Systematic variational method for statistical nonlinear state and parameter estimation

Jingxin Ye, Daniel Rey, Nirag Kadakia, Michael Eldridge, Uriel I. Morone, Paul Rozdeba, Henry D. I. Abarbanel, and John C. Quinn

Phys. Rev. E **92**, 052901 — Published 2 November 2015

DOI: 10.1103/PhysRevE.92.052901

A Systematic Variational Method for Statistical Nonlinear State and Parameter Estimation

Jingxin Ye, Daniel Rey, Nirag Kadakia, Michael Eldridge, Uriel I. Morone, and Paul Rozdeba
Department of Physics, University of California, San Diego, La Jolla, CA 92093-0374, USA

Henry D. I. Abarbanel*
Department of Physics, University of California,
San Diego, La Jolla, CA 92093-0374, USA and
Marine Physical Laboratory (Scripps Institution of Oceanography) University of California, San Diego
La Jolla, CA 92093-0374, USA

John C. Quinn
Intellisis Corporation, 10350 Science Center Drive, Suite 140 San Diego, CA 92121, USA
(Dated: October 15, 2015)

In statistical data assimilation one evaluates the conditional expected values, conditioned on measurements, of interesting quantities on the path of a model through observation and prediction windows. This often requires working with very high dimensional integrals in the discrete time descriptions of the observations and model dynamics, which become functional integrals in the continuous time limit. Two familiar methods for performing these integrals include (1) Monte Carlo calculations and (2) variational approximations using the method of Laplace plus perturbative corrections to the dominant contributions. We attend here to aspects of the Laplace approximation and develop an annealing method for locating the variational path satisfying the Euler-Lagrange equations that comprises the major contribution to the integrals. This begins with the identification of the minimum action path starting with a situation where the model dynamics is totally unresolved in state space, and the consistent minimum of the variational problem is known. We then proceed to slowly increase the model resolution, seeking to remain in the basin of the minimum action path, until a path that gives the dominant contribution to the integral is identified. After a discussion of some general issues, we give examples of the assimilation process for some simple, instructive models from the geophysical literature. Then we explore a slightly richer model of the same type with two distinct time scales. This is followed by a model characterizing the biophysics of individual neurons.

PACS numbers: 05.10.Ln, 05.45.Tp, 92.60.Ry, 87.19.II

I. INTRODUCTION

The transfer of information from observations to quantitative, predictive models of the observed physical system is a challenge in a broad spectrum of fields including meteorology [1, 2], geochemistry [3], and systems biology [4], among many others.

In each temporal observation window, it is usual that only a sparse set of the D-dimensional model dynamical variables $\mathbf{x}(t)$ are measured. From L-dimensional observations $y_l(t_n)$; $l = 1, 2, ..., L$ at times $t_n = \{t_0, t_1, ..., t_m = t_f\}$, we must estimate the full D-dimensional state $x_a(t)$; $a = 1, 2, ..., D$. Typically the measurements are sparse, $L \ll D$; We must also estimate any unknown fixed parameters in the model or in the measurement functions $h_l(\mathbf{x}(t))$, $l = 1, 2, ..., L$ relating the model output $\mathbf{x}(t)$ to the observations $y_l(t)$.

The use of measurements in the window $[t_0, t_f]$ to estimate N_p time-independent parameters $\mathbf{p} = \{p_1, p_2, \ldots, p_{N_p}\}$ and unknown states completes the model [5] and permits us to test or validate the model through predictions for $t > t_f$ where a selected metric compares new observations $\mathbf{y}(t > t_f)$ to new model outputs $\mathbf{h}(\mathbf{x}(t > t_f))$.

Since the measurements are always noisy and the model always has errors, one requires a statistical description of the assimilation of information from the observations into the model. Such a description is based on estimating the conditional distribution of states and parameters $P(\mathbf{x}(t_n)|\mathbf{Y}(n))$ conditioned on measurements up to time t_n : $\mathbf{Y}(n) = \{\mathbf{y}(t_0), ..., \mathbf{y}(t_{n-1}), \mathbf{y}(t_n)\}$ to predict the conditional distribution for $t > t_f$. Our focus is not on the evaluation of $P(\mathbf{x}(t_n)|\mathbf{Y}(n))$ itself, but on the quantities of physical (or biological) interest: the conditional expected value of functions $G(\mathbf{X})$ along the path $\mathbf{X} = \{\mathbf{x}(t_0), ...\mathbf{x}(t_{m-1}), \mathbf{x}(t_m), \mathbf{p}\}$ of the state through the observation window and

*Electronic address: habarbanel@ucsd.edu

beyond. $\mathbf{X}$ is a $(m+1)D+N_p$ dimensional vector. Writing the conditional probability distribution $P(\mathbf{X}|\mathbf{Y}) = P(\mathbf{x}(t_0),...,\mathbf{x}(t_m)|\mathbf{y}(t_0),...,\mathbf{y}(t_m))$ as $\exp[-A_0(\mathbf{X})]$, the expected value of $G(\mathbf{X})$ is

$$E[G(\mathbf{X})|\mathbf{Y}] = \frac{\int d\mathbf{X}\, G(\mathbf{X}) \exp[-A_0(\mathbf{X})]}{\int d\mathbf{X}\, \exp[-A_0(\mathbf{X})]}. \tag{1}$$

This expression for $P(\mathbf{X}|\mathbf{Y})$ defines the "action" $A_0(\mathbf{X})$.

We do not further explicitly show the dependence of the action on the measurements. In discrete time this is an integral of dimension D times the number of discrete time steps in the observation plus prediction windows. In continuous time it is a functional path integral [6, 7].

Important examples of $G(\mathbf{X})$ include the mean or expected path, in which case $G(\mathbf{X}) = \mathbf{X}$, moments about this expected path, and marginal distributions of, say, $x_b(t_k)$ in which case $G(\mathbf{X}) = \delta(\theta - x_b(k))$ giving $P(\theta)$.

The action $A_0(\mathbf{X})$ has the exact representation [5] in discrete time

$$\begin{aligned} A_0(\mathbf{X}) = &- \sum_{n=0}^{m-1} \ln[P(\mathbf{x}(n+1)|\mathbf{x}(n))] - \ln[P(\mathbf{x}(0))] \\ &- \sum_{n=0}^{m} \ln\left[\frac{P(\mathbf{x}(n),\mathbf{y}(n)|\mathbf{Y}(n-1))}{P(\mathbf{y}(n)|\mathbf{Y}(n-1))\, P(\mathbf{x}(n)|\mathbf{Y}(n-1))}\right] \\ = &- \sum_{n=0}^{m-1} \ln[P(\mathbf{x}(n+1)|\mathbf{x}(n))] - \ln[P(\mathbf{x}(0))] \\ &- \sum_{n=0}^{m} \ln\Big[P(\mathbf{y}(n)|\mathbf{x}(n),\mathbf{Y}(n-1))\Big] \\ &+ \text{terms independent of } \mathbf{X} \end{aligned}$$

in which $\mathbf{x}(n) = \mathbf{x}(t_n)$. $P(\mathbf{x}(n+1)|\mathbf{x}(n))$ is the transition probability for the state $\mathbf{x}(n)$ at time t_n to arrive at the state $\mathbf{x}(n+1)$ at time t_{n+1}, and $P(\mathbf{x}(0))$ is the distribution of the state at time t_0 when observations commence. The dynamics moving the model state $\mathbf{x}(n)$ through time resides in $P(\mathbf{x}(n+1)|\mathbf{x}(n))$.

The deterministic dynamics of the model state is taken to satisfy

$$\frac{dx_a(t)}{dt} = F_a(\mathbf{x}(t));\ a = 1, 2, ..., D, \tag{2}$$

in continuous time, and this becomes

$$x_a(n+1) = f_a(\mathbf{x}(n)) \tag{3}$$

in discrete time.

The discretization in time of Equation (2) can be explicit or implicit. Either choice defines an action $A_0(\mathbf{X})$ as a function of the components of $\mathbf{X}$. In this paper we will focus on the extrema of the action in $\mathbf{X}$ varying all components of $\mathbf{X}$ in both the measurement error term and the model error term. Throughout this paper, the trapezoidal rule is used to discretize the model equations,

$$x_a(n+1) = x_a(n) + \frac{\Delta t}{2}[F_a(\mathbf{x}(n+1)) + F_a(\mathbf{x}(n))]. \tag{4}$$

The model equations could then be written as a function of both $\mathbf{x}(n)$ and $\mathbf{x}(n+1)$: $g_a(\mathbf{x}(n), \mathbf{x}(n+1), \mathbf{p}) = 0$. When there is no model error, $P(\mathbf{x}(n+1)|\mathbf{x}(n)) = \delta^{D+N_P}(\mathbf{g}(x(n), \mathbf{x}(n+1), \mathbf{p}))$.

Once a model has been selected by physical considerations, and choices are made for the distributions of errors in the measurements and of errors in the model, the central challenge of statistical data assimilation is the estimation of the path integral Eq. (1).

One approach is to use Monte Carlo methods for evaluating the high dimensional integral [8–11].

In this paper we investigate results associated with the estimation of the integral Eq. (1) using Laplace's method [12]. This is a variational calculation seeking extremum paths $\mathbf{X}^q$ of the action, where the Jacobian

$$\left.\frac{\partial A_0(\mathbf{X})}{\partial \mathbf{X}}\right|_{\mathbf{X}=\mathbf{X}^q} = 0, \quad q = 0, 1, ... , \tag{5a}$$

and the Hessian

$$\left.\frac{\partial^2 A_0(\mathbf{X})}{\partial \mathbf{X}^2}\right|_{\mathbf{X}=\mathbf{X}^q} \text{ is positive definite.} \tag{5b}$$

Laplace's method also allows the evaluation of corrections to Eq. (1) using any of the $\mathbf{X}^q$ as the leading approximation for the expected value. Furthermore, when $\mathbf{f}(\mathbf{x})$ is not linear in the model state variables, there may be multiple solutions to Eq. (5), and we must select which among them provide the most important contributions to the integral. Finding the extremum path is widely known as 4DVar [1, 2] in the geophysical literature.

Paths with distinct values of $A_0(\mathbf{X})$ lead to exponentially different contributions to the expected value path integral Eq. (1). The path which gives the smallest value of the action,

$$A_0(\mathbf{X}^0) = \min_{\mathbf{X}^q} A_0(\mathbf{X}^q), \tag{6}$$

$\mathbf{X}^0$ is the conditional mode of the distribution and also the maximum likelihood estimate. When the action level $A_0(\mathbf{X}^0)$ is much less than any other action level coming from paths $\mathbf{X}^{q\neq 0}$, it exponentially dominates the integral.

In nonlinear problems the search for $\mathbf{X}^0$ requires some care, as we will show in the core of this paper. Our goal here is to expand on the details of an annealing method introduced in Ref.[13] to locate the saddle paths with the smallest action. Further, we estimate the corrections to the approximation of retaining only $\mathbf{X}^0$ as dominating the integral Eq. (5).

Additionally, we use the annealing method to explore several interesting nonlinear dynamical models not addressed in our earlier work. These include a model with combined "fast" and "slow" time scales of the dynamics as well as a standard Hodgkin-Huxley model of an isolated neuron. The latter is a prelude for more complex, biophysically realistic neuron models both in isolation and within functional networks with biophysical connections among them. A final section is devoted to comparing this work with the QSVA method [14], which seeks to solve a related problem by manipulating the length of the assimilation window.

II. FORMULATION OF THE PROBLEM IN DISCRETE TIME AND IN CONTINUOUS TIME

A. Discrete Time

To simplify the discussion we make two familiar assumptions about how the measurement errors and the model errors enter the expression for the action $A_0(\mathbf{X})$: the errors in each are taken to be distributed as a Gaussian,

- the measurement error enters with an inverse covariance matrix $\mathbf{R}_m(l,k,t) = R_m(l,t)\delta_{lk},\ k,l = 1,2,...,L,$
- and the model error enters with an inverse covariance matrix $\mathbf{R}_f(a,b) = R_f(a)\delta_{ab},\ a,b = 1,2,...,D.$

As the dynamics $\mathbf{x}(n+1) = \mathbf{f}(\mathbf{x}(n))$ is nonlinear, the overall path integral is not Gaussian with this choice for the error distributions.

In the presence of Gaussian additive model error, the dynamics satisfies the D-dimensional stochastic discrete time map

$$x_a(n+1) = f_a(\mathbf{x}(n)) + R_f(a)^{-1/2}\eta_a(n), \tag{7}$$

and each component $\eta_a(n)$ is Gaussian distributed as $\mathcal{N}(0,1)$. Here we assume there is no cross-covariance between different state variables, and the case with cross-covariances needs further study. We also assume the measurement function $h_l(\mathbf{x}(t)) = x_l(t_n),\ n = 0,1,\ldots,m.$

The Gaussian error action $A_0(\mathbf{X})$ in discrete time takes the form

$$\begin{aligned} A_0(X) = &\sum_{n=0}^{m}\sum_{l=1}^{L}\frac{R_m(l,n)}{2}[x_l(n) - y_l(n)]^2 - \ln[P(\mathbf{x}(0))] \\ &+ \sum_{n=0}^{m-1}\sum_{a=1}^{D}\frac{R_f(a)}{2}[x_a(n+1) - f_a(\mathbf{x}(n))]^2. \end{aligned} \tag{8}$$

The distribution of initial states $P(\mathbf{x}(0))$ in the action is often assumed to be uniformly distributed or Gaussian distributed. For the uniform distribution case $-\ln[P(\mathbf{x}(0)]$ is a constant and cancels between the numerator and

denominator of expected values Eq. (1). When $P(\mathbf{x}(0))$ is Gaussian, suppose the variation of $\mathbf{x}(0)$ is given about some base state $\mathbf{x}_{\rm base}$, so $-\ln[P(\mathbf{x}(0)] = (\mathbf{x}(0) - \mathbf{x}_{\rm base})^2 R_{\rm base}/2$. This has the form of the measurement term evaluated at $n = 0$. This expression can be incorporated into the term with coefficient R_m in the action.

