can often track the global minimum, but may alternatively confine the search to a particular region of state space around a local minimum with a strong basin of attraction. This deformation process may play a role in collecting the admissible solutions to a small neighborhood of the high-dimensional action manifold early in the annealing process, *before* the finer features of the action materialize at higher R_f.

On the other hand, Hamiltonian annealing minimizes a quadratic function at every value of β, and the solution degeneracies are broken by shrinking the feasible region. The augmented momenta dimensions decouple the minimization procedure from the application of the boundary conditions and the constraints of Hamilton's equations. This facilitates a more thorough exploration of the phase space, albeit at the cost of a lower likelihood of success.

This fuller exploration of the solution space arises from the imposition of symplectic structure. This relaxation is not easily achieved in variational annealing. Although one could perform a similar approach, embedding the action in $\{\mathbf{x}, \dot{\mathbf{x}}\}$ space and applying the Euler-Lagrange equations as constraints, symplectic structure and the boundary conditions are more difficult to enforce in this representation. Furthermore, while $\mathbf{x}$ and $\dot{\mathbf{x}}$ are fundamentally independent variables in the continuous Lagrangian, they are not so in the Euler-Lagrange equations, which are equations only of $\mathbf{x}$ and its time derivatives. In this sense, it seems rather artificial to treat $\dot{\mathbf{x}}$ as independent variables.

While the Hamiltonian annealing algorithm exhibits some limited benefits in nonlinear inference, the nature and role of symplectic structure in nonlinear inference and data assimilation will obviously demand further investigation. Several questions remain unanswered by this work. For instance, why does the benefit of the Hamiltonian formulation disappear at larger Δt? It could be related to the fact that the Lagrangian formulation has the structure of a variational integrator, and it would be interesting to see an analogous variational integrator derived in the Hamiltonian description that also anneals the constraints. This is likely possible with the machinery developed in [Lall and West, 2006]. On the other hand, this is a boundary problem, not an initial value one, and recent work suggests that the benefit of symplectic integrators in the former context is more problem-dependent [Chyba et al., 2009].

Another puzzle: Why is non-symplectic Hamiltonian annealing is nearly as accurate as the symplectic form? This might be related to our choice of discretization: the non-symplectic trapezoidal integrator is a *conjugate* symplectic integrator, which does not exactly conserve bilinear forms, but still bounds their accumulating errors. Therefore, they produce high accuracy long-term integrations, and so in practice may be as reliable as symplectic integrators even for the boundary value problem calculated here. Still, note that the symplectic method consistently finds the global minimum at a higher rate than the nonsymplectic one (Tables 6.2 and 6.3). It is not unreasonable that this difference may manifest itself more drastically in other model systems.

Chapter 7

Final thoughts

We began this thesis with a look at a simple estimation problem: determining the frequency of an ideal oscillator using time-resolved measurements of its position. We have since confronted various methods to estimate dynamical systems that possess the same generic ingredients – underlying dynamical models, sparse observations – but are far more complex, nonlinear, or even chaotic. We saw that the filtering approach – powerful and elegant in its own right – can exhibit limitations susceptible to agnostic priors and dynamical nonlinearities. These two features together often rendered the sequential filter mathematically well-defined but functionally divergent. This feature magnified when the model description contained unknown parameters, which, lacking dynamics, do not fit as neatly into the sequential framework.

We then investigated the sequential Kalman filtering equations in the global, simultaneous time limit. The structure of the resulting high-dimensional Bayesian integral naturally suggested the method of steepest descent as a valid approximation, yet in practice this was hindered by the presence of nearly ubiquitous local minima. We then developed the *variational annealing* algorithm to more systematically track the massively dominant global minimum, relying on the fact that it will lie relatively close to a degenerate minimal valley centered around the measured data. This was shown to very effectively pinpoint the global minima in sparsely observed chaotic dynamical models.

An added benefit of the variational annealing approach is that parameters appear on equal footing as dynamical states, which are collocated in time. The result is that estimations of neural systems, where parameters dictating aspects such as the timescales of synaptic transmission, spiking frequencies, and neural excitability are of central interest. Variational annealing was thus shown to be very effective in determining the underlying parameters of a biophysical model of bursting neurons in the avian vocal pathway.

