

Reconstructing Mammalian Sleep Dynamics with Data Assimilation

Madineh Sedigh-Sarvestani^{1*}, Steven J. Schiff^{1,2,3}, Bruce J. Gluckman^{1,2,4}

1 Center for Neural Engineering, Department of Engineering Science and Mechanics, The Pennsylvania State University, University Park, Pennsylvania, United States of America, **2** Department of Neurosurgery, The Pennsylvania State University, University Park, Pennsylvania, United States of America, **3** Department of Physics, The Pennsylvania State University, University Park, Pennsylvania, United States of America, **4** Department of Bioengineering, The Pennsylvania State University, University Park, Pennsylvania, United States of America

Abstract

Data assimilation is a valuable tool in the study of any complex system, where measurements are incomplete, uncertain, or both. It enables the user to take advantage of all available information including experimental measurements and short-term model forecasts of a system. Although data assimilation has been used to study other biological systems, the study of the sleep-wake regulatory network has yet to benefit from this toolset. We present a data assimilation framework based on the unscented Kalman filter (UKF) for combining sparse measurements together with a relatively high-dimensional nonlinear computational model to estimate the state of a model of the sleep-wake regulatory system. We demonstrate with simulation studies that a few noisy variables can be used to accurately reconstruct the remaining hidden variables. We introduce a metric for ranking relative partial observability of computational models, within the UKF framework, that allows us to choose the optimal variables for measurement and also provides a methodology for optimizing framework parameters such as UKF covariance inflation. In addition, we demonstrate a parameter estimation method that allows us to track non-stationary model parameters and accommodate slow dynamics not included in the UKF filter model. Finally, we show that we can even use observed discretized sleep-state, which is not one of the model variables, to reconstruct model state and estimate unknown parameters. Sleep is implicated in many neurological disorders from epilepsy to schizophrenia, but simultaneous observation of the many brain components that regulate this behavior is difficult. We anticipate that this data assimilation framework will enable better understanding of the detailed interactions governing sleep and wake behavior and provide for better, more targeted, therapies.

Citation: Sedigh-Sarvestani M, Schiff SJ, Gluckman BJ (2012) Reconstructing Mammalian Sleep Dynamics with Data Assimilation. *PLoS Comput Biol* 8(11): e1002788. doi:10.1371/journal.pcbi.1002788

Editor: Boris S. Gutkin, École Normale Supérieure, Collège de France, CNRS, France

Received: May 21, 2012; **Accepted:** October 3, 2012; **Published:** November 29, 2012

Copyright: © 2012 Sedigh-Sarvestani et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This work was supported by NIH/NINDS Grant R01NS065096 to BJG and SJS as well as NIH/NINDS Grant FNS070545A to MSS (<http://www.ninds.nih.gov/>). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: m.sedigh.sarvestani@gmail.com

Introduction

Great strides have been made in understanding the physiological basis for sleep regulation [1] in terms of the interacting neuronal cell groups and their neurotransmitter mediated interactions. This physiology is now increasingly being embodied into complex mathematical models of sleep dynamics [2–6]. But the limits to which these models are either validated or otherwise utilized for insight and prediction is currently limited. Due to physical and technological constraints, simultaneous measurement of the physiology embodied in the models - such as cell group firing rates and neurotransmitter concentrations - is not feasible in freely behaving animals or people. We demonstrate here that such models of the sleep-wake regulatory system can be put into a data assimilation framework that allows for reconstruction and forecasting of unobserved dynamics from limited noisy measurements. We anticipate these tools will help shed light on core brain circuitry implicated in sleep disorders as well as sleep-related neurological disorders such as epilepsy [7], bipolar disorder [8], and generalized anxiety disorder [9].

Data assimilation is an iterative process that couples and synchronizes mathematical models to observed system dynamics

with the purpose of estimating both noisy observed and unobserved variables, as well as forecasting the future system state. Data assimilation algorithms for nonlinear systems often employ the ensemble Kalman filters [10]. One such ensemble filter is the unscented Kalman filter (UKF), used in an iterative prediction-correction scheme in which model-generated predictions are corrected to agree with or track experimental observations [11].

The objectives of this article are to demonstrate data assimilation applicability within the context of relatively high-dimensional nonlinear biological models of the sleep-wake regulatory system, and to investigate the observability properties of these models [4,12]. In the Materials and Methods section, we introduce these models, as well as the basic mathematics of the UKF and parameter estimation algorithms. In the Results section, we demonstrate the use of the UKF to reconstruct data generated from these models. We introduce a reconstruction quantification that allows one to gauge the relative observability of the model variables. We demonstrate how this empirical observability coefficient can be used to optimize UKF parameters such as model covariance inflation, as well as how to select the optimal variables for measurement. We then demonstrate a method for

Author Summary

Mathematical models are developed to better understand interactions between components of a system that together govern the overall behavior. Mathematical models of sleep have helped to elucidate the neuronal cell groups that are involved in promoting sleep and wake behavior and the transitions between them. However, to be able to take full advantage of these models one must be able to estimate the value of all included variables accurately. Data assimilation refers to methods that allow the user to combine noisy measurements of just a few system variables with the mathematical model of that system to estimate all variables, including those originally inaccessible for measurement. Using these techniques we show that we can reconstruct the unmeasured variables and parameters of a mathematical model of the sleep-wake network. These reconstructed estimates can then be used to better understand the underlying neuronal behavior that results in sleep and wake activity. Because sleep is implicated in a wide array of neurological disorders from epilepsy to schizophrenia, we anticipate that this framework will enable better understanding of the link between sleep and the rest of the brain and provide for better, more targeted, therapies.

optimizing model parameters for tracking slowly varying dynamics. Finally, we demonstrate that we can use measurements of discretized sleep-state generated from the model, instead of specific model variables, to reconstruct unobserved model dynamics.

Materials and Methods

Data assimilation is an iterative process that couples and synchronizes mathematical models to observed system dynamics. For illustration of the data assimilation framework and validation of results, we use artificially generated data from the Diniz Behn and Booth (DB) model [4] or its extension by Fleshner, et al. (FBFD model) [12]. We then select a subset of the generated variables, to which we add noise, as our measured data set. This data set, with or without the correct parameters used to generate it, is then passed to the UKF to reconstruct the unobserved states (variables) and forecast future system states. Validation is carried out by quantitatively comparing the reconstructed estimates and parameters with the known original data set. A 4th order Runge-Kutta estimate with an integration time of 0.5 seconds is used for all simulations. The MATLAB code to produce each figure in the Results section is available at ModelDB (<http://senselab.med.yale.edu/modeldb/default.asp>) or can be provided by the authors upon request.

Within this section, we describe both the DB and FBFD models of the sleep-wake regulatory system. We then describe the main features of the UKF and parameter estimation algorithms.

Physiology of Sleep

Recent advances in single and multi-unit recordings have contributed to the growing knowledge of the mammalian sleep-wake regulatory system. The current prevailing hypothesis includes a *flip-flop* switch that regulates transitions between non rapid-eye-movement sleep (NREM) and wakefulness (Wake) [1]. Gamma-aminobutyric acid (GABA)-ergic ventrolateral preoptic nucleus (VLPO) neurons in the hypothalamus promote NREM. Monoaminergic cell groups in the brainstem, including the noradrenergic locus coeruleus (LC) and the serotonergic dorsal

raphe (DR) neurons, promote Wake. Mutual inhibition between these two groups causes each to promote its own activity by inhibiting the other's. McCarley and Hobson [13] described transitions between NREM and rapid eye movement sleep (REM) arising from predator-prey like interactions between cholinergic cell-groups in the brainstem, including the laterodorsal tegmentum (LDT) and pedunculo-pontine tegmentum (PPT), and the monoaminergic cell-groups in LC and DR. For a more in-depth overview of the literature, including controversial hypotheses for REM regulation, see [14].

More recently, orexin and adenosine have been implicated in further regulation of the sleep-wake system. Orexin producing neurons in the lateral hypothalamus have descending projections to all aforementioned monoaminergic and cholinergic cell groups and reinforce arousal, for a review see [15]. Extracellular adenosine has been found to increase during prolonged wakefulness in several cortical and subcortical regions [16], and has been proposed as a homeostatic accumulator of the need to sleep [17].

These dynamics are further modified by the circadian drive [18], regulated by the suprachiasmatic nucleus (SCN) in the hypothalamus, which sets a roughly 24-hour cycle affecting sleep and many other physiological functions. The SCN has indirect projections to the VLPO in the hypothalamus which results in inhibition of sleep during the day [19]. Here day is subjectively defined by species' dependent diurnal behavior, and refers roughly to 12-hour periods consisting mostly of active-wake behavior.

The SCN clock can be modulated by afferent cortical inputs in response to a variety of external cues. Food restriction studies have shown entrainment of the circadian cycle to food availability [20]. Light input from the melatonin expressing ganglion cells in the retina can also affect the SCN [21]. Retrograde trace studies have shown that a number of central nervous system sites innervate the SCN in the rat [22], though further study is needed to fully elucidate the involved circuitry. For instance, it is well known that lesions of the temporal lobe leading to epileptic seizures also affect the circadian clock [23,24], but the relevant brain circuitry has yet to be determined.

Diniz Behn and Booth (DB) Model of Sleep

The DB model [4], depicted in Fig. 1A, describes interactions among five distinct neuronal populations: two Wake-active groups, LC and DR; two groups in the LDT/PPT, one that is REM-active, denoted R; one active both in Wake and REM, denoted W/R; and one group active during NREM in the VLPO. As illustrated in Fig. 1A, these cell groups communicate through various transmitters: LC transmits norepinephrine (NE), DR transmits serotonin (5-HT), the two groups in the LDT/PPT transmit acetylcholine (ACh), and VLPO transmits GABA. Excitatory thalamic input is modeled by the variable δ and the brain's homeostatic sleep drive is represented by h . Sample output of this model's sleep-wake cycles, as well as mutual inhibition between Wake and sleep-active regions is shown in Fig. 1B.