We no longer display $-\ln[P(\mathbf{x}(0))]$ in the following discussion. The resulting action Eq. (8) we call the Gaussian error action. Many other actions may be of physical interest, and they depend in detail on our representation of errors in the measurements and errors in the model.

B. Continuous Time

Although all calculations are performed in discrete time, we use continuous time to gather insight into the saddle paths for the action, and return to discrete time with the lessons in mind. Let time become continuous between the initiation of observations at time t_0 and the completion of measurements at time t_f, we identify the action in continuous time as

$$A_0(\mathbf{x}(t)) = \int_{t_0}^{t_f} dt\, \mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t) \tag{9}$$

where the Lagrangian $\mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t)$, also called the Onsager-Machlup functional [15], is

$$\begin{aligned}\mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t) = &\sum_{l=1}^{L} \frac{R_m(l,t)}{2}[x_l(t) - y_l(t)]^2 \\ &+ \sum_{a=1}^{D} \frac{R_f(a)}{2}[\dot{x}_a(t) - F_a(\mathbf{x}(t))]^2.\end{aligned}$$

$R_m(l,t)$ is nonzero only near the observations times $t \approx t_n$.

The transition from discrete to continuous time has a subtlety which we note here, then bypass as we will return to discrete time for all of our examples and for any applications. As carefully explained in the book of Zinn-Justin [7] and the papers [16, 17] there is a Jacobian involved in the transformation from discrete to continuous time which affects the action through a term involving the divergence of the vector field $\mathbf{F}(\mathbf{x})$. It adds the term

$$\Theta(0) \int_{t_0}^{t_f} dt\, \nabla_x \cdot \mathbf{F}(\mathbf{x}(t)), \tag{10}$$

to the action, where $\Theta(x)$ is the Heaviside function. The value of the quantity $\Theta(0)$ depends on the stochastic discretization scheme we choose. $\Theta(0) = 0, 1/2$ corresponding to the Itô and the Stratonovich scheme, respectively. As noted, we work with the discrete time version of the path, so we do not further consider this term.

C. Laplace's Approximation: Saddle Paths of the Action $A_0(\mathbf{x}(t))$

1. *General Results in Continuous Time*

In continuous time the expected values are written as [6, 7]

$$E[G(\mathbf{x}(t)|\mathbf{Y}] = \frac{\int \mathcal{D}x(t) G(\mathbf{x}(t)) \exp[-A_0(\mathbf{x}(t))]}{\int \mathcal{D}x(t) \exp[-A_0(\mathbf{x}(t))]}, \tag{11}$$

following Eq. (1). There are no restrictions on the variation of the endpoints in this expected value. This is seen directly in the discrete time formulation where we integrate $G(\mathbf{X})$ over all locations on the path including the initial point $\mathbf{x}(t_0)$ and the ending point $\mathbf{x}(t_m)$.

The expansion of the action Eq. (9) about saddle paths $\mathbf{x}^q(t)$; $q = 1, 2, \ldots$, satisfying

$$\frac{\delta A_0(\mathbf{x}(t))}{\delta \mathbf{x}(t)}|_{\mathbf{x}(t)=\mathbf{x}^q(t)} = 0, \tag{12}$$

at $\mathbf{x}(t) = \mathbf{x}^q(t)$ yields, writing $\delta\mathbf{x}(t) = \mathbf{x}(t) - \mathbf{x}^q(t)$,

$$A_0(\mathbf{x}(t)) = A_0(\mathbf{x}^q(t)) + \delta x_a(t)\ \frac{\partial \mathcal{L}(\mathbf{x}, \dot{\mathbf{x}}, t)}{\partial \dot{x}_a}\bigg|_{t_0}^{t_f}$$
$$+ \int_{t_0}^{t_f} dt\ \delta x_a(t)\ \left(\frac{\partial \mathcal{L}(\mathbf{x}, \dot{\mathbf{x}}, t)}{\partial x_a} - \frac{d}{dt}\frac{\partial \mathcal{L}(\mathbf{x}, \dot{\mathbf{x}}, t)}{\partial \dot{x}_a} \right)\bigg|_{\mathbf{X}(t)=\mathbf{X}^q(t)}$$
$$+ \frac{1}{2}\int_{t_0}^{t_f} dt\, \delta x_a(t) M_{ab}(\mathbf{x}^q(t), \dot{\mathbf{x}}^q(t), t, d/dt)\delta x_b(t) + \cdots . \tag{13}$$

The first variation of the action must be zero as a necessary condition for a possible minimum [18]. The integration by parts term leading to the Euler-Lagrange equations requires at the endpoints that

$$\delta x_a(t)\ \frac{\partial \mathcal{L}(\mathbf{x}^q(t), \dot{\mathbf{x}}^q(t), t)}{\partial \dot{x}_a^q(t)}\bigg|_{t_0}^{t_f} = 0, \tag{14}$$

which leads to, as $\delta\mathbf{x}(t)$ is not zero,

$$\frac{\partial \mathcal{L}(\mathbf{x}^q(t), \dot{\mathbf{x}}^q(t), t)}{\partial \dot{x}_a^q(t)}\bigg|_{t_0}^{t_f} = 0, \tag{15}$$

and are boundary conditions on the saddle path $\mathbf{x}^q(t)$. These are known as "natural boundary conditions" [18–21]. This is a necessary condition for the minimum path.

The second variation of the action contains the term in $(\delta\mathbf{x}(t), \delta\dot{\mathbf{x}}(t))$ space

$$\delta x_a(t)\frac{\partial^2 \mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t)}{\partial x_a(t)\, \partial x_b(t)}\delta x_b(t)$$
$$+2\delta x_a(t)\frac{\partial^2 \mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t)}{\partial x_a(t)\, \partial \dot{x}_b}(t)\delta \dot{x}_b(t)$$
$$+\delta \dot{x}_a(t)\frac{\partial^2 \mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t)}{\partial \dot{x}_a(t)\, \partial \dot{x}_b(t)}\delta \dot{x}_b(t). \tag{16}$$

This is familiar from many books on the calculus of variations [18–21]. Sufficient conditions for a minimum of the action are discussed in [19–21].

The Euler-Lagrange equations determining $\mathbf{x}^q(t)$ are

$$R_f(a)\left\{ \frac{d}{dt}[\dot{x}_a(t) - F_a(\mathbf{x}(t))] + [\dot{x}_b(t) - F_b(\mathbf{x}(t))]\frac{\partial F_b(\mathbf{x}(t))}{\partial x_a} \right\}$$
$$= R_m(l, t)[x_l(t) - y_l(t)]\delta_{al}. \tag{17}$$

Since the variations of the locations in state space $\delta\mathbf{x}(t_f)$ and $\delta\mathbf{x}(t_0)$ are unconstrained and independent when we are evaluating the expected value of a function on the path $G(\mathbf{X})$, the Euler-Lagrange equation is a second order differential equation in time with endpoint conditions

$$\frac{\partial \mathcal{L}(\mathbf{x}(t), \dot{\mathbf{x}}(t), t)}{\partial \dot{x}_a} = R_f(a)[\dot{x}_a(t) - F_a(\mathbf{x}(t))] = 0, \tag{18}$$

at $t = t_0, t_f$.

These Euler-Lagrange equations have been considered by Bröcker [22–24] as well as being considered over some years in the work of Bennett [25, 26]. We will discuss the relation of our work with the work by Bröcker below.

2. Boundary Conditions on the Euler-Lagrange Equations

In most treatments of variational problems in data assimilation, one addresses a slightly different set of boundary conditions for the Euler-Lagrange equations defining the saddle path. That is because a slightly different question is asked.

To see the difference we return to discrete time. We are concerned with the expected value of a function on the path which includes integration over both the final point in state space $\mathbf{x}(t_f)$ as well as the initial point in state space $\mathbf{x}(t_0)$. We can rewrite the numerator of the expected value of a function on the path $\mathbf{X} = \{\mathbf{X}', \mathbf{x}(0)\}$ as

$$\int d\mathbf{X}' G(\mathbf{X}', \mathbf{x}(0)) K(\mathbf{X}; t_f; \mathbf{x}(0), t_0)\, d^D x(0). \tag{19}$$

If we hold off in performing the integral over the initial state at the fixed time t_0, then the Laplace approximation which includes a variational principle for the quantity $K(\mathbf{X}; t_f; \mathbf{x}(0), t_0)$ has natural boundary conditions associated with the end point $\mathbf{x}(t_f)$ which is integrated over and not fixed, and it has the restriction that $\delta\mathbf{x}(t_0) = 0$ as $\mathbf{x}(0)$ is not integrated over in the variational principle.

This leads to the common statement that the relevant path has a fixed initial point and a free end point, changing the two point boundary value problem to be addressed to the one stated in the literature [26]. Each choice of boundary conditions is correct for the question posed.

D. Finding Saddle Paths

In continuous time, the saddle path condition is the two point boundary value problem we have described in Eq. (17) and Eq. (18). There are many discussions of how to numerically solve these [27–29], and in a sense we now use the collocation solution method discussed in these references by our return to the discrete time problem.

If we begin with the Gaussian error action

$$\begin{aligned} A_0(\mathbf{X}) = &\sum_{n=0}^{m} \sum_{l=1}^{L} \frac{R_m(l,n)}{2} [x_l(n) - y_l(n)]^2 \\ &+ \sum_{n=0}^{m-1} \sum_{a=1}^{D} \frac{R_f(a)}{2} [x_a(n+1) - f_a(\mathbf{x}(n))]^2, \end{aligned} \tag{20}$$

finding saddle paths, where $\partial A_0(\mathbf{X})/\partial \mathbf{X} = 0$, of this nonlinear function at sizable R_f entails a search in a high dimensional space in which the saddle paths $\mathbf{X}^q$ are located in narrow, possibly deep, valleys not easily found by an arbitrary selection of an initial path for a numerical optimization routine [8, 14].

We have investigated direct searches of saddle paths for this form of the action using a quasi-Newton BFGS method [30] and the public domain optimization program [31] called IPOPT. For the BFGS method, the analytical form of the first-order derivative for the action Eq. (20) is provided to optimization routines; When IPOPT is used, in addition to the analytical gradient, the Hessian matrix is also presented in analytical form to IPOPT, both of which are obtained by using a Python script we developed [32]. In each case the paths found via a direct search from a more or less arbitrary initial selection were not correct. The metric for 'correct' is whether the estimated parameters for the model and the full state of the model at the end of an estimation window, $\mathbf{x}(t_f)$, give good predictions for $t > t_f$. We will see some aspects of this as we begin to explore examples.

We have proposed [13] a strategy for dealing with this 'submersion' of the paths with smallest action. The idea is that when the model errors are forced, by large R_f, to be small, the nonlinearity of the vector field $\mathbf{f}(\mathbf{x})$ manifests itself at the smallest scales in the phase space of the paths $\mathbf{X}$ where we are searching. This causes complicated fine structure seen as multiple local minima [5] in the action, especially when the number of measurements L is too small.

As one reduces R_f, the loss of resolution of the flow of the nonlinear dynamics becomes large, and there is an averaging effect over the multiple local minima encountered. In fact at $R_f = 0$, there is no influence of the nonlinear dynamics of $\mathbf{x}(n+1) = \mathbf{f}(\mathbf{x}(n))$ at all, and the action we are working with then is

$$A_0(X) = \sum_{n=0}^{m} \sum_{l=1}^{L} \frac{R_m(l,n)}{2} [x_l(n) - y_l(n)]^2. \tag{21}$$

The minimum of this consists of $x_l(t) = y_l(t)$ for the observed states with the unobserved states completely unspecified. The huge degeneracy of this minimum action is broken as R_f increases from zero, and we propose to adiabatically track the $R_f = 0$ minimum, which is also the global minimum for $R_f = 0$, by slowly raising the magnitude of R_f from zero.

We call this tracking of extremum path in R_f "annealing" for short as it treats the importance of the deterministic dynamics in a steady, slowly growing manner, as if small R_f corresponds to a high effective temperature in which the nonlinear interaction among "particles" located at $\mathbf{x}(t)$ is initially in a harmonic well centered near the observations

$\mathbf{y}(t)$. As we increase R_f it is as if we cool down a 'temperature' proportional to R_f^{-1} and impose structure on the trajectories $\mathbf{x}(t)$. Note our annealing method is totally different from the well-known simulated annealing method [33].

Our 'annealing schedule' is written as

$$R_f = R_{f0}\alpha^{\beta} \tag{22}$$

with $\alpha > 1$ and $\beta \geq 0$. We proceed in the following manner

- Start at a small value of $R_f = R_{f0}$. We take as an `initial path` for our optimization algorithm the solution at $R_f = 0$ just described with $x_l(t) = y_l(t)$ and the other elements of the path drawn from a uniform random distribution covering the dynamical range of their variation. In practice we have selected R_{f0} to be between 0.001 and 0.01.