The second half of the thesis explored the idea of the estimation as a dynamical system in its own right, exhibiting subtle but important features when the underlying model is itself nonlinear or chaotic. The resulting "dynamics of nonlinear inference" can be cast as a path integral whose classical solutions are described by Euler-Lagrange or Hamilton's equations of motion – the evolution of this system in time relies on both i) the underlying stochastic dynamics *and* ii) the incoming observations. The downside of this formulation is that it is persistently unstable and, being a boundary value problem, is not particularly amenable to shooting methods.

We then explored a partial solution to this problem by utilizing the same notion exploited by variational annealing – that the dominant solutions to this ill-conditioned inverse problem can be tracked by incorporating nonlinearities in a systematic, gradual way. In this method, the optimization was cast in a space of canonical coordinates, with the nonlinearities residing entirely in Hamilton's equations of motion, which were enforced gradually as equality constraints.

Let us close with one proposition that may lead to some fruitful future research directions.

7.1 Hamiltonian variational annealing and chaotic synchronization

There is a notion, not yet explored in this thesis, that positive Lyapunov exponents of chaotic systems can be inverted by partially coupling to systems of similar dynamics. This idea is now known as *Pecora-Carroll synchronization*, after its developers [Pecora and Carroll, 1990]. That is, two Lorenz systems beginning at differing initial conditions can be synchronized together by directly substituting one component of the second system for the corresponding state in the first. A related method of the same vain, more common in the geophysical community, is to estimate the unmeasurable subspaces of a chaotic system by coupling instead to the measured data – the observation stream effectively acts as the "driving" system [Abarbanel, 2013, Lakshmivarahan and Lewis, 2013]. This is known as the *nudging* technique, and inverts the positive Lyapunov exponents in much the same way as the Pecora-Carroll method.

Both of these ideas function by simply integrated the nudged or synchronizing dynamics forward in time. As with the sequential filter, it is unclear how these ideas could be used to probe static parameters without adding them artificially as trivial dynamical states: $\dot{p}_i = 0$. A recent extension has taken the viewpoint that this idea could likewise be used for parameter estimation when couched in an optimization context [Abarbanel et al., 2009]. In fact, this technique, *dynamical state and parameter estimation* or DSPE, exhibits much of the same homotopy continuation-like aspects as variational annealing: the nudging couplings, now promoted to time-dependent parameters, serve as a "smoother" of the cost surface, guiding the estimate systematically to its narrow global minimum within a rough and nonconvex manifold.

Specifically, DSPE optimizes the following cost function:

$$A_{\mathrm{DSPE}}(\mathbf{x}_n, \mathbf{u}_n) = \sum_n^N (\mathbf{x}_n - \mathbf{y}_n)^2 \Big|_{\mathbf{R}^{-1}} + \sum_n^N \mathbf{u}_n^2, \tag{7.1}$$

subject to the hard constraints:

$$\mathbf{x}_{n+1} = \mathbf{g}(\mathbf{x}_n, \mathbf{p}) + \mathbf{u}_n^T \mathbf{H}(\mathbf{y}_n - \mathbf{x}_n), \tag{7.2}$$

where $\mathbf{p}$ are parameters, $\mathbf{g}(\cdot)$ is some discretization of the continuous dynamics $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{p})$, and $\mathbf{H}$ is a projection operator onto the measured subspace, $L < D$. The control variables $\mathbf{u}_n$ are independent for each time and have dimension equal to L. The idea of DSPE is that appreciable values of $\mathbf{u}_n$ smooth the cost function surface in the parameter subspace when walking along the integral curves of the constraint equations. This smoothing can localize the estimate around a single value of the unknown parameter, possibly inaccurate for larger $\mathbf{u}_n$, but as $\mathbf{u}_n$ are driven down by their quadratic penalty, the correct parameter is honed in on.

Interestingly, the idea that Lyapunov exponents can be made negative by coupling or synchronization is somewhat opposed to the running theme in this thesis, where instability is universal. The tradeoff is that stationary paths of the nonlinear inference path integrals are by definition the *maximal* contributions to a conditional expectation that in theory accounts for *all* possible correlations, while the couplings in DSPE do not represent, fundamentally, any optimal or dominant quantity, serving instead as a (powerful) numerical tool to remove instabilities. Still, the similarity of the nudged dynamics to the Kalman filter – specifically, a linear coupling to the residual $(\mathbf{y}_n - \mathbf{x}_n)$ – suggests that Kalman filter is actually an "optimal nudging" technique, where the coupling matrix maximally embodies statistical information about our confidence in the data, the model, and the current estimate. If this were true, however, it would nevertheless only hold when the dynamics are linear anyhow.