Each cell group is described by its firing rate (F) and the concentration (C) of the neurotransmitter that it releases to post-synaptic populations. Cell group firing rates are a function of their input neurotransmitter concentrations, and evolve according to:

$$\dot{F}_X = \frac{F_{X\infty}(I_X) - F_X}{\tau_X} \quad (1)$$

Here I_X is a weighted sum of neurotransmitter i into cell group firing rate F_X , with coupling constants $g_{i,X}$;

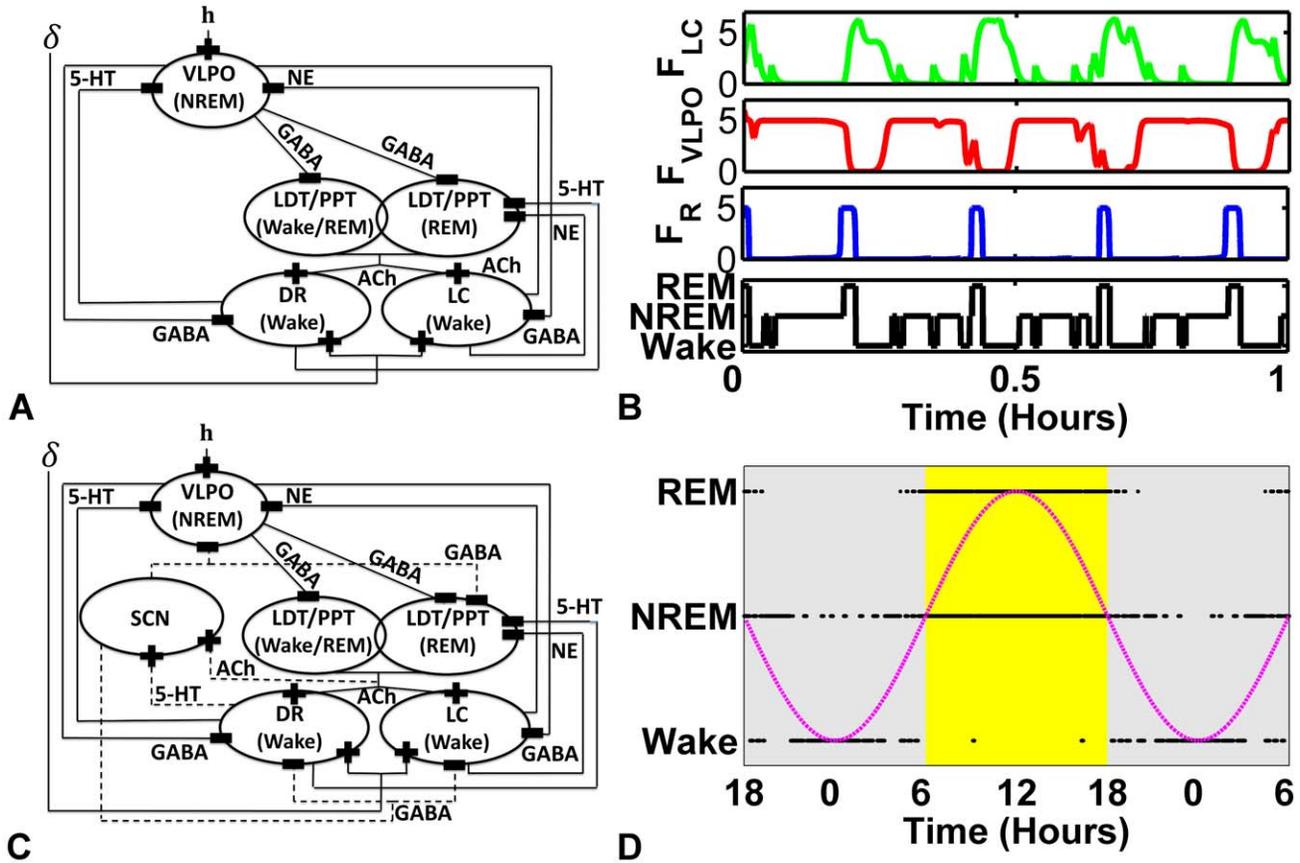


Figure 1. Computational models of the sleep-wake regulatory system and their outputs. A) Diniz Behn and Booth (DB) Model circuit diagram, illustrating the cell groups, their output neurotransmitters, and connections. Inhibitory connections are represented by minus and excitatory connections are represented by plus signs. Locus coeruleus (LC), dorsal raphe (DR), ventrolateral preoptic nucleus (VLPO), laterodorsal/pendunculopontine tegmentum (LDT/PPT), gamma-aminobutyric acid (GABA), serotonin (5-HT), norepinephrine (NE), acetylcholine (ACh), homeostatic sleep-drive (h), non rapid-eye-movement sleep (NREM), rapid-eye-movement sleep (REM). B) Typical output of the DB model for three of the cell group firing rates, plus the scored sleep-state plotted as a hypnogram. C) Fleshner, Booth, Forger and Diniz Behn (FBFD) Model circuit diagram, which expands on the DB model to include circadian modulation by and feedback to the suprachiasmatic nucleus (SCN), and allows for diurnal variations in behavior with light periods dominated by sleep activities and dark periods of extended awake activity. Dashed lines indicate SCN additions to DB model. D) 36 hour hypnogram from the FBFD model, with 24-hour periodic CIRC input input to the SCN superimposed. doi:10.1371/journal.pcbi.1002788.g001

$$\begin{aligned}
 I_{LC} &= g_{A,LC}C_A - g_{N,LC}C_N - g_{G,LC}C_G + \delta \\
 I_{DR} &= g_{A,DR}C_A - g_{S,DR}C_S - g_{G,DR}C_G + \delta \\
 I_{VLPO} &= -g_{N,VLPO}C_N - g_{S,VLPO}C_S - g_{G,VLPO}C_G \quad (2) \\
 I_R &= g_{A,R}C_A - g_{N,R}C_N - g_{S,R}C_S - g_{G,R}C_G \\
 I_{W/R} &= g_{A,W/R}C_A - g_{G,W/R}C_G
 \end{aligned}$$

In addition, τ_X is a first order process time constant. The steady state firing rate, $F_{X\infty}(c)$, a function of input neurotransmitter c , is given by maximum firing rate parameter X_{max} , times a sigmoidal function with midpoint β_X and slope α_X :

$$F_{X\infty}(c) = X_{max} \left(0.5 \left\{ 1 + \tanh \left[\frac{c - \beta_X}{\alpha_X} \right] \right\} \right) \quad (3)$$

where β_X is a constant for all cell groups except VLPO where β_{VLPO} is proportional to the homeostatic sleep drive h . The

concentration of neurotransmitter released by each cell group to the post-synaptic space also evolves according to a first order process with time constant τ_i and steady state neurotransmitter concentration $C_{i\infty}$ given by:

$$\dot{C}_i = \frac{C_{i\infty}(F_X) - C_i}{\tau_i} \quad (4)$$

$$C_{i\infty} = \tanh \left(\frac{F_X}{\gamma_i} \right) \quad (5)$$

where γ_i is an adjustable scale parameter.

Because ACh comes from both the R and W/R cell groups, the total ACh concentration in the post-synaptic space is the sum of the ACh concentrations generated individually from these groups. Random excitatory projections from thalamocortical circuits to the Wake-active populations LC and DR are modeled as Poissonian impulses with a rate of 0.003 Hz, which through a leaky integrator form another input concentration denoted δ with a decay constant of $\tau_{delta} = 10$ seconds:

$$\dot{\delta} = -\delta/\tau_{delta} \quad (6)$$

In addition to firing rate and neurotransmitter concentration variables, the homeostatic drive variable h regulates the duration of sleep and wake bouts by changing β_{VLPO} , the threshold for firing of the NREM-active VLPO cell group. The accumulation of h during Wake, and dissipation during sleep, is given by:

$$\dot{h} = H[(F_{LC} + F_{DR}) - \theta_w] \frac{1-h}{\tau_{hw}} - H[\theta_w - (F_{LC} + F_{DR})] \frac{h}{\tau_{hs}} \quad (7)$$

where H is the Heaviside function, θ_w is the threshold parameter for the onset of increase or decrease in h , and τ_{hs} and τ_{hw} determine the rate of accumulation and dissipation.

Typical output of the DB model is shown in Fig. 1B. The top three traces are the time dynamics of the firing rates for the Wake-active (LC), NREM-active (VLPO), and REM-active (LDT/PPT) cell groups. Note that, following [4] we denote firing rate for the REM-active LDT/PPT cell group as F_R . The state of vigilance (SOV), or sleep state, shown as a hypnogram in the fourth trace, is determined by the rank-ordered comparison of these cell group activities, with LDT/PPT activity dominating the definition.

Fleshner, Booth, Forger and Diniz Behn (FBFD) Model of Sleep

Fleshner, Booth, Forger and Diniz Behn [12] introduced an extension of the DB model that includes the SCN as an additional cell group with GABA as its associated neurotransmitter [12], depicted in Fig. 1C. The firing rate of the SCN cell group follows the same dynamics as the cell groups in the DB model (Eq. 1). The SCN has an inherent 24-hour circadian cycle (*CIRC*), with higher activity during the 12-hour light phase and lower activity during the 12-hour dark phase.

The projections from the sleep-wake network to the SCN provide dynamical feedback that increases the SCN's activity during both Wake and REM and decreases its activity during NREM. The SCN receives 5-HT and ACh synaptic inputs from the core sleep-wake regulatory system through the variable *SYN*. This is modeled by composing I_{SCN} from the sum of *CIRC* and *SYN*. Although the amplitude of *SYN* is smaller than that of *CIRC*, its oscillation time-scale is faster, typically on the order of minutes.

$$I_{SCN} = CIRC(t) + SYN(t) \quad (8)$$

$$CIRC(t) = \sin(2\pi(t - t_{daybreak})/\tau_{day}) \quad (9)$$

$$SYN(t) = g_{A,SCN} C_A + g_{S,SCN} C_S \quad (10)$$

Here $\tau_{day} = 24$ hours. We have shifted the phase of *CIRC* from [12] by adding $t_{daybreak} = 6$ to make the light period (*CIRC* high) start at 6 am. Feed-forward projections of the SCN on the sleep-wake network are mediated through GABAergic transmission, modeled by the additional neurotransmitter concentration $C_{G(SCN)}$, which adds into the dynamics of the LC, DR, VLPO, and R firing rates, modifying Eq. (2) from the DB model to become:

$$\begin{aligned} I_{LC} &= g_{A,LC} C_A - g_{N,LC} C_N - g_{G,LC} C_G - g_{G(SCN),LC} C_{G(SCN)} + \delta \\ I_{DR} &= g_{A,DR} C_A - g_{S,DR} C_S - g_{G,DR} C_G - g_{G(SCN),DR} C_{G(SCN)} + \delta \\ I_{VLPO} &= -g_{N,VLPO} C_N - g_{S,VLPO} C_S - g_{G,VLPO} C_G - \\ &g_{G(SCN),VLPO} C_{G(SCN)} \\ I_R &= g_{A,R} C_A - g_{N,R} C_N - g_{S,R} C_S - g_{G,R} C_G - g_{G(SCN),R} C_{G(SCN)} \\ I_{W/R} &= g_{A,W/R} C_A - g_{G,W/R} C_G \end{aligned} \quad (11)$$

Typical output of the FBFD model on short time-scales is similar to that of the DB model. But, as is typical for real rats, on diurnal time-scales the typical duration times in different states, as well as cycle times through states, changes. The hypnogram of the output of this model is shown in Fig. 1D for a 36 hour period. Rats are nocturnal. In the model, REM and NREM are primarily observed during the putative light phase, while long periods of Wake are observed during the putative dark phases.

Unscented Kalman Filter

The Kalman filter estimates the state of a system from noisy, sparsely measured, variables. Kalman's initial filter derivation [25] was for linear systems. The unscented Kalman filter is an ensemble version developed to tolerate nonlinearities without linearization [26].

The details of the UKF algorithm can be found in many standard textbooks [27,28]. We present here an overview, along with the key equations needed to understand details presented later in this manuscript.