  We have also performed this initial stage of the annealing calculation starting at $R_f = R_{f0}$ with an initial set of N_0 random choices for the components of $\mathbf{X}$. Since the influence of the dynamics is so small at R_{f0}, the paths quickly become those suggested in the previous paragraph. No difference in the subsequent calculations have been seen by us.

- Using a selected optimization procedure, we then utilize N_0 choices for initial paths $\mathbf{X}_0^q;\ q = 1, 2, ..., N_0$ with $x_l(t) = y_l(t)$ and randomly chosen unobserved state variables. Fixed parameters in the model or in an observation function $\mathbf{h}(\mathbf{x})$ are also randomly selected over a finite range. The optimization procedure takes these initial paths and, with $R_f = R_{f0}$ (that is $\beta = 0$), results in N_0 new paths which we call $\mathbf{X}_1^q;\ q = 1, 2, ..., N_0$. We then evaluate the action $A_0(\mathbf{X}_1^q)$ on each of these N_0 paths. At this early stage we usually find the infinitely degenerate action values for the initial path $\mathbf{X}_0^q$ have begun to split.

- We proceed by raising the value of R_f from R_{f0} to $R_{f0}\alpha$, that is $\beta = 1$, and use the N_0 paths $\mathbf{X}_1^q$ as initial paths for this application of the optimization procedure. This results in a new set of N_0 paths $\mathbf{X}_2^q$.

- The paths $\mathbf{X}_2^q$ are now used as N_0 initial paths for $R_f = R_{f0}\alpha^2$, that is $\beta = 2$. This results in a set of N_0 paths resulting from our optimization procedure which we call $\mathbf{X}_3^q$.

- We continue this annealing schedule until by using the N_0 paths $\mathbf{X}_{J-1}^q$ to initialize the optimization procedure at $\beta = J - 1$ we arrive at a set of N_0 paths $\mathbf{X}_J^q$ where we terminate the schedule.

When the annealing schedule is completed, we often encounter in our plots of $A_0(\mathbf{X}^q)$ versus $\log_\alpha[R_f]$ a region for β large enough where some $A_0(\mathbf{X}^q)$ becomes independent of R_f. The lowest action level that splits off in action level value from the action on other paths, will, when the number of measurements L is large enough, provide the dominant contribution to $\langle G(\mathbf{X})\rangle$.

Independence of $A_0(\mathbf{X}^q)$ from R_f indicates that the model output has matched the deterministic dynamics $\mathbf{x}(n+1) = \mathbf{f}(x(n))$ quite well. The remaining term in the action is then

$$A_0(X) = \sum_{n=0}^{m}\ \sum_{l=1}^{L} \frac{R_m(l,n)}{2}[x_l(n) - y_l(n)]^2. \tag{23}$$

As the values $[y_l(n) - x_l(n)]$ are distributed as $\mathcal{N}(0, \sigma^2)$ by our choice, the measurement error term $\sum_{n=0}^{m}\sum_{l=1}^{L}[(x_l(n) - y_l(n))/\sigma]^2/2$ has a χ^2-distribution with $L(m+1)$ degrees of freedom [34]. The mean and RMS variation of this distribution over different choices of noise waveforms are $(m+1)L/2$ and $\sqrt{(m+1)L/2}$, respectively. This level is shown in our action value versus R_f plots by a heavy horizontal line. When the action levels as a function of R_f reach this expected χ^2 lower limit, we have a path $\mathbf{X}^0$ on which the model behavior is consistent with the data within the noise level of the data.

In the examples we will discuss below, we have selected $N_0 = 10 \sim 100$ and taken $20 \sim 50$ annealing steps as a stopping point. We have also used $\alpha = 2$ on the whole, but also selected $\alpha = 1.5$ when we chose to take smaller annealing steps in R_f within our schedule.

III. AN EXAMPLE ILLUSTRATING ANNEALING; LORENZ96 MODEL WITH $D = 5$

We begin by examining the dynamical equations introduced by Lorenz [35]:

$$\frac{dx_a(t)}{dt} = x_{a-1}(t)(x_{a+1}(t) - x_{a-2}(t)) - x_a(t) + \nu \tag{24}$$

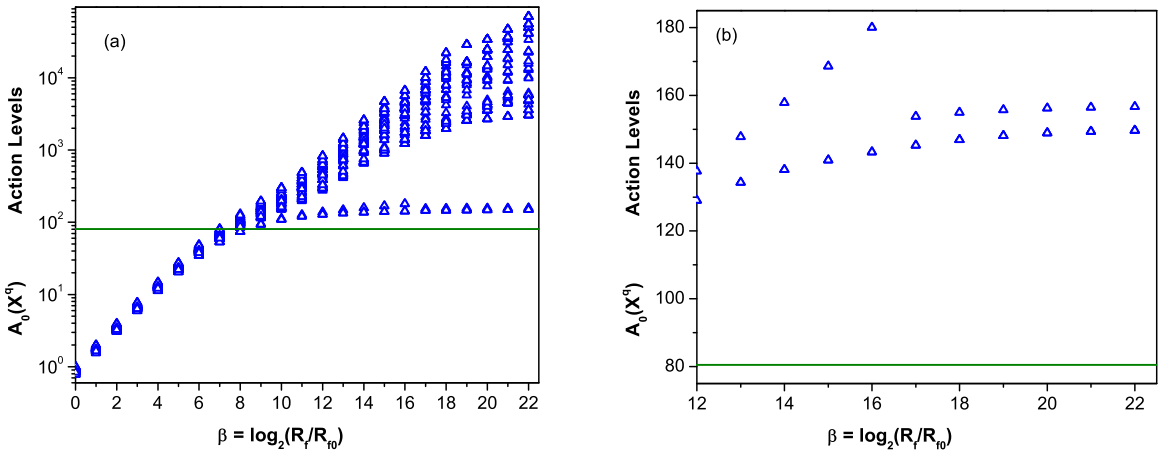

FIG. 1: (color online) **Panel (a)**: Action levels as a function of $\log_2[R_f/R_{f0}] = \beta$ for the Lorenz96 model, $D = 5$, $L = 1$, $R_{f0} = 0.01$. As R_f increases, the model error is decreased. The horizontal line shows the expected value of the measurement error terms in the action. The measurement error term is distributed as χ^2 with this expected value. **Panel (b)**: Zooming in on the Left Panel for large β showing close action levels. The resulting saddle paths all have action levels above the χ^2 expected value for the measurement error action alone. This is an indication that $L = 1$ is not sufficient for identifying a good path for the Lorenz96 $D = 5$ model.

and $a = 1, 2, ..., D$; $x_{-1}(t) = x_{D-1}(t)$; $x_0(t) = x_D(t)$; $x_{D+1}(t) = x_1(t)$. ν is a fixed parameter which we take to be in the range 8.0 to 8.2 where the solutions to the dynamical equations are chaotic [36]. The equations for the states $x_a(t)$; $a = 1, 2, ..., D$ are meant to describe 'stations' on a periodic spatial lattice.

We perform a twin experiment wherein we generate D time series using a standard adaptive fourth order Runge-Kutta algorithm with a time step $\Delta t = 0.025$ with no model error. To these we add Gaussian noise with mean zero and variance $\sigma^2 = 0.25$ to each time series $x_a(t)$. These noisy versions of our model time series constitute our 'data. Our choice of σ^2 means a signal to noise ratio about 60 dB. We selected $\nu = 8.17$ in these calculations.

The measurement window is from $t_0 = 0$ to $t_f = 4.0$, so $m = 160$. L 'measurements' are made at each time step; these are the $\mathbf{y}(t_n)$. The measurement error matrix $\mathbf{R}_m$ is taken to have diagonal elements at each measurement time t_n and is zero at other times. Its magnitude is taken as $R_m = 1/\sigma^2 = 4$. The model error matrix is also taken as diagonal, with elements along the diagonal $R_f = R_{f0}2^\beta$, and we take $\beta = 0, 1, 2, \ldots,$. R_{f0} was chosen 0.01 for these calculations.

We begin with one measurement $y_1(n)$ among the five possible states, i.e. $L = 1$, and in Fig. 1 we display the $\log_{10}$ of the action $A_0(\mathbf{X})$ evaluated at each of the $N_0 = 100$ saddle paths for $R_f = R_{f0}2^\beta$ with $R_{f0} = 0.01$ We begin with $\beta = 0$ and increase it to $\beta = 22$.

The BFGS quasi-Newton method [30] was used as our search algorithm in this example. The optimization stopping conditions for scaled gradient norm, scaled stepsize norm and function change are all 1×10^{-8} and they are fixed throughout the annealing procedure. Low tolerance (larger than 1×10^{-6}) may cause incorrect results or no convergence. We provided the gradient of $A_0(\mathbf{X})$ in an analytical form to the algorithm. We initialized the search at $R_f = R_{f0}$ with N_0 initial paths $\mathbf{X} = \{\mathbf{x}(0), \mathbf{x}(1), ..., \mathbf{x}(m)\}$ as described above. At $R_f = R_{f0}$ we selected the unobserved states at each time step from a uniform distribution in the interval [-10,10]. This is approximately the dynamic range of state variables in the Lorenz96 model. At fixed R_m each search procedure as we slowly increase R_f yields N_0 saddle paths $\mathbf{X}^q$ and associated action levels $A_0(\mathbf{X}^q)$. As R_f increases many initial paths may lead to the same action level.

As one can see in Fig. 1 the degenerate action levels at $R_f = 0$ are split at $\beta = 0$ and then rise until, around $\beta = 12; R_f \approx 100$, two levels split off from the rest and become rather independent of R_f. There are still two quite close levels. Also shown is the expected value of the χ^2-distributed measurement error at 80.5. The distance of the action levels of the paths giving $\mathbf{A}_0(\mathbf{X})$ near values about 150 tell us that these paths are unlikely to give consistency of the model with the data. Using either of these two paths to give us the full model states at the end of the estimation window $t_f = 4$ to predict beyond t_f gives quite inaccurate predictions.

Next we present $L = 2$ measurements, $y_1(n)$ and $y_3(n)$, to the model and again evaluate saddle paths as we vary β. Each path has $(m + 1)D = 805$ components, so the annealing problem is a search for saddle paths of the action Eq. (8) in an 805-dimensional space. We take the distribution of the three unobserved states at t_0, the beginning of

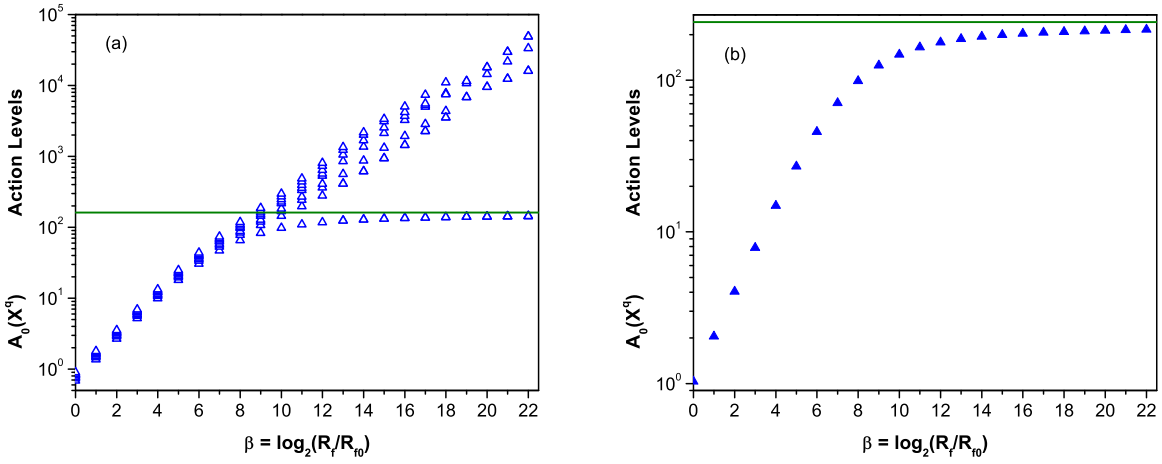

FIG. 2: (color online) **Panel (a)**: Action Levels as a function of R_f for the Lorenz96 model, D = 5, L = 2, $R_{f0} = 0.01$. We used $y_1(t)$ and $y_3(t)$ as data in the action. **Panel (b)**: Action Levels as a function of R_f for the Lorenz96 model, D = 5, L = 3, $R_{f0} = 0.01$. We used $y_1(t)$, $y_3(t)$, and $y_5(t)$ as data in the action.

the observation window, to be uniform over the dynamical range of $\mathbf{x}(t_0)$.

At $\beta = 0$ the degenerate action levels from $R_f = 0$ are split slightly. We follow these to larger values of R_f. (Fig. 2) At low R_f the resolution in path space is very coarse, and our search is successful for finding low lying action levels. Fig. 2 shows quite clearly that there are many paths with similar action level until we reach $\beta \approx 12$, and after that only one remains independent of R_f as the other action levels rise. The expected value of the χ^2-distribution of measurement errors in the action is $N_{\rm data} R_m \sigma^2/2$. This is 161 here, and it is shown in the figure as a heavy horizontal line. The action level for $\mathbf{X}^0$ is very near this χ^2 consistency condition suggesting that the path $\mathbf{X}^0$ expresses consistency of the model and the data.