It would be interesting to ask what would occur if the constraint variables $\mathbf{u}_n$ in DSPE were demanded to be optimal. In fact, the necessary conditions for this situation are given by the Pontryagin maximum principle, a central result in the theory of optimal control [Liberzon, 2012]. When the equations of the maximum principle can be inverted in a proper way, what results is our old friend – the two-point Hamiltonian boundary value problem. It is peculiar that by beginning with control variables specifically designed to stabilize the unstable manifolds, and then demanding that these controls be optimal, the stability is itself compromised. Perhaps this is the price to pay for maximally incorporating global information across the entire observation window in a time-simultaneous fashion.

Fortunately, since there is a natural associated convex cost function, Eq. 7.1, we might instead leverage Hamiltonian annealing to pinpoint this global minimum. Writing the equations of motion for this system is in principle relatively straightforward, a direct application of the maximum principle to Eqs. 7.1 and 7.2. The formulation of these equations bears resemblance to linear-quadratic Gaussian control (LQGC), except that the dynamics are nonlinear [Liberzon, 2012]. There may be a correspondence in the linear limit. There is also a distinction between what occurs in this formulation and in Hamiltonian variational annealing as presented above: here, there is no stochasticity. If anything, the nudging terms now assumes the role of deviations from the deterministic flow, and perhaps the maximum principle provides the equivalent of an "optimal tuning" of the model precision $\mathbf{Q}^{-1}$. All this remains to be seen, but may provide a worthwhile direction for future research.

7.2 Chaos and correlations in Monte Carlo sampling

The intimate connection with instability and optimality is a running theme in this thesis. Locally unstable or globally chaotic trajectories may be of great use in numerical methods that rely on ergodicity. As an example, consider Monte Carlo sampling, where fast convergence to the target distribution relies on the quick decay of sample-to-sample correlations [Metropolis et al., 1953, MacKay, 2003]. One variant of traditional Markov Chain Monte Carlo is *Hybrid Monte Carlo*, also called *Hamiltonian Monte Carlo* (HMC) routine [Neal, 1996]. The idea of HMC is similar in spirit to Hamiltonian variational annealing: the distribution support is augmented with an equal number of artificial momenta variables, and the sampling routine alternates between i) evolution along the Hamiltonian flow in this joint $\mathbf{x}$-$\mathbf{p}$ space and ii) random walks in momentum space alone. This idea works because the acceptance probability is the exponential of the "energy" of this system, which is nearly conserved along Hamiltonian flows, and that these flows are far less correlated than random walks.

An extension of HMC developed recently is to leverage the freedom in the momentum distribution, which determines the “kinetic energy,” to enforce that the system moves along fully Hamiltonian chaotic flows [Kadakia, 2016]. The correlations are designed to fall exponentially, which may not occur if the original Hamiltonian dynamics follows a periodicity of the system. This *Chaotic Hybrid Monte Carlo* scheme (CHMC) was shown to substantially improve convergence time in high-dimensional normal distributions with non-sparse covariances.

Our experience in this thesis with Hamiltonian systems is that if they are not purely oscillatory, instability is guaranteed. The degree of the instability, however, depends on the exact dynamics – the local Lyapunov exponents still sum to zero, but may lie fairly close to the imaginary axis. The resulting dynamics is then largely oscillatory, preserving long-time correlations and slowing the sample convergence. Further extensions to HMC may exploit these ideas by designing acceptance probabilities reliant not just on energy differences but on local Lyapunov spectrums, biasing the sample routine toward unstable directions. The bias can then be removed, possibly at the cost of lower acceptance probability. Since HMC can move in relatively large steps without appreciably sacrificing acceptance probability, there may be an effective compromise between acceptance retention and correlation suppression [Beskos et al., 2013].

Bibliography

[Abarbanel et al., 2009] Abarbanel, H., Creveling, D., Farsian, R., and Kostuk, M. (2009). Dynamical state and parameter estimation. *SIAM Journal on Applied Dynamical Systems*, 8(4):1341–1381.

[Abarbanel, 2013] Abarbanel, H. D. (2013). *Predicting the Future: Completing Models of Observed Complex Systems*. Springer.

[Abramowitz and Stegun, 1964] Abramowitz, M. and Stegun, I. A. (1964). *Handbook of Mathematical Functions*. Dover.

[Allgower and Georg, 1990] Allgower, E. L. and Georg, K. (1990). *Numerical Continuation Methods: An Introduction*. Springer-Verlag.

[Arulampalam et al., 2002] Arulampalam, S., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for on-line nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188.