State estimation with the UKF is carried out recursively using a prediction-correction scheme. Each iteration starts with a best estimate \hat{x}_k of the current state x at iteration time t_k . Included is an estimate of the current uncertainty in state $\hat{P}_{xx,k}$. A prediction or forecast is then generated by iterating an ensemble of points near \hat{x}_k , called sigma points, through the nonlinear model dynamics F . Given a D_x dimensional state space for F , we choose $2D_x$ sigma points such that they have covariance $\hat{P}_{xx,k}$ to represent the state uncertainty. We denote the i^{th} sigma point $\chi_{i,k}$ prior to iteration, and $\tilde{\chi}_{i,k+1}$ after iteration. The model prediction \tilde{x}_{k+1} is then the mean of the forward iterated sigma points:

$$\tilde{x}_{k+1} = \frac{1}{2D_x} \sum_{i=1}^{2D_x} \tilde{\chi}_{i,k+1} \quad (12)$$

The prediction uncertainty $\tilde{P}_{xx,k+1}$ is then the covariance of these points plus an additive covariance inflator matrix CI .

$$\tilde{P}_{xx,k+1} = \frac{1}{2D_x} \sum_{i=1}^{2D_x} (\tilde{\chi}_{i,k+1} - \tilde{x}_{k+1})(\tilde{\chi}_{i,k+1} - \tilde{x}_{k+1})^T + CI \quad (13)$$

CI is nonzero only on the diagonal, and is added to account for underestimates of the forecast error, from the covariance of the sigma points, due to process noise and inadequacies in the filter model [29,30]. The prediction is then corrected to account for a measurement y_{k+1} at time t_{k+1} . y need not contain the same number of variables as x . The correction factor weights the observation and prediction according to the Kalman gain K_{k+1} :

$$\hat{x}_{k+1} = \tilde{x}_{k+1} + K_{k+1}(y_{k+1} - \tilde{y}_{k+1}) \quad (14)$$

where \tilde{y}_{k+1} denotes the prediction mean from the estimated sigma points for the observed variables, $\tilde{Y}_{i,k+1}$:

$$\tilde{y}_{k+1} = \frac{1}{2D_x} \sum_{i=1}^{2D_x} \tilde{Y}_{i,k+1} \quad (15)$$

The Kalman gain is formed from the ratio :

$$K_{k+1} = \tilde{P}_{xy,k+1}(\tilde{P}_{yy,k+1})^{-1} \quad (16)$$

where $\tilde{P}_{xy,k+1}$ and $\tilde{P}_{yy,k+1}$ are formed from averages over the sigma points, either in the full dimensional space of x or in the subspace spanned by the measurements:

$$\tilde{P}_{xy,k+1} = \frac{1}{2D_x} \sum_{i=1}^{2D_x} (\tilde{x}_{i,k+1} - \tilde{x}_{k+1})(\tilde{Y}_{i,k+1} - \tilde{y}_{k+1})^T \quad (17)$$

$$\tilde{P}_{yy,k+1} = \frac{1}{2D_x} \sum_{i=1}^{2D_x} (\tilde{Y}_{i,k+1} - \tilde{y}_{k+1})(\tilde{Y}_{i,k+1} - \tilde{y}_{k+1})^T + R \quad (18)$$

where R is the measurement uncertainty. The Kalman gain is also used to correct - and ideally collapse - the prediction uncertainty:

$$\tilde{P}_{xx,k+1} = \tilde{P}_{xx,k+1} - K_{k+1}\tilde{P}_{xy,k+1} \quad (19)$$

Within this recursive scheme, the UKF synchronizes model state to measurements and thereby improves the estimate of the experimentally inaccessible variables. *The limit to which this succeeds depends in part on the relative observability of the reconstructed model variables from the measured variables. We discuss below an empirical method for assessing this relative observability.* We note that R , the uncertainty of the measurement process usually can be estimated, using the assumption that measurement noise is normally distributed [31–33].

On the other hand, the additive covariance inflation parameter CI is less clearly defined. Some methods have been proposed to estimate its values under the limited case that its source is an additive process noise [34–38]. Within our results - we demonstrate that even with identical system and model dynamics, non-zero CI improves tracking, and present a method of optimally choosing the values of CI for tracking and prediction.

Parameter Estimation

One approach for parameter estimation within the UKF framework is to solve the dual problem of estimating parameters and states at the same time, for instance via an augmented state space approach [11]. The alternative approach is to separate state reconstruction from parameter estimation by iteratively alternating between the two [39]. We found that dual estimation did not work well for our high-dimensional sleep-models, likely in part due to the many degrees-of-freedom when neither parameters nor variables were fixed, and especially because in nonlinear systems the sensitivity of the dynamics to particular parameters can be highly dependent on location in state space. We therefore estimate

parameters iteratively over windows of length T_{win} that are longer than a typical sleep-wake cycle of the dynamics.

Within our method, hidden states are first reconstructed with the UKF using a filter model with initial best-guess parameters. The full-state reconstruction over T_{win} is then used in a parameter estimation step which yields an updated parameter set. This updated parameter set is then passed to the UKF for the next iteration. This process is repeated until the parameter estimate has stabilized.

The parameter estimation step is essentially an application of a multiple shooting method [11,40]. Within each window, we estimate parameters by creating an average cost-function CF that quantifies the divergence between short model-generated trajectories \hat{x}_{traj} and the UKF-reconstructed trajectories *for the measured variables*. We then minimize this cost-function with respect to the parameter of interest. In order to prevent the model-generated trajectories from diverging too far from the reconstructed ones, we reinitialize \hat{x}_{traj} on the reconstructed trajectories at regular intervals dT :

$$\hat{x}_{traj}(k) = \tilde{x}(k) \text{ where } k = 1 * dT, 2 * dT, \dots, T_{win} \quad (20)$$

We then calculate a cost-function averaged over the window T_{win} using the divergence between the model-generated trajectories \hat{x}_{traj} and \tilde{x} :

$$CF = \int_0^{T_{win}} (\hat{x}_{traj} - \tilde{x})^T M (\hat{x}_{traj} - \tilde{x}) dt \quad (21)$$

where M denotes a matrix with non-zero elements on diagonal positions corresponding to measured elements. In order to properly weight the errors for each variable, the non-zero elements of M are set to the inverse of the standard deviation of the associated variable. For our current implementation, we perform a minimization with respect to parameters by explicitly computing $\hat{x}_{traj}(\tilde{p}_{traj})$ for test parameters $\tilde{p}_{traj} = \{\tilde{p}, \tilde{p} \pm \Delta p\}$ and then choosing the one with the minimum CF . We use a constant value of Δp , and restrict our parameter update maximally to Δp per iteration. Though somewhat computationally intensive, this method yields a stable approach to a local minimum in CF . This also limits the resolution to which the parameter can be estimated.

For estimation of non-stationary parameters, we use overlapping windows, with an update period of $T_{param} < T_{win}$. We note that $\Delta p / T_{param}$ should be greater than the maximal expected rate of change of the parameter of interest, to ensure that parameter dynamics are estimated with reasonable fidelity.

Results

Data Assimilation

We can accurately reconstruct unmeasured variables of the DB model of sleep with the UKF framework. To demonstrate this, we generate data from this model, then apply a noisy observation function - the output of which is a noisy subset of the variables - to mimic experimental conditions. We then reconstruct the unobserved variables with the UKF. Finally, we validate this reconstruction by comparing to the original data set.

An example of this procedure is shown in Fig. 2. Long multivariate time series of sleep-wake data were generated from the DB model. The observation function yielded a noisy univariate version of the firing rate of the Wake-active LC region F_{LC} . Explicitly we added random, normally-distributed, zero mean

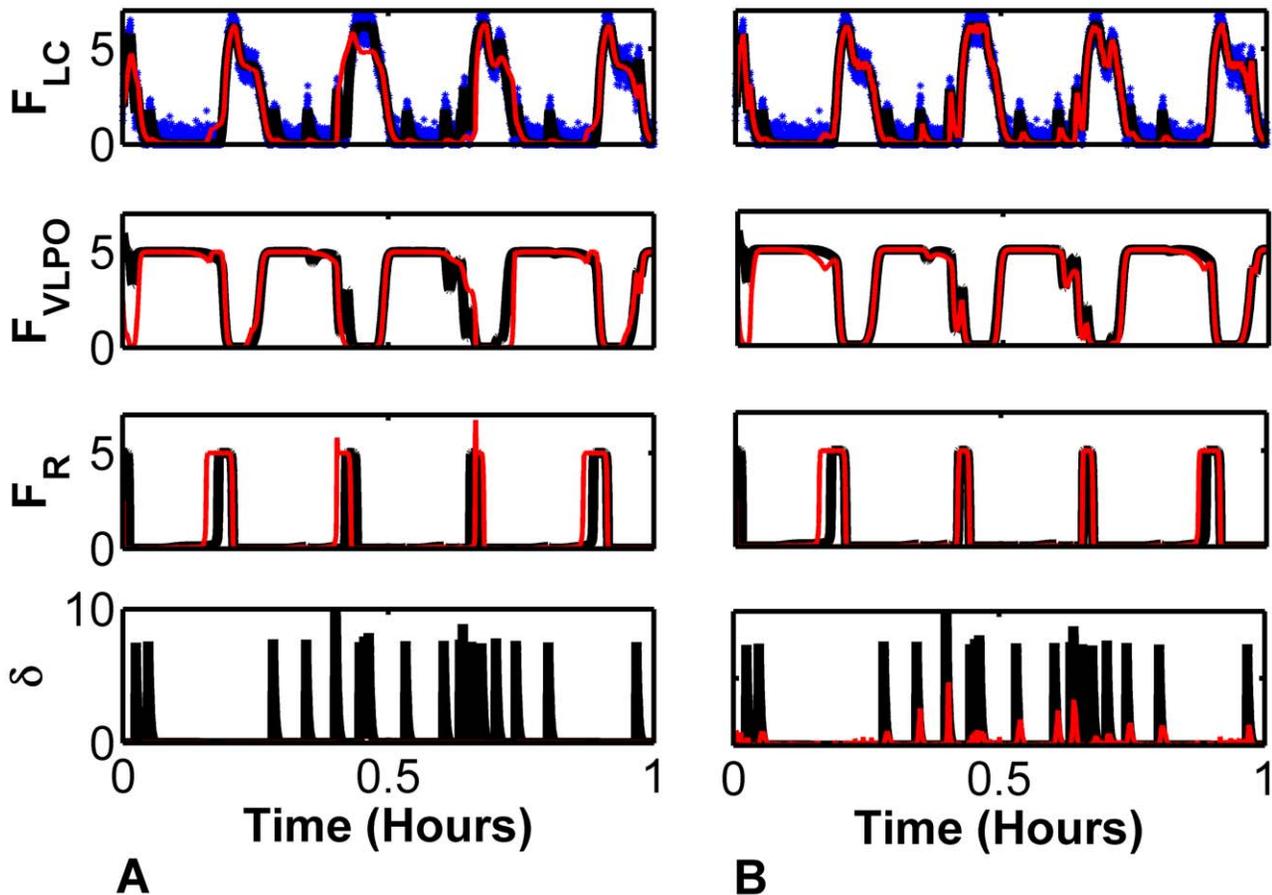


Figure 2. Reconstruction of DB Model Dynamics (A) with default values for covariance inflater CI , and (B) after optimization of CI_{F_R} and CI_δ . Noisy measurements of F_{LC} (blue) were passed to the unscented kalman filter (UKF) framework to track and reconstruct all other variables. Shown are the firing rates for the Wake-active (LC), NREM-active (VLPO), and REM-active (LDT/PPT) cell groups, along with thalamic noise δ . The framework was given the same parameters used to generate the original data. In both A and B the same data was tracked with model initial conditions chosen randomly. After a transient period, reconstructed (red) Wake and NREM dynamics are close to true (black) dynamics. Without CI optimization the dynamics of δ are essentially ignored. After CI optimization at least some of the stochastic δ dynamics - those that measurably affect the dynamics of F_{LC} - are reconstructed and reconstruction of REM dynamics is improved. doi:10.1371/journal.pcbi.1002788.g002

noise with variance of 4% that of the variance of F_{LC} to the true values. We provided the framework the parameters used to generate the original data, and either default values (left panels) or optimized values (right panels) for the covariance inflation parameter CI . Default values of CI were chosen as 10^{-5} times the typical variance of each variable. Additionally, the initial conditions of the model state, \hat{x} , were arbitrarily chosen in each case.