When we increase the number of measured time series to $L = 3$, the results in Fig. 2 show that one path alone emerges from the degeneracy at $R_f = 0$ and after $\beta \approx 12$ is again nearly independent of R_f and close to the expected limit from the χ^2 distribution.

To get some insight into how the annealing procedure proceeds in the sequence of estimates for the observed and unobserved states of the model to which $L = 2$ 'data' time series are presented, we show in Fig. 3 the estimated and the 'data' time courses for both an observed state variable $x_1(t)$ and an unobserved state variable $x_2(t)$ for selected values of $\beta = 0, 12$ and 21.

In Fig. 3 representative time series which are part of the path for different values of β in the lowest action level are plotted to illustrate the annealing process in detail. For very small β, say 0, $R_f = 0.01$, the top two panels of Fig. 3 show the known and estimated components $x_1(t)$, observed, and $x_2(t)$, unobserved, from one of the saddle paths. Since $R_m \gg R_f$, and the measurement error dominates the overall size of the action, paths are forced to follow the noisy measurements almost exactly so as to minimize the measurement error, i.e. $x_l(t) \approx y_l(t)$. The effect of the model error term is quite small with $R_f = 0.01$, the unmeasured states are usually undetermined. Its form depends on the initial random guess path. In this example, the initial path happens to be chosen near the true path, and the unobserved state x_2 is close to the known data at the beginning and end of the window.

As β is increased to 12, $R_f \approx 40$, we have moved from a regime where R_f is quite small to a regime where R_f has become sizable. The role of the model error is no longer insignificant. The trajectory of the observed state $x_1(t)$ is smoother, passing through the middle of the noise fluctuations, but not tracking the noise as was done at $\beta = 0$. The greater R_f, the more information is input from the model. This information from the model helps the observed state filter out the noise to some extent. When β increases up to 21, $R_f \gg R_m$, it enforces the model more and more exactly, $x_a(n+1) \approx f_a(\mathbf{x}(n))$. Both observed states and unobserved states converge to the true path for the lowest action level. The size of action $A_0(\mathbf{X})$ matches the observation error residual $N_{data} R_m \sigma^2/2$.

It is important to note that if we begin our search for the saddle paths $\mathbf{X}^q$ at large values of R_f, we are almost sure to miss the actual path $\mathbf{X}^0$ which gives the lowest action level, since the Hessian matrix of $A_0(\mathbf{X})$ is ill-conditioned when R_f is large and the lowest action level occupies a tiny corner of the large (here 805 dimensional) path space. See Fig. 4.6 in [8].

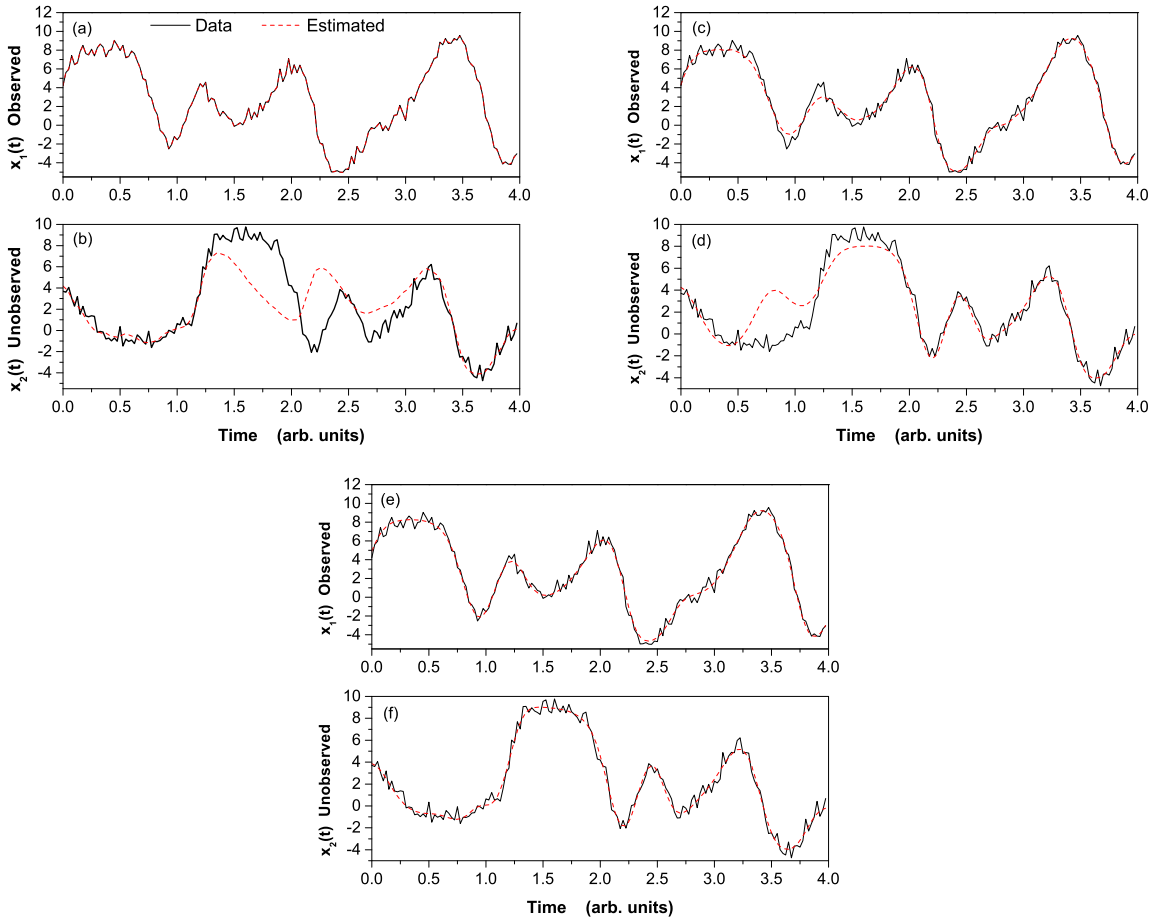

FIG. 3: (color online) Estimation results for one observed state variable $x_1(t)$ and one unobserved (not assimilated) state variable $x_2(t)$ during the annealing procedure. **Panels (a, b)**: $\beta = 0$. **Panels (c, d)**: As β increases to 12, the estimated observed state becomes smoother. **Panels (e, f)**: The observed and unobserved state variables arrive at their true states when β is large enough. Here $\beta = 21$.

IV. CORRECTIONS TO THE APPROXIMATION OF THE DOMINANT SADDLE PATH $\mathbf{X}^q$ TO $\langle G(\mathbf{X})\rangle$

The path integral formulation of $\langle G(\mathbf{X})\rangle$ allows more than just the leading variational approximation to this expected value, as is always the case in the Laplace method. The idea is to identify the path $\mathbf{X}^0$ with the smallest action level and then expand the integral in $\mathbf{X}$ about $\mathbf{X}^0$.

Near $\mathbf{X}^0$ we write

$$A_0(\mathbf{X}) = A_0(\mathbf{X}^0) + (\mathbf{X}-\mathbf{X}^0)_{\alpha_1}\gamma(\mathbf{X}^0)^2_{\alpha_1\alpha_2}(\mathbf{X}-\mathbf{X}^0)_{\alpha_2} + \cdots , \tag{25}$$

and the Hessian matrix $\gamma(\mathbf{X}^0)^2 = A_0^{(2)}(\mathbf{X}^0)/2$ is positive definite, if $\mathbf{X}^0$ gives a minimum of the action.

Changing integration variables to $U_\alpha = \gamma_{\alpha\beta}(\mathbf{X}^0)(\mathbf{X}-\mathbf{X}^0)_\beta$ leads to the numerator of $\langle G(\mathbf{X})\rangle$ in Eq. (1) arising from X^0

$$\int d\mathbf{X}\, \exp[-A_0(\mathbf{X})]G(\mathbf{X}) = \frac{\exp\left[-A_0(\mathbf{X}^0)\right]}{\det\gamma(\mathbf{X}^0)} \int d\mathbf{U}\, \exp\left(-\mathbf{U}^2 - V\right)\left[G(\mathbf{X}^0) + W\right] \tag{26}$$

where

$$V(\mathbf{U},\mathbf{X}^0) = \sum_{r=3} \frac{A^{(r)}(\mathbf{X}^0)_{\alpha_1\ldots\alpha_r}}{r!}[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]_{\alpha_1}\cdots[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]_{\alpha_r},$$
$$W(\mathbf{U},\mathbf{X}^0) = \sum_{k=1} \frac{G^{(k)}(\mathbf{X}^0)_{\alpha_1\ldots\alpha_k}}{k!}[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]_{\alpha_1}\cdots[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]_{\alpha_k}.$$

In the denominator we replace the numerator term $G(\mathbf{X}) = G(\mathbf{X}^0) + W(\mathbf{U},\mathbf{X}^0)$ by unity.

The terms in this integral are evaluated by expanding the Taylor series contributions $V(\mathbf{U},\mathbf{X}^0)$ and $W(\mathbf{U},\mathbf{X}^0)$ in powers of $\mathbf{U}$, and performing the resulting Gaussian integrals in $\mathbf{U}$. Terms with odd powers of $\mathbf{U}$ vanish by symmetry. The contributions to leading order in $1/R_f$ coming from the path X^0 for $E[G(\mathbf{X})|\mathbf{Y}]$ are

$$E[G(\mathbf{X})|\mathbf{Y}] = G(\mathbf{X}^0) + \int \frac{d\mathbf{U}\ \exp\left(-\mathbf{U}^2\right)}{\sqrt{\pi^{(m+1)D}}}\left\{\frac{1}{2}G^{(2)}(\mathbf{X}^0)[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]^2 - \frac{A^{(4)}(\mathbf{X}^0)}{24}G(\mathbf{X}^0)[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]^4 - \frac{A^{(3)}(\mathbf{X}^0)}{6}G^{(1)}(\mathbf{X}^0)[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]^4 + \frac{A^{(3)}(\mathbf{X}^0)^2}{72}G(\mathbf{X}^0)[\gamma(\mathbf{X}^0)^{-1}\mathbf{U}]^6\right\} + O\left(\frac{1}{R_f^2}\right) \tag{27}$$

From the form of the action in Eq. (20) we see that for large R_f each factor of $\gamma(\mathbf{X}^0)$ is of order $\sqrt{R_f}$, and each 'vertex' $A_0^{(4)}(\mathbf{X}^0)$ and $A_0^{(3)}(\mathbf{X}^0)$ is of order R_f. The Gaussian integrals over $\mathbf{U}$ are discussed in the opening chapters of Ref.[7].

Statistics such as the covariance about the saddle path $\mathbf{X}^0$ may be evaluated by selecting $G(\mathbf{X})$ to be the matrix in path $(a,n) = \alpha$ space as

$$G(\mathbf{X}) = (\mathbf{X}-\mathbf{X}^0)_\alpha\,(\mathbf{X}-\mathbf{X}^0)_\beta = U_\eta U_\kappa(\gamma(\mathbf{X}^0)^{-1})_{\kappa\,\alpha}\,(\gamma(\mathbf{X}^0)^{-1})_{\eta\,\beta}.$$

The Gaussian integral

$$\int d\mathbf{U}\ U_\alpha U_\beta\ \exp\left(-\mathbf{U}^2\right), \tag{28}$$

is easily performed, giving

$$\langle(\mathbf{X}-\mathbf{X}^0)_\alpha\,(\mathbf{X}-\mathbf{X}^0)_\beta\rangle = \frac{1}{2}(\gamma^2(\mathbf{X}^0)^{-1})_{\alpha\,\beta}, \tag{29}$$

which behaves as $1/R_f$ for large R_f, implying a steep, narrow minimum in path space.

These results are specific for the form of the action in Eq. (20), and may not apply for other choices for the distribution of noise in the measurements or the distribution of the errors in the models. Distributions with power law or 'fat' tails require additional scrutiny.

V. ENDPOINT CONDITIONS ON THE MINIMUM PATHS

In the integral for conditional expected values Eq. (1), when expressed in discrete time, each of the integrals for states $\int d^D\mathbf{x}(n)$ along the path $\mathbf{X} = \{\mathbf{x}(0),\mathbf{x}(1),...,\mathbf{x}(m)\}$ is unconstrained as we do not specify or hold fixed any of the state values in $\mathbf{X}$. Yet, when we proceed to the continuous time limit, as we have argued from the derivation of the Euler-Lagrange equations the canonical momentum $\partial\mathcal{L}(\mathbf{x},\dot{\mathbf{x}},t)/\partial\dot{\mathbf{x}}(t)$ must vanish at the temporal end points of the integration.

We investigated this by examining the saddle paths for $\beta = 13$ using the Lorenz96 model with $D = 5$ and $L = 2$. In Fig. 4 we show the canonical momentum $p_a(t) = R_f(a)[\dot{\mathbf{x}}_a(t) - F_a(\mathbf{x}(t))]$ for $a = 1,...,5$ scaled by the maximum value within the observation window.