[Auger et al., 2013] Auger, F., Hilairet, M., Guerrero, J., Monmasson, E., Orlowska-Kowalska, T., and Katsura, S. (2013). Industrial applications of the Kalman filter: A review. *IEEE Transactions on Industrial Electronics*, 60(12):5458–5471.

[Bavdekar et al., 2011] Bavdekar, V. A., Deshpande, A. D., and Patwardhan, S. C. (2011). Identification of process and measurement noise covariance for state and parameter estimation using extended Kalman filter. *Journal of Process Control*, 21:585–601.

[Bennett, 1992] Bennett, A. F. (1992). *Inverse Methods in Physical Oceanography*. Cambridge University Press, Cambridge, UK.

[Beskos et al., 2013] Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.-M., and Stuart, A. (2013). Optimal tuning of hybrid Monte Carlo algorithm. *Bernoulli*, 19:1501–1534.

[Betts, 1998] Betts, J. T. (1998). Survey of numerical methods for trajectory optimization. *Journal of Guidance Control and Dynamics*, 21:193–207.

[Betts, 2010] Betts, J. T. (2010). *Practical methods for optimal control and estimation using nonlinear programming*. Siam, Philadelphia, PA, USA.

[Bishop, 1995] Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Oxford University Press, Inc., New York, NY, USA.

[Björck, 1996] Björck, Å. (1996). *Numerical Methods for Least Squares Problems*. Society for Industrial and Applied Mathematics, Philadelphia, PA.

[Brainard and Doupe, 2002] Brainard, M. S. and Doupe, A. J. (2002). What songbirds teach us about learning. *Nature*, 417(6886):351–358.

[Breen et al., 2016] Breen, D., Shirman, S., Armstrong, E., Kadakia, N., and Abarbanel, H. (2016). HVC interneuron properties from statistical data assimilation. *arXiv:1608.04433 [q-bio.NC]*.

[Broyden, 1970] Broyden, C. G. (1970). The convergence of a class of double-rank minimization algorithms 1. General considerations. *IMA Journal of Applied Mathematics*, 6(1):76–90.

[Cardin et al., 2004] Cardin, J. A., Raksin, J. N., and Schmidt, M. F. (2004). The sensorimotor nucleus NIf is necessary for auditory processing but not vocal motor output in the avian song system. *Journal of Neurophysiology*, 93:2157–2166.

[Catterall et al., 2012] Catterall, W. A., Raman, I. M., Robinson, H. P. C., Sejnowski, T. J., and Paulsen, O. (2012). The Hodgkin-Huxley heritage: From channels to circuits. *The Journal of Neuroscience*, 32(41).

[Chaichian and Demichev, 2001] Chaichian, M. and Demichev, A. (2001). *Path Integrals in Physics: Volume I Stochastic Processes and Quantum Mechanics*. CRC Press.

[Channell and Scovel, 1990] Channell, P. J. and Scovel, C. (1990). Symplectic integration of Hamiltonian systems. *Nonlinearity*, 3(2):231–259.

[Chow and Buice, 2015] Chow, C. C. and Buice, M. A. (2015). Path integral methods for stochastic differential equations. *The Journal of Mathematical Neuroscience*, 5(1):8.

[Chyba et al., 2009] Chyba, M., Hairer, E., and Vilmart, G. (2009). The role of symplectic methods in optimal control. *Optimal Control Applications and Methods*, 30:367–382.

[Coleman and Vu, 2005] Coleman, M. J. and Vu, E. T. (2005). Recovery of impaired songs following unilateral but not bilateral lesions of nucleus uvaeformis of adult zebra finches. *Journal of Neurobiology*, 63(1):70–89.

[Cormen et al., 2001] Cormen, T. H., Stein, C., Rivest, R. L., and Leiserson, C. E. (2001). *Introduction to Algorithms*. McGraw-Hill Higher Education, 2nd edition.

[Covey and Carter, 2015] Covey, E. and Carter, M. (2015). *Basic Electrophysiological Methods*. Oxford University Press.

[Daou et al., 2013] Daou, A., Ross, M. T., Johnson, F., Hyson, R. L., and Bertram, R. (2013). Electrophysiological characterization and computational models of HVC neurons in the zebra finch. *Journal of Neurophysiology*, 110:1227–1245.

[Dayan and Abbott, 2005] Dayan, P. and Abbott, L. F. (2005). *Theoretical Neuroscience: Computational And Mathematical Modeling of Neural Systems*. Massachusetts Institute of Technology Press.