Shown in Fig. 2 are the reconstructed (red) and true (black) values of the NREM-active firing rate variable F_{VLPO} , the REM-active firing rate variable F_R and the stochastic thalamic noise variable δ . In both cases of tracking, reconstruction of the observed variable F_{LC} is good. This can be seen from the closeness of the reconstructed traces to both the observation points, shown in blue, and the true values in black. Likewise, the reconstruction of F_{VLPO} also tracks the true state quite well.

However, for the default values of CI , F_R is not reconstructed as well, and δ is not reconstructed at all. These errors extend to lower reconstruction fidelity of F_{VLPO} and even of F_{LC} . On the other hand, when we use optimized CI values the reconstruction of F_R is improved. In addition, much of the thalamic noise input through δ - which is stochastically driven and receives no input from the

other variables - is now represented. For these reconstructions, we initialized the model state far from that of the true system state. Therefore, there is a transient period during which reconstruction is poor. In our experience, once the model state comes close to that of the true system, this data assimilation framework keeps the model relatively close to the system state.

Quantification of reconstruction fidelity and observability. Visual inspection of the similarity between the reconstructed and true dynamics is only a qualitative result. We therefore use the mean square difference between the reconstructed (\hat{x}_i) and true (x_i) values for each variable to quantify the accuracy of state reconstruction. We normalize this error by the variance of each variable's dynamics to form ε_i^2 , a normalized mean square error for the i^{th} variable :

$$\varepsilon_i^2 = \frac{\langle (x_i - \hat{x}_i)^2 \rangle}{\text{var}(x_i)} \quad (22)$$

For perfect reconstruction, $\varepsilon_i^2 \rightarrow 0$. Its maximum depends on the ratio of the full range of the variable to the square root of its variance. For typical variables of the DB model this is of order 3,

though for some such as δ it is ~ 15 . For visualization and regularization purposes, we therefore use the inverse of $1 + e_{ij}^2$, which is bounded between $[0,1]$ as a reconstruction fidelity metric.

The reconstruction fidelity of any particular variable may depend on framework parameters such as integration or reconstruction time step, covariance inflation, as well as qualities that are inherent to the model dynamics such as its observability. Observability is a structural property of a model defined as the ability to recover the model state through the observation of one or more of its outputs [41]. It is well known that not all variables can be used as observables to reconstruct the full dynamics. Nonetheless, information regarding the partial observability of each variable can be used to choose the optimal variable for measurement in the UKF framework.

In Fig. 3, we show the reconstruction fidelity for each variable (down the columns) as a function of observation variable (across rows), in matrix format. We have used constant default values for CI , the covariance inflation parameter, relative observation noise R of 4% of the variable's variance, and no thalamic noise δ . In this color coded plot, red indicates good reconstruction, and blue indicates poor reconstruction.

The diagonal values in this matrix indicate how well a particular variable is reconstructed from itself. Though these tend toward the maximum, they are limited both by the noisy observation as well as the influence of the unobserved variables. Better reconstruction of the full dynamics is indicated by a column that has more red in it. To that extent the best observables for reconstruction are either F_R , the firing rate of the REM-active group, or its synaptic output $C_{A(R)}$. Likewise, the relative reconstruction of a variable from other variables can be gauged by the colors across its row.

Given observation of F_{LC} , the best reconstructed variable is $F_{W/R}$, the firing rate of the Wake/REM-active cell group, although several other variables are reconstructed quite well according to high values down the column marked F_{LC} . The worst reconstructed variables are F_R and its synaptic output $C_{A(R)}$. We interpret this to mean that F_R is relatively *less observable* from F_{LC} than is $F_{W/R}$. Furthermore, we assert that this is a useful empirical metric for gauging the partial observability of the state space from a measured variable, and name it the *Empirical Observability Coefficient (EOC)*:

$$EOC_{ij} = \frac{1}{1 + e_{ij}^2} \quad (23)$$

where e_{ij}^2 is the normalized reconstruction error for variable i given measurement j . Note for figure labels, we have left off the conditional reference to the measurements.

From the fourth column of the EOC matrix shown in Fig. 3, we observe that F_R is reconstructed well only if we observe F_R or $C_{A(R)}$, the concentration of ACh that this group transmits to the synaptic space. The poor reconstruction of F_R is in part due to symmetries in its dynamics. $C_{A(R)}$ is summed together with $C_{A(W/R)}$, the concentration of ACh transmitted from the Wake/REM-active cell group. As seen in Eq. (2), the total concentration of ACh C_A , then appears as one of the input transmitters to the Wake-active, REM-active, and Wake/REM-active cell groups. For further discussion of the effect of symmetry on observability and time series reconstruction, see [42].

We did not include dynamics of the thalamic noise variable δ in the model used to generate the *EOC* to highlight patterns in the *EOC* that would be obscured in the presence of this variable.

We note that this empirical observability is related to the *model* used in the UKF framework, not the true dynamics of the system

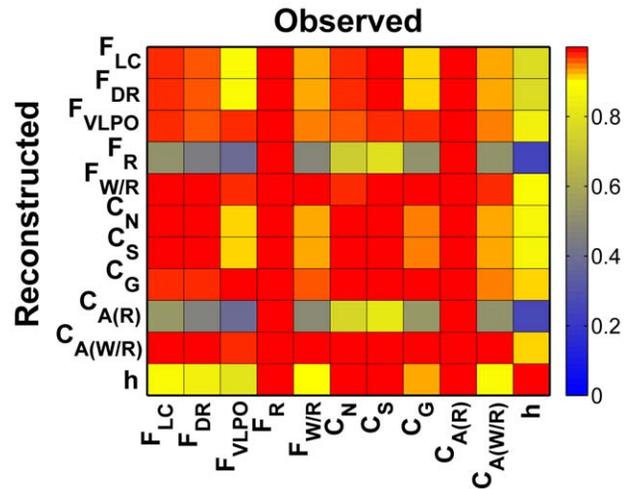


Figure 3. Empirical Observability Coefficient (EOC) matrix. *EOC* for the DB model with no thalamic noise and default values of CI . EOC_{ij} is an empirical measure of how well variable i is reconstructed from measurement of variable j . $EOC_{ij} \in [0,1]$ with perfect reconstruction being 1. Here *EOC* was computed using 12 hours of data. From the *EOC* matrix, we observe that F_R (row) is poorly observed - poorly reconstructed - from most variables, although its measurement (column) yields good reconstruction of almost all other variables. doi:10.1371/journal.pcbi.1002788.g003

being tracked. In other words, we assess observability in the model. Assuming that the model represents some of the underlying true system dynamics, then those aspects of the true system will also be observable. The computation of the *EOC* should be done from data generated from the model, not from observed data of the true system.

EOC Optimization of Covariance Inflation Parameter

We can find optimal framework parameters, such as the UKF covariance inflation parameter matrix CI , by maximizing *EOC* values. Although the matrix CI only has nonzero diagonal terms CI_i , for the full DB model including the thalamic noise output variable δ , there are 12 CI_i 's. So blind simultaneous optimization is inefficient. But we can use the full *EOC*, and the ranked partial observability, as a guide to this optimization.

Note from Eq. 13 that CI_i adds to the diagonal elements of the covariance of the sigma points. This inflation has the effect of widening the sigma points on the next iteration step, which results in an increase in the Kalman gain. Larger values for the Kalman gain bias the correction towards the measurements.

Our general rule therefore is that if measurement of a variable k yields poor reconstruction of other variables - i.e. low values of $EOC_{i,k}$ down a column - then we should favor measurement derived values of other variables over model derived ones, and therefore should use *increased* values of CI_k . On the other hand, if a variable k is not reconstructed well from other variables - i.e. low values of $EOC_{k,j}$ across the row - we should favor model derived values over measurement derived values for this variable by decreasing CI_k .

We iteratively compute the *EOC*, then choose the variable k with the lowest scores down a row or column, and change its corresponding CI_k appropriately. We then recompute the *EOC* and repeat. This prevents us from optimizing with respect to CI_k 's that have only modest impact on reconstruction fidelity.

There is a finite usable range for CI_k . As an inflator for the covariance matrix P_{xx} , CI_k must be greater than or equal to zero. The standard deviation and range of the dynamics of variable k

are two natural scales that can be used to define the usable range of CI_k . We use the square of the former, multiplied by a proportionality constant, as a default starting value for CI_k . The square of the latter forms the maximum for CI_k .

We now demonstrate this algorithm to optimize reconstructions given measurements of F_{LC} , as in Fig. 2B. The full EOC for the DB model - including δ - with default CI is shown in Fig. 4A. From the far right column, we observed that *no* variables are reconstructed well from measurements of δ . This is understandable, since the dynamics of δ receive no input from any of the other variables. Therefore we start our optimization of CI by adjusting CI_δ , and explicitly expect to increase it.

Shown in Fig. 4D is $EOC_{\delta, F_{LC}}$ as a function of increasing CI_δ . Although only the trace for δ is shown, $EOC_{i, F_{LC}}$ increases for most variables as a function of increasing CI_δ . We pick optimal values for CI_δ based on the average peak reconstruction of all variables from measurement of F_{LC} , found with a value of $CI_\delta = 0.06 * var(\delta)$. The EOC matrix after this first CI optimization iteration is shown in Fig. 4B.

Notably, although F_{LC} is the variable measured from the real system, its reconstruction improves when CI_δ is increased. This effect can be further understood by inspection of Fig. 2. The brief increases in F_{LC} from its low value - interpreted behaviorally as brief awakenings that correlate with spikes in δ in Fig. 2, are better reconstructed with optimized CI_δ . Indeed, the EOC

matrix values overall, shown in Fig. 4B, have increased with increasing CI_δ .

Now the row/column with the lowest values, on average, corresponds to reconstruction of F_R . Therefore we expect to need to decrease CI_{F_R} to improve reconstruction. Reconstruction fidelity of F_R from measurement of F_{LC} , as measured by $EOC_{F_R, F_{LC}}$ is shown in Fig. 4E as a function of CI_{F_R} . Reconstruction improves with decreasing values over the potential usable range. As shown by the black horizontal line, the best reconstruction is achieved using the minimum value of 0 for CI_{F_R} , although values of CI_{F_R} smaller than the default value of 10^{-5} result only in marginal reconstruction improvement. This second optimization step yields only marginal improvement in the overall EOC matrix shown in Fig. 4C. In part, this lack of improvement in reconstruction is due to the poor observability of REM dynamics through other variables as apparent from the row marked F_R in Fig. 3.