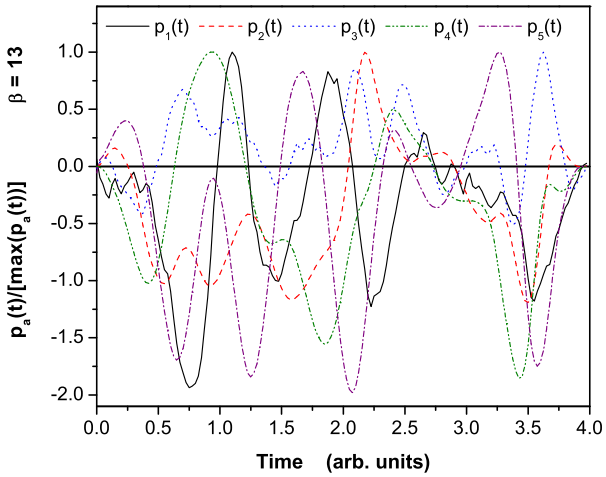

FIG. 4: (color online) Five canonical momenta $p_a(t) = R_f[\dot{x}_a(t) - F_a(\mathbf{x}(t))]$ of one local minimum path for the Lorenz96 model, $D = 5$, $L = 2$. $\beta = 13$. $R_f(a) = R_f$ for all a = 1, ...,5. At the endpoints $t = 0$ and $t = 4$, all the $p_a(t)$ go to zero, satisfying the boundary conditions associated with the vanishing of the first variation of the action. Each $p_a(t)$ is scaled by its maximum magnitude over $0 \leq t \leq 4$

We randomly picked one of the saddle paths at $\beta = 13$, and we evaluated the $\dot{\mathbf{x}}(t)$ in the canonical momentum using the second-order central finite difference scheme

$$\dot{\mathbf{x}}(n) = \frac{\mathbf{x}(n+1) - \mathbf{x}(n-1)}{2\Delta t} + O(\Delta t^2),$$

for interior points in the observation window. We maintained second-order accuracy for the evaluation of $\dot{\mathbf{x}}(t)$ at the boundary by using

$$\dot{\mathbf{x}}(0) = \frac{-3\mathbf{x}(0) + 4\mathbf{x}(1) - \mathbf{x}(2)}{2\Delta t} + O(\Delta t^2)$$
$$\dot{\mathbf{x}}(m) = \frac{3\mathbf{x}(m) - 4\mathbf{x}(m-1) + \mathbf{x}(m-2)}{2\Delta t} + O(\Delta t^2).$$

The scaled trajectories in $t_0 \leq t \leq t_f$ of the five components of canonical momentum $p_a(t)$ are plotted in Fig. 4, which shows the required boundary condition is satisfied rather well. As this boundary condition is necessary for an extremum of the action when the end points are not constrained, this result may be only a consistency check on the accuracy of our calculation. Of course, it is good that the result is positive.

VI. FURTHER EXAMPLES

To further explore the annealing method we examine an example from the atmospheric sciences literature [35] and an example involving a standard Hodgkin-Huxley neuron model. The first has both fast and slow variables, namely fast and slow time scales in the governing differential equations, representing small scale (fast) and large scale (slow) atmospheric variations. The challenge to a variational method is to reliably capture both time scales in identifying an accurate saddle path. The second moves away from the simple vector fields in the Lorenz96 model, and fluid dynamical models in a general sense, which are at most quadratic in their nonlinearities. The kinetics of gating variables associated with voltage dependent conductances in neurons involves parameters and states entering the vector fields through exponentials reflecting the underlying statistical properties of the cellular processes.

A. Lorenz96 model with both fast and slow variables

Dynamical systems varying with several distinct time scales are commonly seen in earth system models. In the same paper where Lorenz introduced the Lorenz96 model we reported on earlier, he also introduced a modified model

with both fast and slow variables to study the local instability responsible for convective activity. [35] This model is given by

$$\begin{aligned}\frac{dx_k(t)}{dt} =& - x_{k-1}(t)(x_{k-2}(t) - x_{k+1}(t)) - x_k(t) + \nu \\ & - \frac{hc}{b} \sum_{j=J(k-1)+1}^{kJ} z_j(t); \\ \frac{dz_j(t)}{dt} =& - cbz_{j+1}(t)(z_{j+2}(t) - z_{j-1}(t)) - cz_j(t) \\ & + \frac{hc}{b} x_{\text{floor}[(j-1)/J]+1}(t). \end{aligned} \tag{30}$$

where $k = 1, 2, ..., D_{\text{slow}} = K$; $j = 1, 2, ..., D_{\text{fast}} = JK$ and floor$[x]$ is the floor function. $x_0(t) = x_K(t), x_{-1}(t) = x_{K-1}(t)$, etc. and likewise for the $\mathbf{z}(t)$ variables. The first equation describes the linked dynamics of a set of K slow, large-amplitude variables $x_k(t)$, each of which is associated with J fast, small-amplitude variables $z_j(t)$ whose dynamics are described by the second equation. $z_j(t)$ represents a convective-scale quantity coupling with $x_k(t)$ that favors the convective activity. They can be visualized as sectors on a lattice circle (Figure 1 in Ref.[37]). Each $x_k(t)$ sector contains many $z_j(t)$ sectors. We chose $D_{\text{slow}} = K = 5$ and $D_{\text{fast}} = J = 5$; this means that five $z_j(t)$ sectors are contained in one $x_k(t)$ sector. The detailed time scales for x_k and z_j are determined by the parameters h, c, b. We follow Lorenz and select $h = 1, c = 10, b = 10$, so the fast variables $z_j(t)$ vary approximately 10 times more rapidly than the slow variables $x_k(t)$, while their amplitudes are about one tenth of those of the $x_k(t)$. The forcing parameter ν is taken to be 18 following Wilks' work to make both $x_k(t)$ and $z_j(t)$ chaotic [37].

To generate our data, we integrate the slow equations with a time step $\Delta t = 0.001$ for the temporal window $t = [0, 4] = 4000\Delta t$ using a standard adaptive fourth order Runge-Kutta scheme. From these time series 'data' are obtained by adding white Gaussian measurement noise to the computed time series: $\mathcal{N}(0, 0.5)$ for $x_k(t)$ and $\mathcal{N}(0, 0.05)$ for $z_j(t)$.

Fixing the parameters h, c, b and f at the values used to generate the data, we seek to estimate the unobserved state variables as $L < K(J+1)$ time series are presented to the model. As above we perform our calculation using the annealing method. Following our definition, R_m equals the inverse of the variance of measurement noise, i.e. $R_m = 4$ for $x_k(t)$ and $R_m = 400$ for $z_j(t)$ when the corresponding variables are measured, and 0 otherwise.

Instead of choosing the same R_f value for every variable as we did in our illustrative example Lorenz96 model, the values of R_f need to be determined by how rapidly the variables changes. The R_f works as a penalty parameter during the annealing process. A more rapidly varying variable requires a larger R_f to regulate its fluctuations, and vice versa. By manipulating the values of R_f, we can ensure the model error terms are of the similar scale. Therefore, during the variational process, all the elements of each path from different variables will be well adjusted simultaneously to the path.

As we stated above, the amplitude of the $dz_j(t)/dt$ is approximately ten times smaller than that of the $dx_k(t)/dt$, so the variance of model error for $x_k(t)$ is about 100 times larger. The R_f's, as the inverse of model error variance, are chosen to be

$$R_{f0} = \begin{cases} 0.01 & \text{for } x_k(t) \\ 1 & \text{for } z_j(t). \end{cases}$$

According to our results above in the Lorenz96 $D = 5$ model the action level plots suggest that when measurements of $y_1(t), y_3(t)$ are presented, all the unmeasured states can be accurately estimated and lead to excellent prediction. We proceed then by presenting the noisy $x_1(t)$ and $x_3(t)$, namely, $y_1(t), y_3(t)$ to the fast/slow model, and then gradually increase the number of noisy measurements of $z_j(t)$ until we can find the consistent lowest action level we expect. We call the total number of measurements L and recognize it is comprised of 2 observations of the $x_k(t)$ and $L - 2$ observations of the $z_j(t)$.

The action level plots are shown in Fig. 5. We start the annealing calculation with $N_0 = 100$ initial random paths, so that whether there are paths located at the expected lowest action level is a stochastic event. When $L \leq 13$, it is rare to find a path finally reaching the expected lowest action level. The same calculations for $L = 12, 13$ are repeated 10 times, and no action level near the overall expected values is observed. Fig. 5 shows one example of $L = 13$ such that no action level levels out with increasing R_f. This suggests none of the paths is consistent with the model and that model errors are amplified as R_f increases. It is a sign that indicates the information provided by the selected measurements is not enough. Another sign denoting lack of information in the measurements is that there are many close action levels near the expected lowest action level as shown in Figure 1 of Ref.[13].

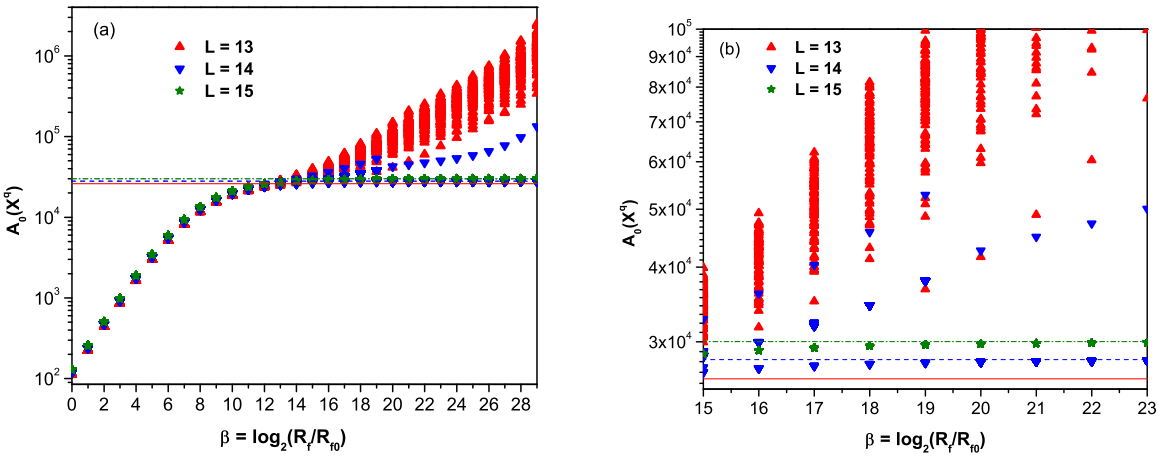

FIG. 5: (color online) **Panel (a)** Action Levels as a function of R_f for the Lorenz96 Fast/Slow model, $K = 5$, $J = 5$: when $L = 13$ we used noisy measurements of $x_1(t), x_3(t)$ and $z_{2j-1}(t)$'s with $j = 1, 2, \ldots, 11$ as measured variables. At $L = 14$, $z_{23}(t)$ is added; at $L = 15$, $z_{25}(t)$ is also added. **Panel (b)** Zooming in on the action levels in the Left Panel for large β. Note that when $L = 14$ and $L = 15$, the lowest action level splits off from the other allowed action levels corresponding to other paths meeting the saddle path condition.

It is worth pointing out that not having enough measurement information is not the only possible cause of having no levels becoming independent of R_f for large R_f. Anything introducing significant non-zero model errors can cause this as well. For example, numerical discretization error in action may appear to be important in fast/slow dynamical systems when one doesn't use a small enough observation time step to meet the requirement of the resolution of their fastest dynamics. In the Lorenz96 fast/slow model $K = 5$, $J = 5$ if we use $\Delta t > 0.0025$, one can observe this even when $L = 14$.

We added another measurement to the action, and the complicated action levels at $L = 13$ are immediately reduced to three distinct levels when $12 \leq \beta \leq 19$ and then collapses into two levels, including one associated with the expected value, after $\beta \geq 20$. Adding in $z_{25}(t)$, we found all of the 100 random initial paths converge to the solution of the consistent action minimum.

The estimated state variables of the Lorenz96 fast/slow model at $L = 14$ are displayed in Fig. 6. Predictions are obtained by advancing the model forward in time using the estimated state variables at $t = 4.0$ as an initial condition. (Fig. 7) Both estimation and prediction show excellent agreement with the data generated in this twin experiment. In time the prediction loses its accuracy because of the chaotic nature of solutions to this dynamical system. This is especially apparent for the fast variables $z_j(t)$.

The annealing method is successful in locating a dominant lowest action path even in the presence of distinct times scales in the dynamics. One must sample the observations fast enough to capture the higher frequency variations of the fast variables. In the annealing search for lowest action levels, one must also select the ratios of the maximal values of R_f to reflect the different dynamical time scales.

B. Standard Hodgkin-Huxley neuron model: NaKL model

We selected a fairly standard Hodgkin-Huxley (HH) neuron model [38, 39] consisting of four state variables: the voltage $V(t)$ across the cell membrane as well as three voltage dependent gating variables for Na^+ and K^+ channels $m(t), h(t)$ and $n(t)$. The equation governing changes in voltage across the cell membrane is current conservation with conductances for Na^+ and K^+ ions through the membrane that depend on the voltage $V(t)$. This reflects the change in permeability to these ions of proteins that transect the membrane and change their conformation as a function of the voltage across the membrane. The specific forms of the voltage dependent conductivities in this HH model are taken from textbook descriptions based on the 1940s and 1950s work of Hodgkin, Huxley, Katz, and many others. The reversal potentials are determined by the competition of diffusion associated with ion concentration differences within and without the cell and transport of charged ions by the electric field associated with the difference in voltage across the membrane. The Nernst equation which determines these reversal potentials is directly from statistical physics.