[Dekker, 1978] Dekker, H. (1978). Proof of the path integral representation of the nonlinear Fokker-Planck equation by means of Fourier series. *Physics Letters*, 65A(5-6):388–390.

[Dempster et al., 1977] Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B (Methodological)*, 39(1):1–38.

[Doucet et al., 2001] Doucet, A., de Freitas, N., and Gordon, N. (2001). *Sequential Monte Carlo methods in practice*. Springer-Verlag.

[Doucet and Johansen, 2009] Doucet, A. and Johansen, A. M. (2009). A tutorial on particle filtering and smoothing: fifteen years later. In Crisan, D. and Rozovskii, B., editors, *The Oxford Handbook of Nonlinear Filtering*.

[Duane et al., 1987] Duane, S., Kennedy, A., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Physics Letters B*, 195(2):216 – 222.

[Ermentrout and Terman, 2010] Ermentrout, G. B. and Terman, D. H. (2010). *Mathematical Foundations of Neuroscience*. Interdisciplinary Applied Mathematics. Springer, New York, NY.

[Evensen, 1994] Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99:10,143–10,162.

[Evensen, 2003] Evensen, G. (2003). The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343–367.

[Evensen, 2009] Evensen, G. (2009). *Data Assimilation: The Ensemble Kalman Filter*. Springer.

[Fee et al., 2004] Fee, M. S., Kozhevnikov, A. A., and Hahnloser, R. H. R. (2004). Neural mechanisms of vocal sequence generation in the songbird. *Annals of the New York Academy of Sciences*, 1016(1):153–170.

[Feng et al., 2014] Feng, B., Fu, M., and Ma, H. (2014). Kalman filter with recursive covariance estimation–sequentially estimating process noise covariance. *IEEE Transactions on Industrial Electronics*, 61(11):6253–6263.

[Feynman and Hibbs, 1965] Feynman, R. P. and Hibbs, A. R. (1965). *Quantum Mechanics and Path Integrals*. Dover Publications.

[FitzHugh, 1961] FitzHugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445.

[Fletcher, 1970] Fletcher, R. (1970). A new approach to variable metric algorithms. *The Computer Journal*, 13(3):317.

[Fletcher, 1987] Fletcher, R. (1987). *Practical Methods of Optimization*. John Wiley and Sons, 2 edition.

[Gardiner, 1985] Gardiner, C. (1985). *Handbook of Stochastic Methods*. Springer-Verlag, Berlin.

[Gelman et al., 2013] Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. Chapman and Hall/CRC, 3 edition.

[Gibb et al., 2009] Gibb, L., Gentner, T. Q., and Abarbanel, H. D. (2009). Inhibition and recurrent excitation in a computational model of sparse bursting in song nucleus hvc. *Journal of Neurophysiology*, 102(3):1748–1762.

[Gilks et al., 1995] Gilks, W., Richardson, S., and Spiegelhalter, D. (1995). *Markov Chain Monte Carlo in Practice*. Chapman and Hall.

[Goldin et al., 2013] Goldin, M. A., Alonso, L. M., Alliende, J. A., Goller, F., and Mindlin, G. B. (2013). Temperature induced syllable breaking unveils nonlinearly interacting timescales in birdsong motor pathway. *PLOS ONE*, 8(6):1–10.

[Goldstein et al., 2002] Goldstein, H., Poole, C., and Safko, J. (2002). *Classical Mechanics*. Addison Wesley.

[Gordon et al., 1993] Gordon, N., Salmond, D., and Smith, A. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEEE Proceedings F, Radar and Signal Processing*, 140(2):107–113.

[Grewal and Andrews, 2014] Grewal, M. S. and Andrews, A. P. (2014). *Kalman Filtering: Theory and Practice with MATLAB*. Wiley, 4 edition.

[Hahnloser et al., 2002] Hahnloser, R. H. R., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419(6902):65–70.

[Hairer et al., 2006] Hairer, E., Wanner, G., and Lubich, C. (2006). *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer Series in Computational Mathematics; Volume 31, 2nd edition.

[Hansen, 2010] Hansen, P. C. (2010). *Discrete Inverse Problems: Insight and Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA.

[Hille, 2001] Hille, B. (2001). *Ion Channels of Excitable Membranes*. Sinauer Associates.

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

[Houtekamer and Zhang, 2016] Houtekamer, P. L. and Zhang, F. (2016). Review of the ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 144(12):4489–4532.