Empirical Observability and Choice of Measured Variables

We investigated pairings of two or more variables with respect to their relative partial observability. We found that for the DB and FBFD models, the empirical observability of variable i given measurements of variables j, k is always at least as good as the individual EOC : $EOC_{i, (j, k)} \geq \max\{EOC_{i, j}, EOC_{i, k}\}$. We also observed that good reconstruction of all variables requires some

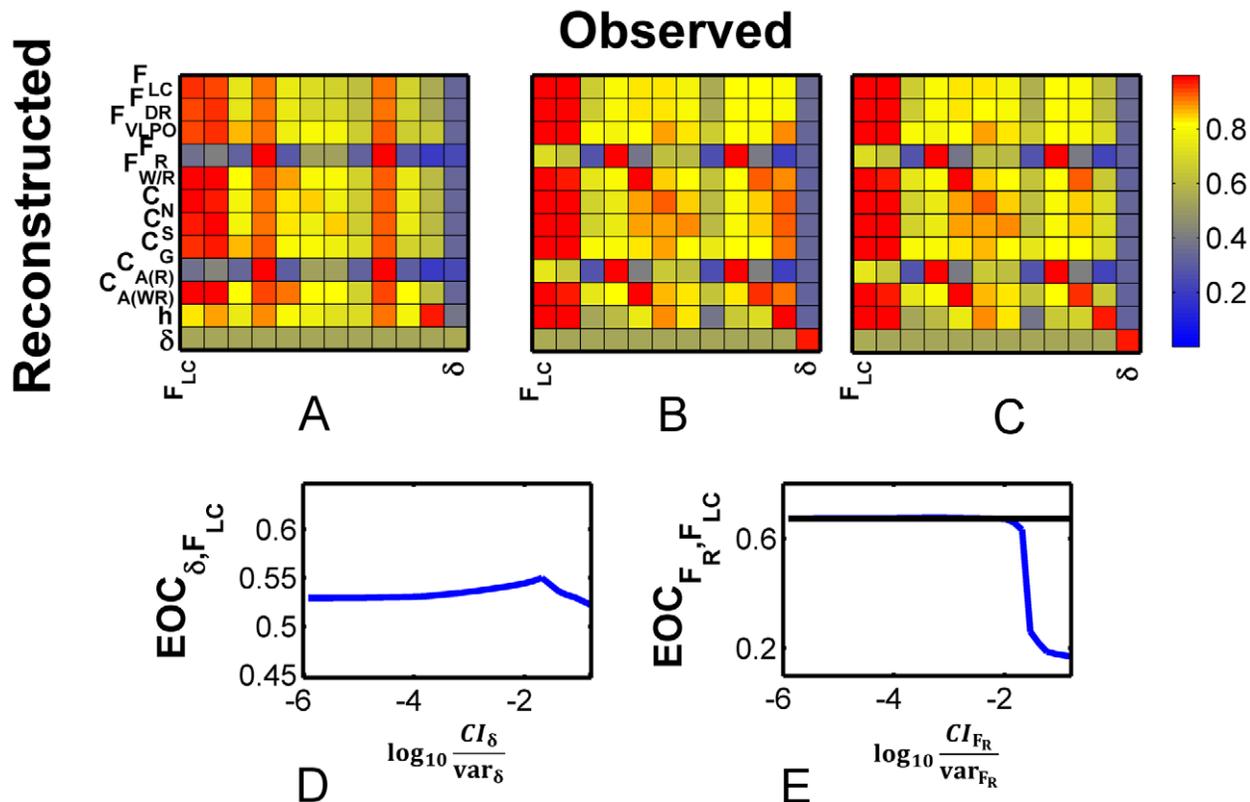


Figure 4. Optimization of Covariance Inflator CI . Although the individual EOCs are metrics of reconstruction fidelity, the ranked observability, from the full EOC can be used to guide optimization of the covariance inflator CI_k : Poorly observed variables across their rows - low $EOC_{k,j}$ - should have decreased CI_k . Variables whose measurement yields poor reconstruction columnwise - low $EOC_{i,k}$ - should have increased CI_k . Algorithmically, we iteratively adjust CI_k for the variable k with the overall lowest mean row or column. In A–C are shown the EOC matrix after each optimization iteration for the full DB model with thalamic noise. A) EOC computed with default values for CI , i.e. $CI_k = 10^{-5} var(k)$. Note that the lowest mean row/column corresponds to the measurement of δ , therefore we optimize CI_δ first. B) EOC after optimization of CI_δ . C) EOC after optimizing CI_{F_R} . Shown are $EOC_{i, F_{LC}}$ as a function of D) CI_δ for optimization step between A and B and E) CI_{F_R} for optimization steps between B and C. Optimal values of CI_i are chosen from the peaks of these plots. doi:10.1371/journal.pcbi.1002788.g004

measurement of both Wake and REM dynamics. These states are readily observed from real biological systems from external physiological measurements such as power bands in the EEG, muscle tone, and eye movement. Therefore, for the subsequent computations, we assimilate noisy measurements of both Wake-active F_{LC} and REM-active F_R dynamics, and use them to reconstruct the full system state and derive parameter values.

Parameter Fitting

As applied here, the UKF framework requires both a model for the dynamics as well as the model's parameters. We have implemented a version of a multiple shooting method [11] for optimizing the choice of parameters. The performance of this method is illustrated in Fig. 5.

For illustration purposes we generated data with fixed parameters and assimilated noisy measurements of F_{LC} and F_R to reconstruct the dynamics. Initially, all model parameters in the UKF were set to the same values used to generate the true data set - except the parameter g_{ALC} that couples ACh into F_{LC} dynamics. To this we supplied an arbitrary initial value.

Parameter estimation is performed by minimizing the distance between the UKF reconstructed traces and short model-generated trajectories that originate on the reconstructed traces. For these computations, we set the length of these short trajectories at

2 minutes. This is long enough that differences in parameters result in measurable divergence between the short computed trajectories and the reconstructed dynamics. Here *measurable* is much larger than the measurement noise, but not so large that the distance between the computed and reconstructed trajectories becomes comparable to the range of the state space.

To sample the full state space, each step of this minimization averages this divergence over time windows longer than the sleep-wake cycle time of the dynamics. As seen in Fig. 5A, our estimation of g_{ALC} converges to the true value. In Fig. 5B, we plot trajectories for the short model-generated (magenta), reconstructed (red), and true (black) F_{LC} dynamics for different periods of the convergence of g_{ALC} . Note that initially, for g_{ALC} significantly different than the true value, the short trajectories diverge quickly from the reconstructed values, and the reconstructed values of F_{LC} are different from the true ones. When g_{ALC} approaches the true value, both the short model-generated and reconstructed trajectories approach the true dynamics.

As coded, the parameter estimation essentially optimizes short model-generated forecasts. To investigate the effect on reconstruction fidelity, we compute the normalized mean square reconstruction error ε_i^2 for each variable, averaged over each parameter estimation window. This is shown for variables F_{LC} , $F_{W/R}$, and the homeostatic drive h . We note that for initial values

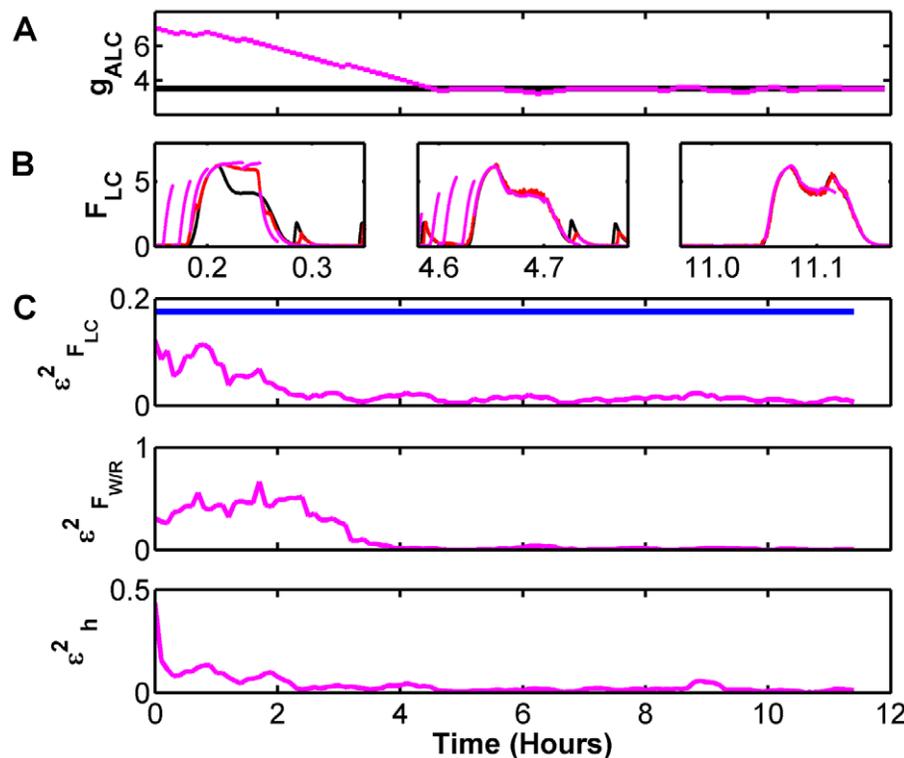


Figure 5. Parameter estimation with multiple shooting method for reconstruction of DB model from measurement of F_{LC} and F_R with unknown value for parameter g_{ALC} . Parameter estimation is performed by minimizing the divergence between the UKF reconstructed dynamics and short model-generated trajectories that originate on the reconstructed trajectories. To sample the full state space, each step of this minimization averages this divergence over time windows longer than the cycle time of the dynamics. Here we use half hour windows, with 80% overlap. A) Convergence of the estimated parameter g_{ALC} to the true value. B) Trajectories for the short model generated (magenta), reconstructed (red), and true (black) F_{LC} dynamics for different periods of the convergence of g_{ALC} . Note that initially, for g_{ALC} significantly different than the true value, the short trajectories diverge quickly from the reconstructed values, and the reconstructed values of F_{LC} are different from the true values. When g_{ALC} approaches the true value, both short model-generated and reconstructed trajectories approach the true values. C) Reconstruction metric ε_i^2 computed for each data assimilation window for three of the variables. As a reference point, the reconstruction metric for the original noisy observation of F_{LC} is shown in blue. Note that although the parameter estimation essentially optimizes short model generated forecasts, it has the effect of optimizing hidden variable reconstruction. doi:10.1371/journal.pcbi.1002788.g005

of this parameter, even reconstruction of the measured variable F_{LC} is quite poor - with typical errors $>30\%$ of its standard deviation. As a reference point, the initial measured data - a noisy version of F_{LC} - has a normalized mean squared error $\varepsilon_{F_{LC,meas}}^2 = 0.04$, shown as a horizontal blue line. As the estimated parameter converges, $\varepsilon_{F_{LC}}^2$ falls well below $\varepsilon_{F_{LC,meas}}^2$, and the reconstruction metric improves for all variables.