The cell responds to external currents as a driving force by its cross membrane voltage $V(t)$ rising if the current

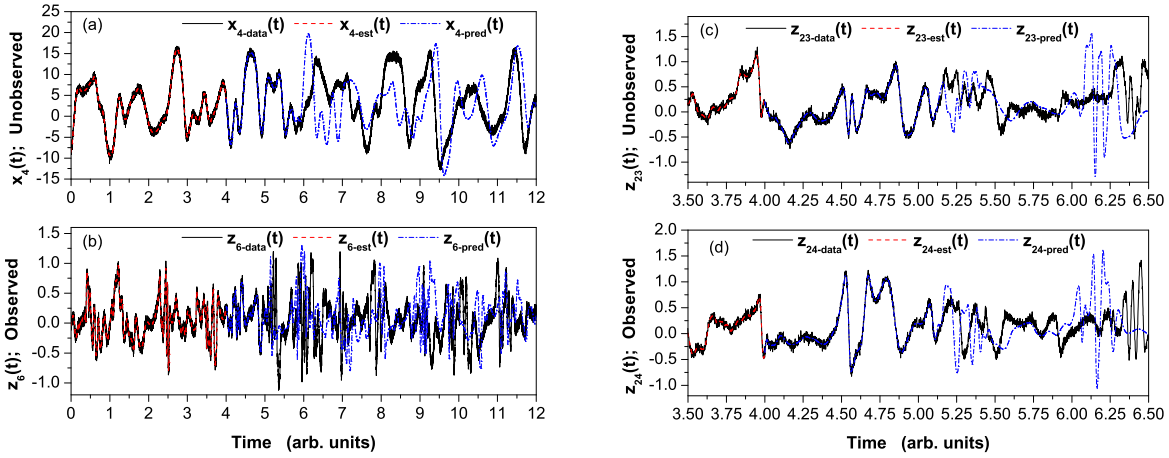

FIG. 6: (color online) Data, estimated and predicted time series for the Lorenz96 Fast/Slow model with $K = 5$, $J = 5$ and $L = 14$. f = 18. We used noisy versions of the $x_1(t)$, $x_3(t)$ and of the $z_{2j-1}(t), j = 1, 2, \ldots, 12$ as measured variables. In the **Panel (a)** $x_4(t)$ is an `unobserved` slow variable. In the **Panel (c)** $z_{23}(t)$ is an `unobserved` fast variable. In the **Panel (b)** and **(d)** both of the fast variables $z_6(t)$ and $z_{24}(t)$ are `observed`.

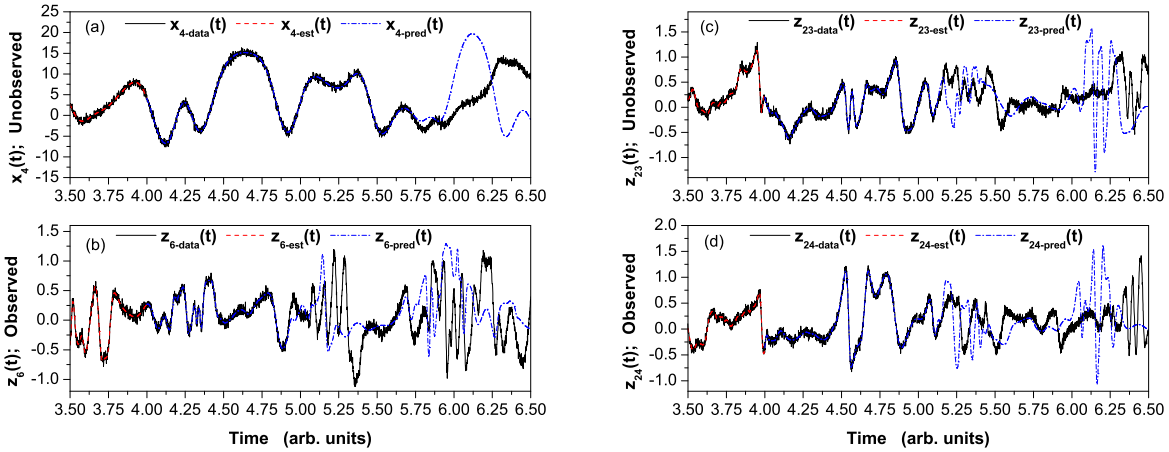

FIG. 7: (color online) Zoom in on the data, estimation, and prediction of the $x_4(t), z_6(t), z_{23}(t)$ and $z_{24}(t)$ variables in the Lorenz96 fast/slow model.

causes depolarization of the cell, or the voltage decreases when the cell becomes more polarized. (Fig. 8) The rise in voltage triggers an instability in the phase space of the HH model associated with a sudden influx of Na^+ ions which is then counterbalanced by a flow of K^+ ions out of the cells as the voltage rises to order + 50 mV. All this takes place on the order of 5 ms and is seen as a 'spike' in the voltage time series.

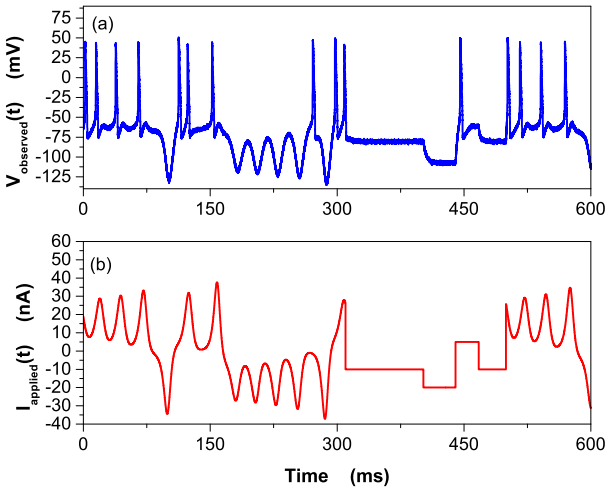

FIG. 8: (color online) **Panel (a)** Voltage response of the standard Hodgkin-Huxley neuron model with Na^+, K^+ and leak channels, our NaKL model, in response to the applied (injected) stimulus current shown in the **Panel (b)**

The model is governed by the following four first-order differential equations:

$$
\begin{aligned}
C\frac{dV(t)}{dt} =& I_{inj}(t) + g_{Na}m(t)^3h(t)(E_{Na} - V(t)) \\
&+ g_K n(t)^4(E_K - V(t)) + g_L(E_L - V(t)) \\
\frac{da(t)}{dt} =& \frac{a_0(V(t)) - a(t)}{\tau_a(V(t))} \qquad a(t) = \{m(t), h(t), n(t)\} \\
a_0(V) =& \frac{1}{2} + \frac{1}{2}\tanh\left(\frac{V - V_a}{\Delta V_a}\right) \\
\tau_a(V) =& \tau_{a0} + \tau_{a1}\left(1 - \tanh^2\left(\frac{V - V_a}{\Delta V_a}\right)\right)
\end{aligned}
$$

In these equations the $g_{\rm ion}$'s are maximum conductances for the ion channels, the $E_{\rm ion}$ are reversal potentials for those ion channels, $I_{\rm inj}(t)$ is the external stimulating current injected into the neuron. This current is selected by the experimenter and has no independent dynamics.

The gating variables $a(t)$ are taken to satisfy first order kinetic equations and range between zero and unity. The overall strength of an ion channel is set by the maximal conductances, and this represents the number of individual ion channels. These are phenomenological choices.

The quantities $a_0(V)$ and $\tau_a(V)$ are the voltage dependent activation function and the voltage dependent time constant of the gating variable $a(t)$. The forcing to the cell $I_{\rm inj}(t)$ is known to us. In our parametrization of the cell dynamics there are 19 fixed parameters and three unobserved state variables $a(t) = \{m(t), h(t), n(t)\}$ to be determined. All $a(t)$ lie between zero and one.

Only the voltage across the cell membrane is measurable in real neurobiological experiments, however, successful data assimilation, in effect, 'measures' the gating variable time series as well as the unknown parameters. We present only noisy time series of $V(t)$ to the model; these are our $\mathbf{y}(t)$ in the notation we have used for the general discussion above.

The parameters used to generate data are listed in Table I. The waveform of the injected current is chosen to be a combination of step functions and segments of a chaotic time series taken from one of the variables of the Lorenz63 model [40]. This current is displayed in the bottom panel of Fig. 8. A standard adaptive fourth order Runge-Kutta solver is used to produce the data using time steps of $\Delta t = 0.025$ms, and white Gaussian noise with an RMS level of 1mV is added to the $V(t)$ time series to represent the noise accompanying the measurements in laboratory biological experiments. This voltage time course is in the top panel of Fig. 8.

The previous two examples, the Lorenz96 $D = 5$ model and the Lorenz96 fast/slow model contain only quadratic nonlinear terms in their differential equations. The difficulties of state and parameter estimation result from their chaotic trajectories. The NaKL model in the selected parameter region is not chaotic, and the challenge of data

TABLE I: Known and estimated parameters for the NaKL model. We also display the bounds used for the nonlinear search algorithm.

Parameters	Known	Estimated	Search Lower Bound	Search Upper Bound
g_{Na}	120.0	108.4	50.0	200.0
E_{Na}	50.0	49.98	0.0	100.0
g_K	20.0	21.11	5.0	40.0
E_K	-77.0	-77.09	-100.0	-50.0
g_L	0.3	0.3028	0.1	1.0
E_L	-54.0	-54.05	-60.0	-50.0
C	0.8	0.81	0.5	1.5
V_m	-40.0	-40.24	-60.0	-30.0
ΔV_m	0.0667	0.0669	0.01	0.1
τ_{m0}	0.1	0.0949	0.05	0.25
τ_{m1}	0.4	0.4120	0.1	1.0
V_h	-60.0	-59.43	-70.0	-40.0
ΔV_h	-0.0667	-0.0702	-0.1	-0.01
τ_{h0}	1.0	1.0321	0.1	5.0
τ_{h1}	7.0	7.76	1.0	15.0
V_n	-55.0	-54.52	-70.0	-40.0
ΔV_n	0.0333	0.0328	0.01	0.1
τ_{n0}	1.0	1.06	0.1	5.0
τ_{n1}	5.0	4.97	2.0	12.0

assimilation comes from the richer nonlinearity in the dynamics of the gating variables and the sensitivity of the model behavior to changes in parameter values.

In the numerical optimization used to find the saddle paths for any model one must specify search bounds for each parameter and each state variable. The goal is to find appropriate choices for these values that constrain the model states and parameters to biologically acceptable regions. The bounds for the voltage $V(t)$ are taken as -150mV and $+70$mV based upon our solutions to the equations. The gating variables are bounded between 0 and 1, since they represent the probability whether ion channels are open or closed.

The optimization is implemented with IPOPT using an interior-point method [31]. We found the interior-point method both more stable and substantially faster than the L-BFGS-B method [41].

In this twin experiment only a noisy voltage $V(t)$ is 'measured' and presented to the model; so $L = 1$. As the dynamical range of voltage is a hundred times larger than that of the gating variables, we first calculated the action levels with $R_m^{(V)} = 1, R_{f0}^{(V)} = 10^{-3}, R_{f0}^{(m)} = 10, R_{f0}^{(h)} = 10, R_{f0}^{(n)} = 10$ and $\alpha = 3/2$. The largest β was taken as 50. Also we decreased α from 2 to 3/2 so that the pace of increasing resolution in model state space is slower than in our earlier examples as we increment changes in R_f. This allows us to stay well within the basin of attraction of the lowest action level.

The top panel of Fig. 10 displays the action level plot with the configuration above. When $25 \leq \beta \leq 39$, there are several different levels, and they reveal the expected action level after $\beta \geq 40$. The action level plot suggests that the voltage measurement alone is sufficient to determine the three unobserved state variables as well as the 19 parameters. The estimates of these parameters are displayed in Table I.

We can select the values of R_{f0}'s for each state variable according to our knowledge about the amplitudes and time scales of the state variables by looking at the time series of solutions of the model. The time constants of gating variables characterize the their response to the change of voltage. The sodium activation variable $m(t)$ is the fastest, with a time constant of several hundreds of microseconds, which is a little bit slower than $V(t)$. $h(t)$ and $n(t)$ are much slower having a time constant of a few milliseconds. We set the ratio of $R_{f0}^{(m)}/R_{f0}^{(V)} = 5 \times 10^4$ and raised the ratio of $R_{f0}^{(n)}/R_{f0}^{(m)}$ and $R_{f0}^{(h)}/R_{f0}^{(m)}$ from 1 to 10 to compensate for the effects induced by different time constants. With $R_m = 1, R_{f0}^{(V)} = 10^{-3}, R_{f0}^{(m)} = 50, R_{f0}^{(h)} = 500, R_{f0}^{(n)} = 500$, the action level plot in the bottom panel of Fig. 10 shows this configuration of R_{f0} can effectively enforce that most saddle paths stay near the expected lowest action level.