[Hunt and Ross, 1981] Hunt, K. L. and Ross, J. (1981). Path integral solutions of stochastic equations for nonlinear irreversible processes: The uniqueness of the thermodynamic lagrangian. *Journal of Chemical Physics*, 52(2):976–984.

[Huys and Paninski, 2009] Huys, Q. J. M. and Paninski, L. (2009). Smoothing of, and parameter estimation from, noisy biophysical recordings. *PLOS Computational Biology*, 5:1–16.

[Izhikevich, 2007] Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience : The Geometry of Excitability and Bursting*. Computational neuroscience. MIT Press, Cambridge, MA, London, UK.

[Jin et al., 2007] Jin, D. Z., Ramazanoglu, F. M., and Seung, H. S. (2007). 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.

[Johnson and Wichern, 2007] Johnson, R. A. and Wichern, D. W., editors (2007). *Applied Multivariate Statistical Analysis*. Pearson, 6 edition.

[Johnston and Wu, 1994] Johnston, D. and Wu, S. M.-S. (1994). *Foundations of Cellular Neurophysiology*. MIT Press/Bradford Books, 1 edition.

[Jordan and Smith, 1999] Jordan, D. and Smith, P. (1999). *Nonlinear Ordinary Differential Equations*. Oxford University Press, 3 edition.

[Julier et al., 2000] Julier, S., Uhlmann, J., and Durrant-Whyte, H. F. (2000). A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic Control*, 45(3):477–482.

[Julier and Uhlmann, 1997] Julier, S. J. and Uhlmann, J. K. (1997). New extension of the Kalman filter to nonlinear systems. In Kadar, I., editor, *Signal Processing, Sensor Fusion, and Target Recognition VI*, volume 3068 of *Proceedings of SPIE*, pages 182–193.

[Kadakia, 2016] Kadakia, N. (2016). Hybrid Monte Carlo with chaotic mixing. *arXiv:1604.07343 [physics.data-an]*.

[Kadakia et al., 2016] Kadakia, N., Armstrong, E., Breen, D., Morone, U., Margoliash, D., and Abarbanel, H. D. (2016). Nonlinear statistical data assimilation for $\mathrm{HVC_{RA}}$ nuerons in the avian song system. 110(6):417–434.

[Kadakia et al., 2017] Kadakia, N., Rey, D., Ye, J., and Abarbanel, H. D. I. (2017). Symplectic structure of statistical data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 143(703):756–771.

[Kalman, 1960] Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45.

[Kalnay, 2002] Kalnay, E. (2002). *Atmospheric Modeling, Data Assimilation, and Predictability*. Cambridge University Press.

[Kantas et al., 2015] Kantas, N., Doucet, A., Singh, S. S., Maciejowski, J., and Chopin, N. (2015). On particle methods for parameter estimation in state-space models. *Statistical Science*, 30(3):328–351.

[Karasalo and Hu, 2011] Karasalo, M. and Hu, X. (2011). An optimization approach to adaptive Kalman filtering. *Automatica*, 47:1785–1793.

[Kirk, 1970] Kirk, D. E. (1970). *Optimal Control Theory: An Introduction*. Dover Publications, Inc.

[Koch and Segev, 1998] Koch, C. and Segev, I. (1998). *Methods in Neural Modeling: From Ions to Networks*. MIT Press.

[Kosche et al., 2015] Kosche, G., Vallentin, D., and Long, M. A. (2015). Interplay of inhibition and excitation shapes a premotor neural sequence. *Journal of Neuroscience*, 35(3):1217–1227.

[Kostuk et al., 2012] Kostuk, M., Toth, B. A., Meliza, C. D., Margoliash, D., and Abarbanel, H. D. (2012). Dynamical estimation of neuron and network properties II: path integral Monte Carlo methods. *Biological Cybernetics*, 106(3):155–167.

[Kubota and Taniguchi, 1998] Kubota, M. and Taniguchi, I. (1998). Electrophysiological characteristics of classes of neuron in the HVc of the zebra finch. *Journal of Neurophysiolology*, 80(2):914–923.

[Lakshmivarahan and Lewis, 2013] Lakshmivarahan, S. and Lewis, J. M. (2013). Nudging methods: A critical overview. In Park, S. K. and Xu, L., editors, *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications, Volume 2*. Springer.

[Lall and West, 2006] Lall, S. and West, M. (2006). Discrete variational Hamiltonian mechanics. *Journal of Physics A*, 39(19):5509.

[Laplace, 1986] Laplace, P.-S. (1986). Memoir on the probability of the causes of events. *Statistical Science*, 1(3):364–378. Original work published 1774.