Dynamical Parameter Tracking

We can also estimate parameters which change slowly over time. We demonstrate this by using a slightly modified DB model, which lacks any circadian dynamics, to reconstruct dynamics observed from the expanded FBFD model which specifically includes SCN driven circadian oscillations. We use this modified DB model to assimilate noisy measurements of F_{LC} and F_R from the full FBFD model, and use it within the multiple shooting method to estimate the value of $C_{G,SCN}$.

An example of the output is shown in Fig. 6 for a 1.5 day period. We have skipped the initial 12 hours which includes a transient period of convergence of both the filter and the parameter estimate. The effect of the SCN is to modulate the overall sleep cycles, with frequent sleep periods that include REM in the light period and dominant, longer Wake periods in the dark period. Short example time series for F_{LC} and F_{VLPO} are shown in the panels in Fig. 6A for

different phases of the circadian cycle. The filter model in the UKF, also used in parameter estimation, is missing these SCN associated variables and the fast feedback oscillations resulting from their interaction with the sleep network. However, we replace the input contribution of the SCN's feed forward GABAergic projections on to the sleep network to a single quasi-static parameter $C_{G,SCN}$ that gets added to other neurotransmitter variables in Eq. (11). We then estimate this parameter which represents the presumed SCN drive.

The estimated value for $C_{G,SCN}$ (magenta) is shown in Fig. 6B, along with the true input from SCN in the generating model (black). Though the reconstructed parameter is an estimate with inherent averaging over half-hour periods, and therefore does not reproduce the fast dynamics of the real SCN input, it tracks the mean value quite well. In addition, it yields good reconstruction of the model variables. Examples of the normalized reconstruction error, averaged over the fitting windows, are shown in Fig. 6C for sample variables. Here again, as a reference point, we plot the mean squared error for the noisy measurement of F_{LC} (blue line) in the top panel of Fig. 6C. Note that even reconstruction of the homeostatic sleep drive h , which has no direct coupling to the observed variables, is quite good over most of the day.

Reconstruction from the Hypnogram

So far, we have implemented the data assimilation framework using measurements that amount to noisy versions of the true

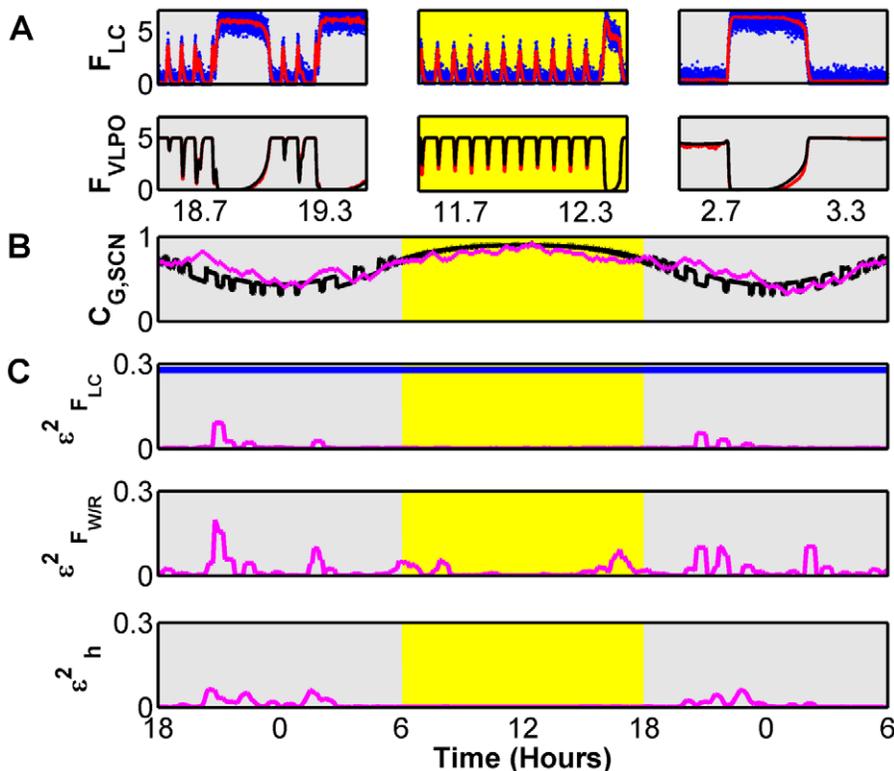


Figure 6. Parameter Tracking to accommodate circadian dynamics. Noisy measurements of F_{LC} and F_R from the full FBFD model were assimilated with a version of the DB model that represented input from the SCN as a quasi-static parameter $C_{G,SCN}$ whose value was estimated and tracked in 80% overlapping half hour windows. SCN activity imposes circadian and light-driven dynamics that modulate sleep-wake cycles and prevalence of either sleep or wake activity. A) Short excerpts of reconstructed dynamics for various phases of the circadian cycle. B) Estimated (magenta) and true (black) value of the tracked parameter $C_{G,SCN}$. Note that the tracked value is an estimate, with inherent smoothing on the time scale of a half hour, and therefore does not reconstruct all of the detailed dynamics of the true value which oscillates due to the interplay between the core sleep-wake regulatory cell groups and the SCN. C) Normalized reconstruction error for various variables. As a reference, the reconstruction error for the noisy F_{LC} measurement is shown in blue. The reconstruction of unobserved variables $F_{W/R}$ and homeostatic sleep drive h is quite good as indicated by small ε^2 values.

doi:10.1371/journal.pcbi.1002788.g006

variables. In real applications, when one uses observations from real systems, the actual system measurements might only remotely resemble variables in the tracking model. But even in this case, data assimilation methods can still be used. To this end, we demonstrate that we can use measurements of state-of-vigilance (SOV) generated from the model and illustrated in Fig. 1B, to reconstruct the unobserved model dynamics with reasonable fidelity.

The method we have implemented is illustrated in Fig. 7. We sleep-score the model-generated data, also used in Figs. 2–5, by assigning an SOV to each point as a function of time. The SOV is determined based on relative values of F_{LC} , F_{VLPO} , and F_R . We then take the filter model, and generate example data, which we also sleep-score. From this scored filter-model data, we compute the probability distribution functions (pdf) for the variables F_{LC} , F_{VLPO} , and F_R conditioned on SOV. These are illustrated in Fig. 7A. Note that these state-dependent distributions are highly skewed, and have small variance around the mean.

The observation function from the measured values - here SOV as a function of time shown in Fig. 7B - must provide values and error estimates of variables in the filter model to the UKF. To translate the observed SOV to inputs to the UKF, we use the state-conditioned medians from the above-generated pdfs, and then use the state-conditioned standard deviations as the measurement uncertainties. In this way, we use observations of SOV to infer observations of the model variables. We then use these observations, shown for F_{LC} in blue in Fig. 7C, as inputs to the UKF.

Note that in this case, the measurement noise estimates are time dependent. After a short transient convergence time, the reconstructed dynamics converge close to the true dynamics. However, certain details such as brief awakenings and transitions into NREM are not reconstructed well.

We can likewise apply all the other tools described here to assimilation of SOV data through this inferred observation function. Shown in Fig. 8 is the same parameter estimation procedure as shown in Fig. 5, with the same initial conditions for unknown parameter g_{ALC} . Although the convergence is not as good as with direct observation of F_{LC} , the estimated parameter does approach the parameter used to generate the data. The reconstruction error in F_{LC} decreases as the parameter approaches its correct value, however neither converge all the way. This can be understood because the UKF attempts to constrain the observed variables to the median values mapped from the SOV. Likewise, the parameter estimation algorithm attempts to minimize the error between model forecasts and reconstructions for the observed variables.

As a supplemental performance metric, we also consider the reconstruction error if we simply use the median observation map for all variables as our reconstruction. These are plotted as horizontal dashed lines for each variable in Fig. 8C. The UKF reconstruction error for the observed variable F_{LC} improves beyond this reference point as parameter estimation improves. In contrast, the UKF reconstruction errors for unobserved variables such as $F_{W/R}$ and h are overall far better.

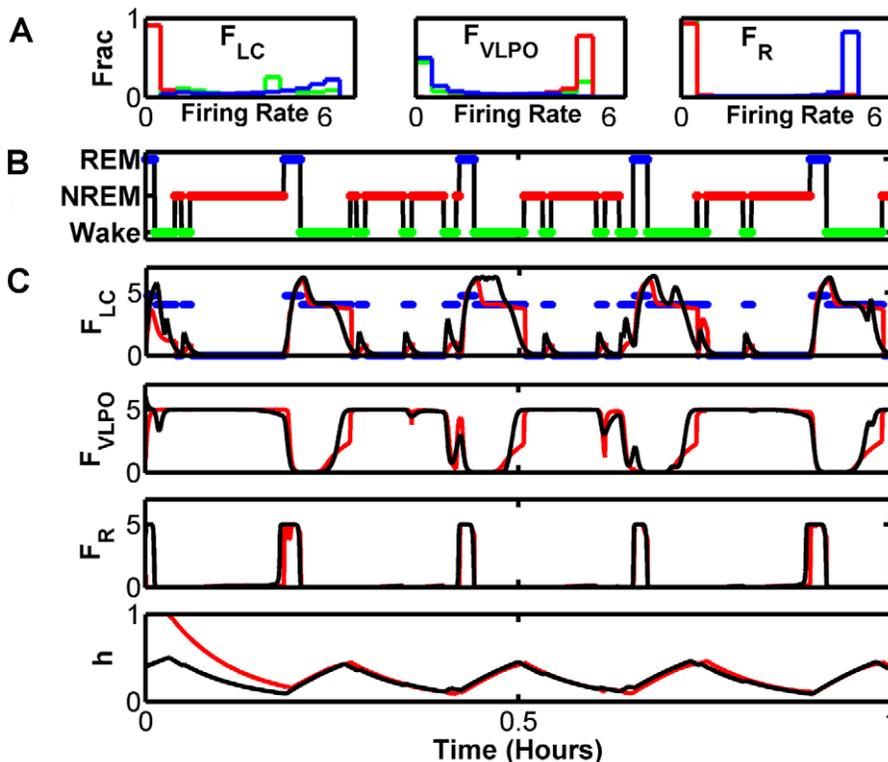


Figure 7. Reconstruction of DB dynamics from measured hypnogram. SOV is used along with an inferred observation function to translate an observed hypnogram into state conditioned observations for F_{LC} , F_{VLPO} , and F_R , and their variances. We use the UKF to reconstruct the full variable state space from these observations. A) Probability distributions of firing rates for F_{LC} , F_{VLPO} and F_R during Wake (black), NREM (red), and REM (blue). These firing rates were generated from the filter-model. B) Hypnogram of observed SOV for a 1 hour time series, with colors to match (A). C) Reconstructed (red) and true (black) traces for F_{LC} , F_{VLPO} , F_R , and h . The inferred observation for F_{LC} is also shown (blue). After a transient period, the reconstruction converges to the true value, even for the homeostatic drive variable h which was not observed. However, details of the dynamics that are not accounted for by the state-of-vigilance (SOV) such as brief awakenings and transitions into and out of NREM are not reconstructed well. doi:10.1371/journal.pcbi.1002788.g007