The detailed action levels can depend on the choice of R_{f0}. In the NaKL example when the ratios of $R_{f0}^{(n)}/R_{f0}^{(m)}$

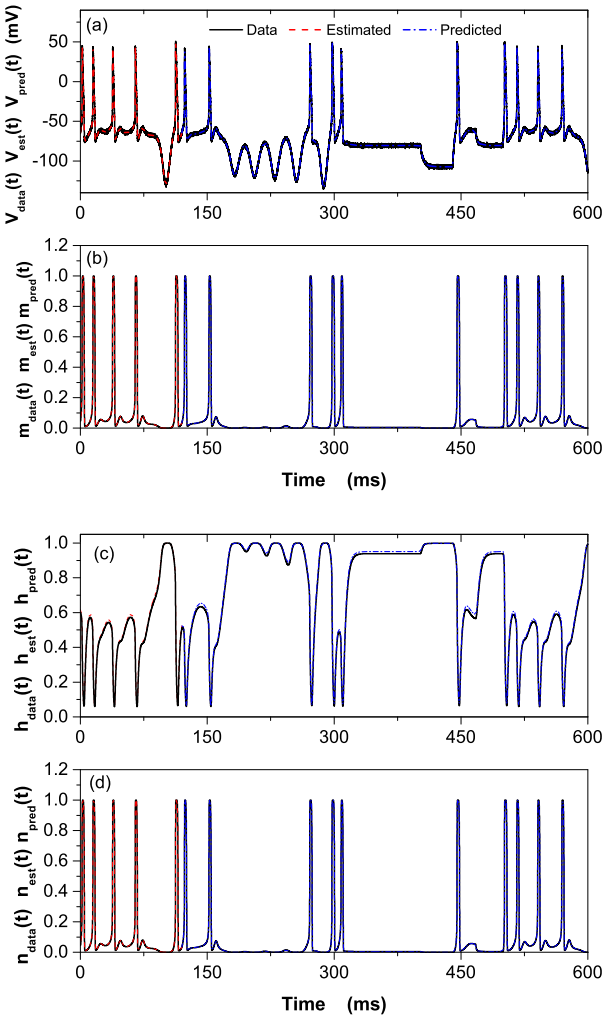

FIG. 9: (color online) **Panels (a, b)**: Data (black), estimated (red) and predicted (blue) state variables $V(t), m(t)$ for the NaKL model when only the noisy membrane voltage $V(t)$ is measured and presented to the model. **Panels (c,d)**: Data (black), estimated (red) and predicted (blue) state variables $h(t), n(t)$ when only the noisy membrane voltage $V(t)$ is measured and presented to the model

and $R_{f0}^{(h)}/R_{f0}^{(m)}$ are as large as 100, we often observed another action level with a value close to the lowest one. This can also depend on the specific choice of the N_0 initial paths with which we start the annealing.

A lesson we learn from this example is that we should take both the amplitude and the time scale of the state variables into consideration when selecting the scale of R_f values. The proper configurations of R_f values enlarge the probability to have the candidate paths converge to the expected action level, and also accelerate the convergence rate to the optimal paths.

VII. DISCUSSION

A. Connection with Bröcker's Results

The annealing method has a close relationship, and in some places a significant overlap with two very nice papers [22, 23] and a quite valuable and pedagogical Summer School presentation by J. Bröcker [24]. He considers the formulation of the assimilation problem in continuous time and identifies an action which corresponds to our $A_0(\mathbf{X})$.

Missing in his discussions is the context of the path integral to be used to evaluate expected values of functions on the path, so no method is presented to evaluate the accuracy of the variational result as a function, say, of R_f.

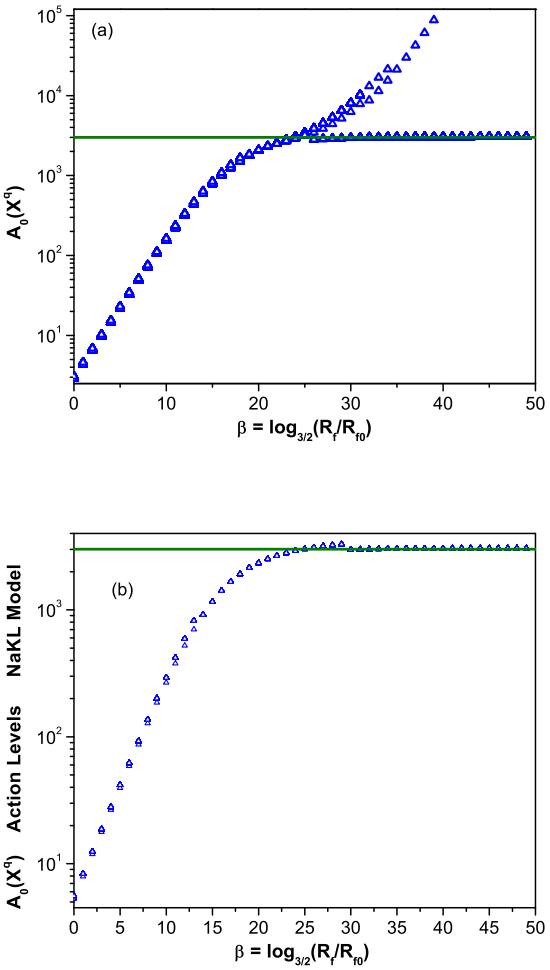

FIG. 10: (color online) Action levels as a function of R_f for the NaKL model with only the noisy membrane voltage $V(t)$ measured and presented to the model. **Panel (a)**: we selected $R_{f0}^{(V)} = 10^{-3}, R_{f0}^{(m)} = 10, R_{f0}^{(h)} = 10, R_{f0}^{(n)} = 10$ and $\alpha = 3/2$.; **Panel (b)**: $R_m = 1, R_{f0}^{(V)} = 10^{-3}, R_{f0}^{(m)} = 10, R_{f0}^{(h)} = 1000, R_{f0}^{(n)} = 1000$, and $\alpha = 3/2$

He identifies the boundary conditions on the Euler-Lagrange equations for the extrema of the action, and recognizes with great care how one selects among the possibilities.

His formulation of the actual action differs from ours in that he adds to the deterministic equations $\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t))$ a control term which is employed in moving the output of a model toward the observations, then he imposes this control via a Lagrange multiplier. The Euler-Lagrange equations in Eq. (11) of Ref.[24] are, happily, in his notation the same as our Eq. (17). In Section 3, especially Exercise 3.1 of Bröcker's Summer School notes [24] he shows that his approach and our yield the same extremum conditions on the action.

Importantly he recognizes and explores with insight the manner in which the model error term in the action and the measurement error term in the action 'balance' to direct the solution to a chaotic model equation to those regions of phase space where the observations provide information about the model solutions.

Bröcker does not consider the question of how many measurements are required to allow the search for extrema of the action to be achievable [42] or whether any of the allowed extrema are, in fact, the overall minimum of the action. He does stress the importance of the model error term in the action and formulates its appearance in a very clear and useful fashion.

Our use of an annealing method to address these latter two questions has roots in our own work [5] and has connections with Bröcker's work. The route we follow now is to return to discrete time formulations of the action with a focus on the questions one wishes to answer in the use of the data assimilation methods in physical and biological problems: (1) What is the expected value of the state and parameters of the model system at the end of an

observation window–this includes unmeasured as well as measured states. (2) What are the RMS errors about this expected state ? (3) What is the accuracy, including RMS errors, of the predictions of the model, conditioned on information transferred to it by the data, for times after the observation window.

B. Comparison with the quasi-static variational assimilation (QSVA) Method

There is another strategy for determining the best path for a selected model with chaotic trajectories given observed data known as the quasi-static variational assimilation (QSVA) method [14]. In this approach the known equations of motion are initialized with some $\mathbf{x}^{(0)}(t=0)$ and integrated forward a small step in time of length τ. A cost function comparing the known observations with the model output over that time step τ is minimized to adjust the initial condition $\mathbf{x}^{(0)}(t=0) \to \mathbf{x}^{(1)}(t=0)$. Then the time interval is extended to 2τ and starting with $\mathbf{x}^{(1)}(t=0)$ the equations of motion are integrated forward to 2τ. The difference between the orbit from the initial condition $\mathbf{x}^{(1)}(t=0)$ is compared via the cost function to data in the longer interval 2τ and the cost function is minimized taking $\mathbf{x}^{(1)}(t=0) \to \mathbf{x}^{(2)}(t=0)$. This repeated adding increments of time to the observation window until it is $N\tau$ long, and a final $\mathbf{x}^{(N-1)}(t=0)$ is found via minimizing the cost function over the interval up to $N\tau$. Starting with this final selection of initial condition $\mathbf{x}^{(N-1)}(t=0)$ a trajectory to time $N\tau$ and beyond is generated using the equations of motion. This gives a path through the estimation window $[0, N\tau]$ and beyond for prediction. By choosing the original initial condition $\mathbf{x}^{(0)}(t=0)$ within a resolution ball of uncertainty and drawing N_0 different selections, one can generate N_0 different paths during both the estimation and prediction windows. In the paper [14] the Lorenz 1963 model was used with parameters set to produce chaotic orbits, and it was assumed all the state variables were observed.

We have made a direct comparison of the QSVA approach with our annealing method. We used the Lorenz96 model with $D = 5$ and $L = 3$ to generate 'data', and then we added noise of the same level with variance $\sigma^2 = 1/4$ to the observations. We observed L = 3, $y_1(t), y_3(t)$, and $y_5(t)$ and in choosing $\mathbf{x}^{(0)}(t=0)$ for QSVA we selected the unobserved components $y_2(0), y_4(0)$ from a uniform distribution in the interval $[-10, 10]$. We then evaluated (1) $N_0 = 100$ initial conditions arriving at 100 final selections of initial conditions using time steps of $\tau = 4\Delta t = 0.1$ where $\Delta t = 0.025$ and (2) a second numerical trial with $\tau = 4\Delta t = 0.004$ where $\Delta t = 0.001$ with $N_0 = 60$ initial conditions. We generated the data with a fourth-order Runge-Kutta integrator for the $D = 5$ Lorenz96 equations with forcing $f = 8.17$. The minimization of the cost function

$$C(\mathbf{x}(0), N) = \sum_{n=0}^{N} \sum_{l=1}^{L} (x_l(n) - y_l(n))^2 \tag{31}$$

at each step to find an improved $\mathbf{x}^{(K)}(0)$ was performed using a quasi-Newton BFGS method [30]. The estimation window was taken as 160 $\Delta t = 4$, and the prediction window ranged from 4 to 8.

We compared this calculation of initial conditions and estimated and predicted orbits with the annealing method described in Sec.III using annealing up to $\beta = 30$ as shown in the right panel of Fig. 2.

To compare the outcome of the two approaches we display the following results for the first set of calculations with $\Delta t = 0.025$ and $N_0 = 100$:

- In the autonomous Lorenz96 model $D = 5$ there is one unstable direction with positive Lyapunov exponent $\lambda = 0.53$ and one neutral direction as it is a differential equation. We projected the $N_0 = 100$ solutions from QSVA into the plane of these two directions along with the location in the same plane of the first state $\mathbf{x}(0)$ from the annealing process. These are shown in Fig. 11.

- The outcome of the annealing estimations gives the whole path in the time interval [0,4], including $\mathbf{x}(t=0)$ and $\mathbf{x}(t=4)$ at the start and end of the interval, respectively. The predictions of annealing method are obtained in two ways: one integrates dynamical model Eq. (24) forward from $\mathbf{x}(t=0)$ and the other one starts from the end of the estimation window $\mathbf{x}(t=4)$ directly which can largely reduce the numerical errors introduced by chaos. (Fig. 13) However, QSVA only allows the first option.

- We evaluated the RMS error for both the annealing and QSVA calculations by comparing the model output in all state variables both in the prediction window. The prediction RMS are defined as

$$\text{RMS} = \sqrt{\frac{1}{DN} \sum_{n=N+1}^{2N} \sum_{a=1}^{D} (x_a(n) - y_a(n))^2}.$$

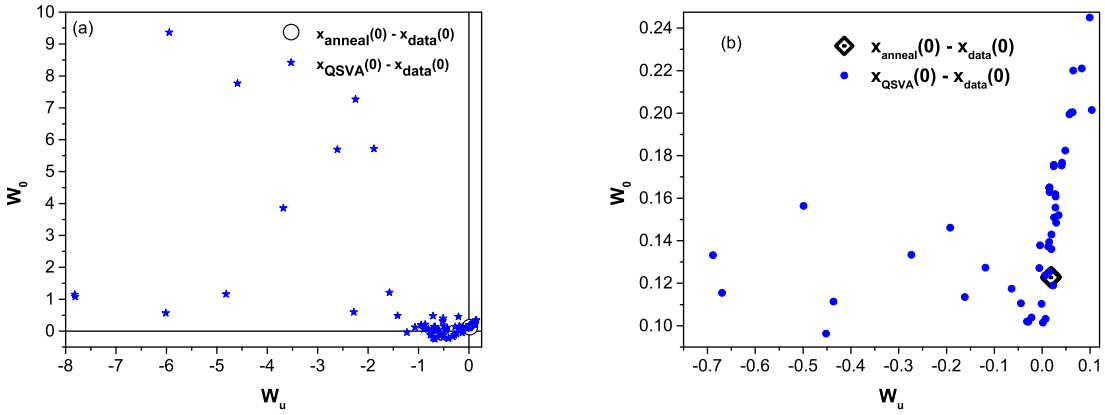

FIG. 11: (color online) Lorenz96, D = 5, L = 3, projection of $\mathbf{x}(0)$ on the neutral W_0 and unstable W_u manifolds for annealing with $\beta = 30$ and for QSVA; $N_0 = 100$ and $\Delta t = 0.025$. The **Panel (b)** zooms a portion of the **Panel (a)** near the annealing solution.