[LeCun et al., 2015] LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436 – 444.

[Leok and Zhang, 2011] Leok, M. and Zhang, J. (2011). Discrete Hamiltonian variational integrators. *IMA Journal of Numerical Analysis*, 31(4):1497–1532.

[Li, 1997] Li, T. (1997). Numerical solution of multivariate polynomial systems by homotopy continuation methods. *Acta Numerica*, 6:399–436.

[Liberzon, 2012] Liberzon, D. (2012). *Calculus of Variations and Optimal Control Theory*. Princeton University Press, Princeton, NJ, USA.

[Long and Fee, 2008] Long, M. A. and Fee, M. S. (2008). Using temperature to analyse temporal dynamics in the songbird motor pathway. *Nature*, 456(7219):189–194.

[Long et al., 2010] Long, M. A., Jin, D. Z., and Fee, M. S. (2010). Support for a synaptic chain model of neuronal sequence generation. *Nature*, 48(7322):394–399.

[Lorenz, 2006] Lorenz, E. N. (2006). Predictability: A problem partly solved. In Palmer, T. and Hagedorn, R., editors, *Predictability of Weather and Climate*. Cambridge.

[MacKay, 2003] MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.

[Marsden and West, 2001] Marsden, J. E. and West, M. (2001). Discrete mechanics and variational integrators. *Acta Numerica*, 10:357–514.

[Meliza et al., 2014] Meliza, C., Kostuk, M., Huang, H., Nogaret, A., Margoliash, D., and Abarbanel, H. (2014). Estimating parameters and predicting membrane voltages with conductance-based neuron models. *Biological Cybernetics*, 108(4):495–516.

[Meng et al., 2011] Meng, L., Kramer, M. A., and Eden, U. T. (2011). A sequential Monte Carlo approach to estimate biophysical neural models from spikes. *Journal of Neural Engineering*, 8(6):065006.

[Metropolis et al., 1953] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6).

[Mitchell et al., 2002] Mitchell, H. L., Houtekamer, P., and Pellerin, G. (2002). Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Monthly Weather Review*, 130:2791âĂŞ2808.

[Mooney, 2000] Mooney, R. (2000). Different subthreshold mechanisms underlie song selectivity in identified HVc neurons of the zebra finch. *Journal of Neuroscience*, 20:5420–5436.

[Mooney and Prather, 2005] Mooney, R. and Prather, J. F. (2005). The hvc microcircuit: the synaptic basis for interactions between song motor and vocal plasticity pathways. *Journal of Neuroscience*, 25(8):1952–1964.

[Nagumo et al., 1962] Nagumo, J., Arimoto, S., and Yoshizawa, S. (1962). An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50:2061–2070.

[Nakano et al., 2007] Nakano, S., Ueno, G., and Higuchi, T. (2007). Merging particle filter for sequential data assimilation. *Nonlinear Processes in Geophysics*, 14:395–408.

[Natvik and Evensen, 2003] Natvik, L.-J. and Evensen, G. (2003). Assimilation of ocean colour data into a biochemical model of the North Atlantic Part 1. data assimilation experiments. *Journal of Marine Systems*, 40–41:127–153.

[Neal, 1996] Neal, R. M. (1996). *Bayesian Learning for Neural Networks*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.

[Nocedal and Wright, 2006] Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer, New York, NY, 2 edition.

[Northrop, 1961] Northrop, T. G. (1961). The guiding center approximation to charged particle motion. *Annals of Physics*, 15:79–101.

[Okanoya, 2004] Okanoya, K. (2004). The Bengalese Finch: A window on the behavioral neurobiology of birdsong syntax. *Annals of the New York Academy of Sciences*, 1016(1):724–735.

[Onsager and Machlup, 1953] Onsager, L. and Machlup, S. (1953). Fluctuations and irreversible processes. *Physical Review*, 91(6):1505.

[Pecora and Carroll, 1990] Pecora, L. M. and Carroll, T. L. (1990). Synchronization in chaotic systems. *Physical Review Letters*, 64(8):821–824.

[Pissanetzky, 1984] Pissanetzky, S. (1984). *Spare Matrix Technology*. Elsevier.

[Rao, 2010] Rao, A. V. (2010). A survey of numerical methods for optimal control. *Advances in the Astronautical Sciences*, 135(1):497–52.

[Rauch et al., 1965] Rauch, H., Tung, F., and Striebel, C. (1965). *AIAA Journal*, 3(8):1445–1450.