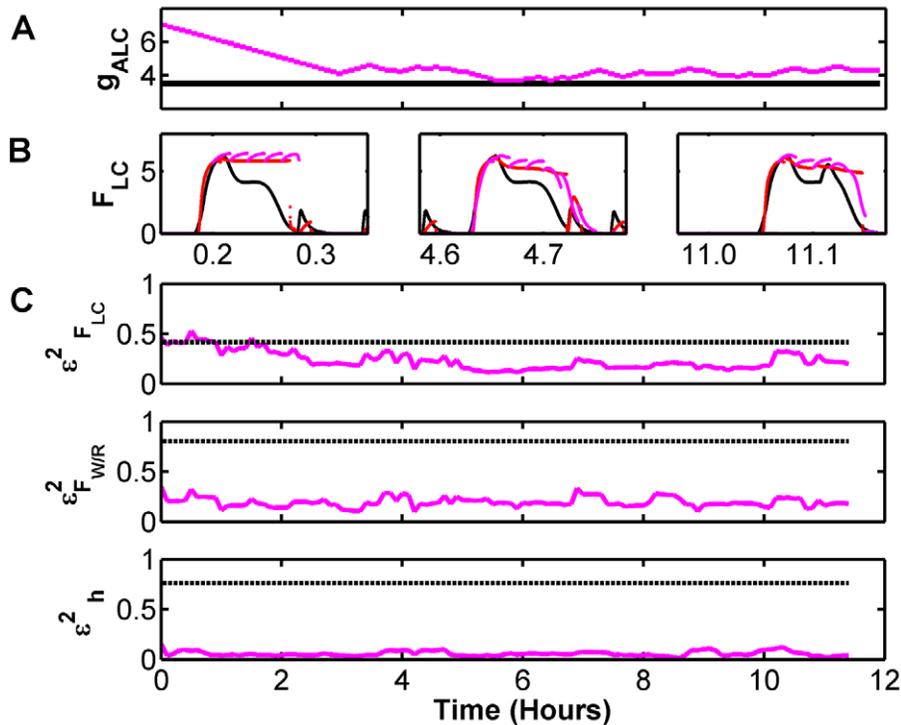


Figure 8. Parameter estimation from observed hypnogram for reconstruction of DB model from inferred measurement of F_{LC} , F_{VLPO} , and F_R with unknown value for parameter g_{ALC} . A) Convergence of the estimated parameter (magenta) to the true value (black). B) Trajectories for the short model-generated (magenta), reconstructed (red), and true (black) F_{LC} dynamics for different periods of the convergence of g_{ALC} . C) Reconstruction metric ϵ_F^2 computed for each data assimilation window for three of the variables. Horizontal dashed lines correspond to ϵ_F^2 computed from the state-conditioned discrete map used to translate the SOV to model space. Note that once the parameter is optimized, the UKF reconstruction far outperforms the observation map. doi:10.1371/journal.pcbi.1002788.g008

Discussion

Data assimilation is a valuable tool in the study of any complex system, where measurements are incomplete, uncertain, or both. It enables the user to take advantage of all available information including experimental measurements and short-term model forecasts of a system. Since the introduction of the UKF to neuronal dynamics by Voss et al. in 2004 [11], a few investigators have applied these methods to the study of biological systems [43–48]. Other data assimilation techniques have also been successfully applied to study neuronal dynamics [49]. Nevertheless, the sleep modeling community has yet to utilize these resources. Several important advantages of data assimilation in sleep modeling are : 1) access to unmeasured variables to create a more complete estimate of model state 2) subject-specific parameter estimation even when the parameter is associated with an unobserved variable 3) allowance for uncertainty in model structure or measurements and 4) prediction of future dynamics.

Not all variables perform equally in reconstructing the state space. In biological experiments utilizing data assimilation it would be beneficial to have some insight into the relative performance of each variable so that we can choose the best one or ones for measurement. A natural metric to guide this choice is the observability of each variable. Letellier et al. showed in [50] that observability and ability to synchronize are related. Since the UKF is basically a synchronization scheme, it follows that reconstruction-performance by any variable should be a function of its observability. Thus we propose using observability based metrics in the study of partially observed biological systems.

Analytical methods to determine observability for nonlinear systems are mathematically rigorous, require rational models, and generally do not produce graded values for partial observability. Letellier et al. [51,52] proposed a simple algebraic solution to rank all variables of a system according to their relative partial observability. Although their approach works well for low-dimensional systems, we found it problematic for our high-dimensional sparsely connected system, where many variables are directly coupled to just one or two other variables, and where the coupling is effectively on only in highly localized regions of state space.

Inspired by their work, we developed an empirical metric, the *EOC*, to rank the partial observability of each variable based on reconstructed error. The *EOC* can be used to select the optimal observed variable to obtain the best estimate of a particular unobserved variable. The absolute optimal observed variable receives as input to its dynamics unambiguous invertible information about the state of the unobserved variable. Here invertible implies a one-to-one (bijetvie) relationship between the unobserved and observed variables. In complex networks, this observability is modulated by the number and relative weights of additional unobserved variables in the system that couple into the dynamics of the observed variable [50–52].

Because the *EOC* is a measure of reconstruction fidelity, we demonstrate that the reconstruction framework parameters can be optimized by improving it. Importantly, we described an intuitive approach to use the *EOC* to optimize the covariance inflation parameters *CI*. Although some analytical methods have been proposed for this task in nonlinear systems [34–38], we are

unaware of an observability-based metric for covariance inflation optimization.

Correct parameter estimates aid the prediction of future dynamics and model selection and verification and can provide useful biological information. A common method for parameter estimation in nonlinear models utilizes a feedback-synchronization scheme, developed by [53] and extended by [54–58] and many others. Within such a scheme, two identical - except for unknown parameters - systems are unidirectionally coupled, and are continuously synchronized through error feedback. The parameters of the responder are allowed to vary - often using a gradient-descent approach - to minimize a cost function based on driver-responder distance. Although these methods have been shown to work well for systems with smooth variables we found that the sharp transitions in our firing rates, and the highly variable sensitivity of the dynamics to particular parameters as a function of position in state space, resulted in unstable and inaccurate parameter estimates.

We therefore adopted a multiple shooting parameter estimation method [11,40] that estimates divergence of short model forecasts from the UKF reconstructed trajectories over time windows long enough to explore the state space. This estimation step involves the minimization of a least-squared error, and can therefore be cast as a maximum-likelihood step. This is done in an iterative fashion to update parameter estimates by minimizing divergence of trajectories reconstructed using previous parameter estimates. Therefore this method becomes an expectation-maximization method, with all the associated global optimization implications [59,60].

Estimation of one or more parameters with any parameter estimation method will be inherently limited by the identifiability of the state space. Identifiability is a structural property of a model defined as the ability to identify a unique set of parameter values given error-free observations of the dynamics [61]. A comparable experimental or empirical version of identifiability has also been discussed by [62–65]. If some parameters are not structurally identifiable no parameter estimation method will prevail. Our experience and expectation is that the multiple shooting method will converge reasonably for combinations of identifiable parameters, but the convergence time increases with the number of parameters.

A key advantage of using the UKF for state reconstruction is allowance for uncertainties in the model and/or measurements. As noted, the Kalman filter is an iterative prediction-correction scheme. By altering the elements of the covariance inflation parameters CI and measurement uncertainty R , we can guide the Kalman filter to favor either the observations or model predictions. Higher values of CI downgrade the model-based forecasts during the correction step. We utilized this when developing the method for optimizing choice of CI values based on the *EOC*. For those variables that are poorly observed from others, we more heavily weight prediction over measurement; for those that yield poor reconstruction of other variables, we more heavily weight measurements.

Furthermore, as we showed in Fig. 6, inadequate models - which omit the full dynamics of certain variables - can be used to successfully assimilate experimental data and estimate unknown dynamics. In this example, we used a model that lacked any circadian dependencies to correctly estimate a 24-hour cycle and the mediated interaction with the SCN. Therefore our data assimilation framework can tolerate inadequate models and uncover dynamics outside the scope of the model's governing equations.

Several issues must be considered for assimilation of biological measurements. First, initial values for the filter parameters should be estimated off-line via the iterative reconstruct state/estimate

parameter approach. During this off-line learning process, non-arbitrary initial values for the covariance matrices as well as model parameters can be determined. Second, we will not have access to many of the state variables for validation. Previously, we developed a system that can automatically stage the behavioral state of a freely moving animal in real time [66], based on measurements of EEG and head acceleration with a resolution of a few seconds. This process can validate the UKF's predictions of sleep-state transitions. We can also use the scored behavioral state to infer the value of the Wake-active, NREM-active, and REM-active firing rate variables. As we have shown in Fig. 7 and Fig. 8, we can then use these inferred measurements to reconstruct hidden variables and estimate unknown parameters.

It is technically feasible to measure extracellular neurotransmitter concentrations using either dialysis or electrochemical sensors. Dialysis measurements do not have the temporal resolution to resolve REM dynamics, which occur on the order of one minute or less in the rodent, or the spatial specificity to localize dynamics to sleep-wake nuclei in the rodent brain. However these measurements could be used to track and validate slow systemic dynamics such as the circadian variations that modulate the sleep-wake nuclei. In contrast, off-the-shelf electrochemical sensor technology [67] allows for highly localized measures of neurochemicals such as ACh and 5-HT with sub-second temporal resolution and sub-mm spatial resolution. Such measurements can and should be used to establish and validate models used within the data assimilation framework. In addition, they can be used to identify the subset of measurements that can be accurately reconstructed from less costly observations. An appropriate cost-function for biological data assimilation would balance the degree of reconstruction inaccuracy against the cost of obtaining risky or hard-to-access measurements.

We also note that this framework could potentially be used to choose among model dynamics. Our parameter estimation methods rely on a minimization of prediction error. A similar metric or cost function could be utilized to differentiate between UKF-based tracking and prediction of system dynamics utilizing different models with such a filter framework.

In conclusion, we have presented a data assimilation framework for combining sparse measurements together with a relatively high-dimensional nonlinear computational model to reconstruct unmeasured variables, and have demonstrated its use in the context of a model of the sleep-wake regulatory system. We have demonstrated with simulation studies that once the tracked state approaches the true system state, it reliably reconstructs the unobserved system state (Fig. 2). We have introduced a metric for ranking relative observability for computational models (Fig. 3) that allows us not only to assess reconstruction performance based on choice of measurement, which can serve as a guide to which system variables to measure, but also provides a methodology for optimizing filter framework parameters such as the covariance inflation (Fig. 4). In addition, we have demonstrated a parameter estimation method (Fig. 5) that allows us to track non-stationary model parameters and accommodate slow dynamics not included in the UKF model such as circadian-dependent input from the SCN (Fig. 6). Finally, we have demonstrated that we can even use observed discretized SOV, which is not one of the model variables, to successfully reconstruct model state (Figs. 7–8).

These key features will aid in the transition of this framework to the experimental bench. Our long-term plan is to develop an observer-predictor system that will track and predict sleep-wake cycles as well as the underlying state of the neural cell groups and their neurochemical environment. Because these system dynamics are implicated in and interact with numerous neurological diseases

from epilepsy to schizophrenia, we anticipate that these tools will enable better understanding of the detailed interactions and provide for better, more targeted, therapies.

Acknowledgments

The authors would like to thank Cecilia Diniz Behn and Victoria Booth for extensive discussions on modeling of the sleep-wake regulatory network.