We performed a second set of calculations with $\Delta t = 0.001$ and $N_0 = 60$ initial conditions. We projected the $N_0 = 60$ solutions from QSVA into the plane of these two directions, unstable and neutral, along with the location in the same plane of the first state $\mathbf{x}(0)$ from the annealing process. The projection and a zooming in on the small values of the error are shown in Fig. 12. The histogram of RMS errors during the prediction window are then shown in Fig. 13.

Our conclusion is that the QSVA method, while straightforward to implement, may not do as well in the realistic case of noisy sparse data, $L < D$, which is likely to be encountered in realistic situations. It is clear that as one decreases the time step between iterations of the QSVA protocol, the error as seen in the projection onto the unstable and neutral directions at the end of the assimilation window decreases, yet there are quite a few instances when the error is quite large. Similarly, in the histogram of the RMS errors in the prediction window, as shown in Fig. 12 shows much better results than for the larger value of Δt in the first calculation.

The annealing results are rather the same in each case with a clear clustering of errors in the neutral/unstable plane near small errors, and a tight cluster of RMS prediction errors. We can conclude that the annealing method produces a much narrower distribution of candidates for the path with smallest action level when dealing with noisy, sparse data than the QSVA method in the same conditions. QSVA does produce a selection of paths with excellent initial conditions for prediction, but these come along with paths quite far from those with small errors. How one is to choose among the paths with large deviations from small errors is not entirely clear in the QSVA algorithm applied to the circumstances presented here.

The two critical differences in this comparison are that the annealing methods seeks `paths` as the outcome of the variational principles, and it surveys those paths and selects that which gives rise to the minimum observed value of the action. It should also be noted, as we mentioned earlier in the paper, that when one looks for a path with minimum action in the deterministic case, called strong 4DVar or large R_f, the procedure does not find the desired minimum [8]. QSVA utilizes models with no model error ($R_f \to \infty$) throughout.

VIII. SUMMARY

Transferring information from noisy, sparse measurements to a quantitative, predictive model of complex systems has been examined here in the context of variational principles based on Laplace's method [12] for approximating high dimensional integrals. Our attention is on the conditional expected value of functions $G(\mathbf{X})$ on the path $\mathbf{X} = \{\mathbf{x}(0), \mathbf{x}(1), ..., \mathbf{x}(m), \mathbf{p}\}$ of a D-dimensional model dynamical system when presented with L measurements $\mathbf{y}(t_n)$ at times within an observation window $\{t_0, t_1, ..., t_m\}$. The variational principle seeks extrema of the action $A_0(\mathbf{X}) = -\ln[P(\mathbf{X}|\mathbf{Y})]$, where $\mathbf{Y}$ is the collection of observations. In the presentation of the variational calculations, we provided a comparison of the formulations in discrete time and continuous time. Calculations are performed in the latter, though insights are often achieved in the former as well. By casting the question in terms of the expected value of functions along the path Eq. (1), we are able to both estimate the dominant constibutions to the expected value as well as the corrections associated with an expansion of the integrand about the allowed extremum paths where

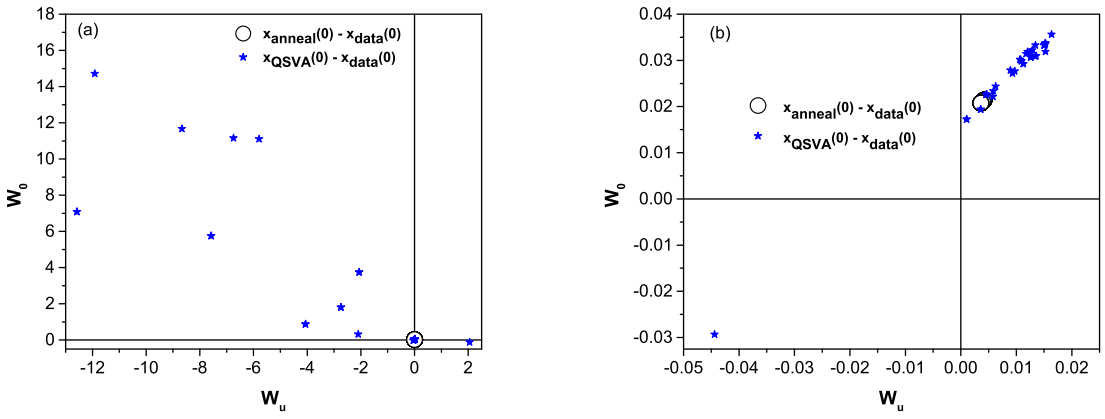

FIG. 12: (color online) Lorenz96, D = 5, L = 3, projection of $\mathbf{x}(0)$ on the neutral W_0 and unstable W_u manifolds for annealing with $\beta = 30$ and for QSVA; $N_0 = 60$ and $\Delta t = 0.001$. The **Panel (b)** zooms in on a portion of the **Panel (a)** near the annealing solution.

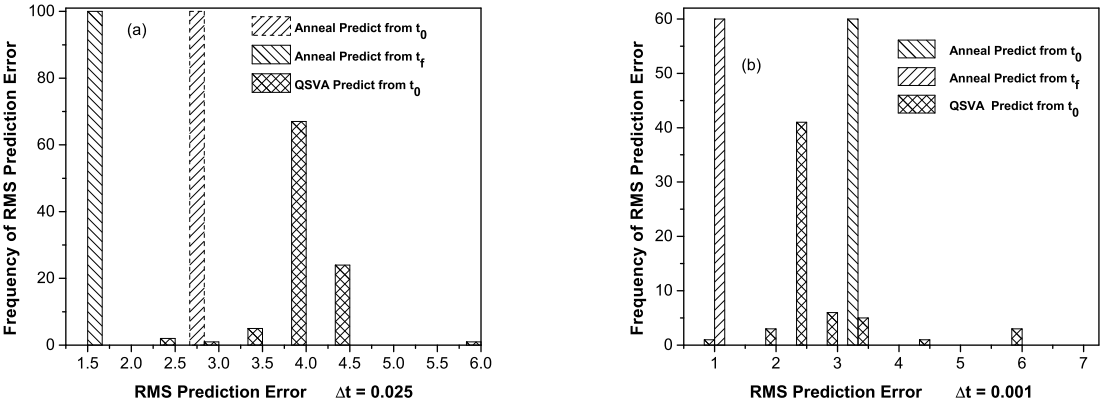

FIG. 13: Lorenz96, D = 5, L = 3, **Panel (a)** Histogram of RMS prediction error for annealing with $\beta = 30$ and for QSVA with $\Delta t = 0.025$, $N_0 = 60$. **Panel (b)** Histogram of RMS prediction error for annealing with $\beta = 30$ and for QSVA with $\Delta t = 0.025$, $N_0 = 100$

$\partial A_0(\mathbf{X})/\partial \mathbf{X} = 0$.

In this paper we also discussed some details of a procedure [13] for annealing, through the slow introduction of model accuracy, which yields the path having the smallest value of the action. This path dominates the integral with exponentially larger contributions than paths with larger action levels. In the case of a Gaussian error action Eq. (8) where the measurement errors are Gaussian and the model errors are Gaussian, we argued the as the accuracy of the model increases, namely R_f becomes large, the corrections to the contribution of the path with the lowest action level decrease as powers of R_f^{-1}. We do not have a similar argument for other distributions of model or measurement error.

We investigated several instructive models of dynamical systems: (1) the Lorenz96 model [35] with $D = 5$, (2) the Lorenz96 model with fast and slow variables [35], and (3) the standard Hodgkin-Huxley model with Na^+, K^+, and leak channels [38, 39]. In each case we found that the annealing method which starts at zero model resolution $R_f = 0$ then slowly increases R_f using the variational solution at any value of R_f to initialize the variational calculation at the next larger value of R_f shows action levels associated with the allowed saddle paths that split as R_f/R_m is about 100 leaving a lowest action level to dominate the expected value integrals.

A section of the paper compares the QSVA method [14] to the annealing method explored here. The Lorenz96 model with $D = 5$ and L = 3 noisy measurements was used in the comparison. The annealing method consistently produces allowed extremum paths with small errors and good predictions while the QSVA approach produces paths with a large range of errors, the best of which allow excellent predictions beyond the observation window.

ACKNOWLEDGMENTS

Partial support has come from the ONR MURI program (N00014-13-1-0205). We appreciate discussions with O. Talagrand and M. Jardak on QSVA.

[1] G. Evensen, *Data Assimilation: The Ensemble Kalman Filter* (Springer, 2009).
[2] A. C. Lorenc and T. Payne, Quarterly Journal of the Royal Meteorological Society **133**, 607 (2007).
[3] H. Eibern and H. Schmidt, Journal of Geophysical Research: Atmospheres (1984–2012) **104**, 18583 (1999).
[4] J. Liepe, P. Kirk, S. Filippi, T. Toni, C. P. Barnes, and M. P. Stumpf, Nature protocols **9**, 439 (2014).
[5] H. D. Abarbanel, *Predicting the Future: Completing Models of Observed Complex Systems* (Springer, 2013).
[6] I. M. Gelfand and A. M. Yaglom, Journal of Mathematical Physics **1**, 48 (1960).
[7] J. Zinn-Justin, *Quantum field theory and critical phenomena* (Oxford University Press, 2002).
[8] J. C. Quinn, PhD Dissertation in Physics, University of California, San Diego (2010), URL `http://escholarship.org/uc/item/0bm253qk`.
[9] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Baysian Data Analysis: Second Edition* (Chapman & Hall, CRC Press, 2004).
[10] R. H. Landau, M. J. Paez, and C. C. Bordeianu, *A Survey of Computational Physics: Introductory Computational Science* (Princeton University Press, 2010).
[11] A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, and A. Stuart, Bernoulli **19**, 1501 (2013).
[12] P. S. Laplace, Mémoires de Mathématique et de Physique, Tome Sixième pp. 621–656 (1774).
[13] J. Ye, N. Kadakia, P. J. Rozdeba, H. D. I. Abarbanel, and J. C. Quinn, Nonlinear Processes in Geophysics Discussions **1**, 1603 (2014).
[14] C. Pires, R. Vautard, and O. Talagrand, Tellus A **48**, 96 (1996).
[15] L. Onsager and S. Machlup, Physical Review **91**, 1505 (1953).
[16] B. Jouvet and R. Phythian, Physical Review A **19**, 1350 (1979).
[17] D. Hochberg, C. Molina-Paris, J. Perez-Mercader, and M. Visser, Physical Review E **60**, 6343 (1999).
[18] D. Liberzon, *Calculus of Variations and Optimal Control Theory* (Princeton University Press, 2012).
[19] D. E. Kirk, *Optimal Control Theory: An Introduction* (Dover Publications, Inc., 1970).
[20] M. Kot, *A First Course In the Calculus of Variations* (American Mathematical Society; Providence, Rhode Island., 2014).
[21] I. M. Gelfand and S. V. Fomin, *Calculus of Variations* (Dover Publications, Inc., 1963).
[22] J. Bröcker, Quarterly Journal of the Royal Meteorological Society **136**, 1906 (2010).
[23] J. Bröcker and I. G. Szendro, Quarterly Journal of the Royal Meteorological Society **138**, 785 (2012).
[24] J. Bröcker, Summer School/Creative Workshop: Data Assimilation & Inverse Problems (2013), URL `http://www.inverseproblems.info/reading:summer$\_$school$\_$2013`.
[25] A. F. Bennett, *Inverse Methods in Physical Oceanography* (Cambridge University Press, 1992).
[26] B. S. Chua and A. F. Bennett, Ocean Modelling **3**, 137 (2001).
[27] U. M. Ascher, R. M. M. Mattheij, and R. D. Russell, *Numerical solution of boundary value problems for ordinary differential equations*, vol. 13 (SIAM, 1995).
[28] D. S. Kendall Atkinson, Weimin Han, *Numerical Solution of Ordinary Differential Equations* (John Wiley & Sons, Inc., Hoboken, New Jersey, 2009).
[29] E. Süli, people.maths.ox.ac.uk/suli/nsodes.pdf pp. 1–82 (2014).
[30] W. H. Press et al., *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press, 2012).
[31] A. Wächter and L. T. Biegler, Mathematical programming **106**, 25 (2006).
[32] J. Ye, B. Toth, and C. Knowlton (2015), URL `https://github.com/yejingxin/minAone`.
[33] S. Kirkpatrick, M. Vecchi, et al., science **220**, 671 (1983).
[34] W. Feller, *An introduction to probability theory and its applications*, vol. 2 (John Wiley & Sons, 2008).
[35] E. N. Lorenz, in *Predictability of weather and climate*, edited by T. Palmer and R. Hagedorn (Cambridge, 2006).
[36] M. Kostuk, PhD Dissertation in Physics, University of California, San Diego (2012), URL `http://escholarship.org/uc/item/2fh4d086`.
[37] D. S. Wilks, Quarterly Journal of the Royal Meteorological Society **131**, 389 (2005).
[38] D. Johnston and S. M.-S. Wu, *Foundations of Cellular Neurophysiology* (Bradford Books, MIT Press, 1995).
[39] D. Sterratt, B. Graham, A. Gillies, and D. Willshaw, *Principles of Computational Modelling in Neuroscience* (Cambridge University Press, 2011).
[40] E. N. Lorenz, Journal of the atmospheric sciences **20**, 130 (1963).
[41] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, ACM Transactions on Mathematical Software (TOMS) **23**, 550 (1997).
[42] H. D. I. Abarbanel, M. Kostuk, and W. Whartenby, Quarterly Journal of the Royal Meteorological Society **136**, 769 (2010).