[Sakmann and Neher, 1984] Sakmann, B. and Neher, E. (1984). Patch clamp techniques for studying ionic channels in excitable membranes. *Annual Review of Physiology*, 46:455–472.

[Särkkä, 2013] Särkkä, S. (2013). *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics textbooks. Cambridge University Press, Cambridge, UK, New York.

[Snyder et al., 2008] Snyder, C., Bengtsson, T., Bickel, P., and Anderson, J. (2008). Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136:4629–4640.

[Sompolinksy et al., 1988] Sompolinksy, H., Crisanti, A., and Sommers, H. (1988). Chaos in random neural networks. *Physical Review Letters*, 61(3).

[Sprott, 2010] Sprott, J. C. (2010). *Elegant Chaos: Algebraically Simple Chaotic Flows*. World Scientific.

[Sterratt et al., 2011] Sterratt, D., Graham, B., Gillies, A., and Willshaw, D. (2011). *Principles of Computational Modelling in Neuroscience*. Cambridge University Press, Cambridge, UK.

[Strogatz, 2001] Strogatz, S. H. (2001). *Nonlinear Dynamics And Chaos*. 1 edition.

[Tenenbaum and Pollard, 1985] Tenenbaum, M. and Pollard, H. (1985). *Ordinary Differential Equations*. Dover Publications.

[Toth et al., 2011] Toth, B. A., Kostuk, M., Meliza, C. D., Margoliash, D., and Abarbanel, H. D. (2011). Dynamical estimation of neuron and network properties I: variational methods. *Biological Cybernetics*, 105(3-4):217–237.

[Ullah and Schiff, 2009] Ullah, G. and Schiff, S. J. (2009). Tracking and control of neuronal hodgkin-huxley dynamics. *Physical Review E*, 79:040901.

[van Vreeswijk and Sompolinsky, 1996] van Vreeswijk, C. and Sompolinsky, H. (1996). Chaos in neural networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726.

[Vavoulis et al., 2012] Vavoulis, D. V., Straub, V. A., Aston, J. A. D., and Feng, J. (2012). A self-organizing state-space-model approach for parameter estimation in Hodgkin-Huxley-type models of single neurons. *PLOS Computational Biology*, 8(3):e1002401.

[Vogel, 2002] Vogel, C. R. (2002). *Computational Methods for Inverse Problems*. Society for Industrial and Applied Mathematics, Philadelphia, PA.

[Vogelstein et al., 2009] Vogelstein, J. T., Watson, B. O., Packer, A. M., Yuste, R., Jedynak, B., and Paninski, L. (2009). Spike inference from calcium imaging using sequential Monte Carlo methods. *Biophysical Journal*, 97(2):636–655.

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

[Wan and van der Merwe, 2002] Wan, E. A. and van der Merwe, R. (2002). The unscented Kalman filter. In Haykin, S., editor, *Kalman Filtering and Neural Networks*, pages 221 – 280. Wiley.

[Wendlandt and Marsden, 1997] Wendlandt, J. M. and Marsden, J. E. (1997). Mechanical integrators derived from a discrete variational principle. *Physica D*, 106:223–246.

[Wisdom and Holman, 1991] Wisdom, J. and Holman, M. (1991). Symplectic maps for the N-body problem. *Astronomical Journal*, 102:1528.

[Ye et al., 2015a] Ye, J., Kadakia, N., Rozdeba, P. J., Abarbanel, H. D. I., and Quinn, J. C. (2015a). Improved variational methods in statistical data assimilation. *Nonlinear Processes in Geophysics*, 22:205–213.

[Ye et al., 2015b] Ye, J., Rey, D., Kadakia, N., Eldridge, M., Morone, U., Rozdeba, P., Abarbanel, H. D. I., and Quinn, J. C. (2015b). A systematic variational method for statistical nonlinear state and parameter estimation. *Physical Review E*, 92(5):052901.

[Yoshida, 1990] Yoshida, H. (1990). Construction of higher-order symplectic integrators. *Physics Letters A*, (5-7):262–268.

[Yu and Margoliash, 1996] Yu, A. C. and Margoliash, D. (1996). Temporal hierarchical control of singing in birds. *Science*, 273:1871–1875.

[Zelinka et al., 2014] Zelinka, I., Zmeskal, O., and Saloun, P. (2014). Fractal analysis of fitness landscapes. In *Recent Advances in the Theory and Application of Fitness Landscapes*, pages 427–456. Springer.