References

- Saper CB, Scammell TE, Lu J (2005) Hypothalamic regulation of sleep and circadian rhythms. *Nature* 437: 1257–63.
- Tamakawa Y, Karashima A, Koyoma Y, Katayama N, Nakao M (2006) A quartet neural system model orchestrating sleep and wakefulness mechanisms. *J Neurophysiol* 95: 2055–69.
- Diniz Behn CG, Brown EN, Scammell TE, Kopell NJ (2007) Mathematical model of network dynamics governing mouse sleep. *J Neurophysiol* 97: 3828–40.
- Diniz Behn CG, Booth V (2010) Simulating microinjection experiments in a novel model of the rat sleep-wake regulatory network. *J Neurophysiol* 103: 1937–1953.
- Phillips JK, Robinson PA (2007) A quantitative model of sleep-wake dynamics based on the physiology of the brainstem ascending arousal system. *J Biological Rhythms* 22: 167–179.
- Rempe MJ, Best J, Terman D (2010) A mathematical model of the sleep/wake cycle. *J Mathematical Biology* 60: 615–44.
- Kothare SV, Zarowski M (2011) Sleep and epilepsy: common bedfellows. *J Clinical Neurophysiol* 28: 101–2.
- Coogan AN, Thome J (2011) Chronotherapeutics and psychiatry: setting the clock to relieve the symptoms. *World J Biol Psychiatry* 12: 40–43.
- Chase RM, Pincus DB (2011) Sleep-related problems in children and adolescents with anxiety disorders. *Behav Sleep Med* 9: 224–36.
- Kalnay E (2003) Atmospheric modeling, data assimilation and predictability. UK: Cambridge University Press. 364 p.
- Voss HU, Timmer J (2004) Nonlinear dynamical system identification from uncertain and indirect measurements. *International Journal of Bifurcation and Chaos* 14: 1905–1933.
- Fleshner M, Booth V, Forger DB, Diniz Behn CG (2011) Circadian regulation of sleep-wake behaviour in nocturnal rats requires multiple signals from suprachiasmatic nucleus. *Philos Transact A Math Phys Eng Sci* 369: 3855–83.
- McCarley RW, Hobson JA (1975) Neuronal excitability modulation over the sleep cycle: a structural and mathematical model. *Science* 189: 58–60.
- Fuller PM, Saper CB, Lu J (2007) The pontine REM switch: past and present. *The Journal of Physiology* 584: 735–41.
- Ohno K, Sakurai T (2008) Orexin neuronal circuitry: Role in the regulation of sleep and wakefulness. *Frontiers in Neuroendocrinology* 29: 70–87.
- Porkka-Heiskanen T, Strecker RE, McCarley RW (2000) Brain site-specificity of extracellular adenosine concentration changes during sleep deprivation and spontaneous sleep: an in vivo microdialysis study. *Neuroscience* 99: 507–517.
- Huang ZL, Urade Y, Hayaishi O (2011) The role of adenosine in the regulation of sleep. *Curr Top Med Chem* 11: 1047–57.
- Borbély AA (1982) A two-process model of sleep regulation. *Human Neurobiology* 1: 195–204.
- Deurveilher S, Semba K (2005) Indirect projections from the suprachiasmatic nucleus to major arousal-promoting cell groups in rat: Implications for the circadian control of behavioural state. *Neuroscience* 130: 165–183.
- Vervey M, Amir S (2009) Food-entrainable circadian oscillators in the brain. *European Journal of Neuroscience* 30: 1650–1657.
- Hattar S, Liao HW, Takao M, Berson DM, Yau KW (2002) Melanopsin-containing retinal ganglion cells: Architecture, projections, and intrinsic photosensitivity. *Science* 295: 1065–1070.
- Krout KE, Kawano J, Mettenleiter TC, Loewy AD (2002) CNS inputs to the suprachiasmatic nucleus of the rat. *Neuroscience* 110: 73–92.
- Quigg M (2000) Circadian rhythms: interactions with seizures and epilepsy. *Epilepsy Research* 42: 43–55.
- Hofstra WA, de Weerd AW (2009) The circadian rhythm and its interaction with human epilepsy: A review of literature. *Sleep Medicine Reviews* 13: 413–420.
- Kalman RE (1960) A new approach to linear filtering and prediction problems. *Transactions of the ASME Journal of Basic Engineering* 82: 35–45.
- Julier SJ, Uhlmann JK (1997) A new extension of the Kalman filter to nonlinear systems. *P SPIE* 3068: 182–193.
- Schiff SJ (2012) *Neural Control Engineering*. MIT Press. 384 p.
- Simon D (2006) Optimal state estimation: Kalman, H [infinity] and nonlinear approaches. Hoboken, NJ: Wiley- Interscience. 552 p.
- Miyoshi T (2011) The Gaussian approach to adaptive covariance inflation and its implementation with the local ensemble transform. *Monthly Weather Review* 139: 1519–35.
- Anderson JL, Anderson SL (2011) A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *American Meteorological Society* 127: 2741–58.

We thank Bjoern Schelter for explaining the link between our parameter estimation method and expectation-maximization.

Author Contributions

Conceived and designed the experiments: MSS SJS BJG. Performed the experiments: MSS BJG. Analyzed the data: MSS SJS BJG. Wrote the paper: MSS SJS BJG.

- Mehra R (1970) On the identification of variances and adaptive Kalman filtering. *IEEE Transactions on Automatic Control* 15: 175–184.
- Mohamed AH, Schwarz (1999) Adaptive Kalman filtering for INS/GPS. *Journal of Geodesy* 73: 193–203.
- Wang J (2000) Stochastic modeling for real-time kinematic GPS/GLONASS position. *Navigation* 46: 297–305.
- Korniyenko OV, Sharawi MS, Aloï DN (2005) Neural network based approach for tuning KALMAN filter. In: *Electro Information Technology, 2005 IEEE International Conference on*. pp 1–5.
- Odelson BJ, Lutz A, Rawlings JB (2006) The autocovariance least-squares method for estimating covariances: application to model-based control of chemical reactors. *IEEE Transactions on Control Systems Technology* 14: 532–540.
- Akesson BM, Jrgensen JB, Poulsen NK, Jrgensen SB (2008) A generalized autocovariance least-squares method for Kalman filter tuning. *Journal of Process Control* 18: 769–779.
- Jatoh RK, Kumar TK (2009) Particle swarm optimization based tuning of unscented Kalman filter for bearings only tracking. In: *Advances in Recent Technologies in Communication and Computing, 2009. ARTCom '09. International Conference on*. pp 444–448.
- Rajamani MR, Rawlings JB (2009) Estimation of the disturbance structure from data using semidefinite programming and optimal weighting. *Automatica* 45: 142–148.
- Deng B, Wang J, Che Y (2009) A combined method to estimate parameters of neuron from a heavily noise-corrupted time series of active potential. *Chaos* 19: 015105.
- van Domselaar B, Hemkar P (1975) Nonlinear parameter estimation in initial value problems. Technical report. Mathematical Centre Amsterdam.
- Kalman RE (1960) On the general theory of control systems. *Proc IFAC 1st International Congress* 1: 481–92.
- Letellier C, Aguirre LA (2002) Investigating nonlinear dynamics from time series: The influence of symmetries and the choice of observables. *Chaos* 12: 549–558.
- Hu X, Nenov V, Bergsneider M, Glenn TC, Vespa P, et al. (2007) Estimation of hidden state variables of the intracranial system using constrained nonlinear Kalman filters. *IEEE transactions on Bio-Medical Engineering* 54: 597–610.
- Quach M, Brunel N, D'Alché-Buc F (2007) Estimating parameters and hidden variables in nonlinear state-space models based on ODEs for biological networks inference. *Bioinformatics* 23: 3209–16.
- Eberle C, Ament C (2010) The unscented Kalman filter estimates the plasma insulin from glucose measurement. *Bio Systems* 103: 67–72.
- Ullah G, Schiff SJ (2010) Assimilating seizure dynamics. *PLoS Computational Biology* 6: e1000776.
- Schiff SJ (2010) Towards model-based control of Parkinson's disease. *Philos Transact A Math Phys Eng Sci* 368: 2269–308.
- Freestone DR, Aram P, Dewar M, Scerri K, Grayden DB, et al. (2011) A data-driven framework for neural field modeling. *Neuro Image* 56: 1043–58.
- Toth B, Kostuk M, Meliza C, Margoliash D, Abarbanel H (2006) Dynamical estimation of neuron and network properties I: variational methods. *Biological Cybernetics* 105: 1–21.
- Letellier C, Aguirre L (2010) Interplay between synchronization, observability, and dynamics. *Phys Rev E* 82: 1–11.
- Letellier C, Aguirre LA, Maquet J (2005) Relation between observability and differential embeddings for nonlinear dynamics. *Phys Rev E* 71: 066213.
- Letellier C, Aguirre L (2009) Symbolic observability coefficients for univariate and multivariate analysis. *Phys Rev E* 79: 066210.
- Pecora LM, Carroll TL (1990) Synchronization in chaotic systems. *Phys Rev Lett* 64: 821–824.
- Parlitz U (1996) Estimating model parameters from time series by autosynchronization. *Phys Rev Lett* 76: 1232–1235.
- Maybhatte A, Amritkar RE (1999) Use of synchronization and adaptive control in parameter estimation from a time series. *Phys Rev E* 59: 284–293.
- Konnur R (2003) Synchronization-based approach for estimating all model parameters of chaotic systems. *Phys Rev E* 67: 027204.
- Huang D (2004) Synchronization-based estimation of all parameters of chaotic systems from time series. *Phys Rev E* 69: 6–9.
- Abarbanel HDI, Creveling DR, Jeanne JM (2008) Estimation of parameters in nonlinear systems using balanced synchronization. *Phys Rev E* 77: 016208.
- Vaida F (2005) Parameter convergence for EM and MM algorithms. *Statistica Sinica* 15: 831–840.

60. Wu JCF (1983) On the convergence properties of the EM algorithm. *The Annals of Statistics* 11: 95–103.
61. Jacquez JA, Greif P (1985) Numerical parameter identifiability and estimability: Integrating identifiability, estimability, and optimal sampling design. *Mathematical Biosciences* 77: 201–227.
62. Raue A, Kreutz C, Maiwald T, Bachmann J, Schilling M, et al. (2009) Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* 25: 1923–9.
63. Raue A, Becker V, Klingmüller U, Timmer J (2010) Identifiability and observability analysis for experimental design in nonlinear dynamical models. *Chaos* 20: 045105.
64. Raue A, Kreutz C, Maiwald T, Klingmüller U, Timmer J (2011) Addressing parameter identifiability by model-based experimentation. *IET Systems Biology* 5: 120.
65. Margaria G, Riccomagno E, White IJ (2004) Structural identifiability analysis of some highly structured families of statespace models using differential algebra. *J Math Biol* 49: 433–454.
66. Sunderam S, Chernyy N, Peixoto N, Mason JP, Weinstein SL, et al. (2007) Improved sleep-wake and behavior discrimination using MEMS accelerometers. *Journal of Neuroscience Methods* 163: 373–83.
67. Michael AC, Borland LM, editors (2007) *Electrochemical Methods for Neuroscience*, CRC Press, chapter 19.