A data-driven phase and isostable reduced modeling framework for oscillatory dynamical systems

Dan Wilson

View Online   Export Citation   CrossMark

A data-driven phase and isostable reduced modeling framework for oscillatory dynamical systems

View Online   Export Citation   CrossMark

Dan Wilson[a] 🆔

AFFILIATIONS

Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, Tennessee 37996, USA

[a]**Author to whom correspondence should be addressed:** dwilso81@utk.edu

ABSTRACT

Phase-amplitude reduction is of growing interest as a strategy for the reduction and analysis of oscillatory dynamical systems. Augmentation of the widely studied phase reduction with amplitude coordinates can be used to characterize transient behavior in directions transverse to a limit cycle to give a richer description of the dynamical behavior. Various definitions for amplitude coordinates have been suggested, but none are particularly well suited for implementation in experimental systems where output recordings are readily available but the underlying equations are typically unknown. In this work, a reduction framework is developed for inferring a phase-amplitude reduced model using only the observed model output from an arbitrarily high-dimensional system. This framework employs a proper orthogonal reduction strategy to identify important features of the transient decay of solutions to the limit cycle. These features are explicitly related to previously developed phase and isostable coordinates and used to define so-called data-driven phase and isostable coordinates that are valid in the entire basin of attraction of a limit cycle. The utility of this reduction strategy is illustrated in examples related to neural physiology and is used to implement an optimal control strategy that would otherwise be computationally intractable. The proposed data-driven phase and isostable coordinate system and associated reduced modeling framework represent a useful tool for the study of nonlinear dynamical systems in situations where the underlying dynamical equations are unknown and in particularly high-dimensional or complicated numerical systems for which standard phase-amplitude reduction techniques are not computationally feasible.

Published under license by AIP Publishing. https://doi.org/10.1063/1.5126122

Many methods exist for studying the perturbed dynamics of a periodic oscillator in a reduced coordinate framework, but few are well suited for implementation in experimental settings where the underlying model dynamics are typically unknown. In this work, a reduction framework is proposed for inferring phase-amplitude reduced equations using only observed model output. In contrast to comparable reduction strategies, the proposed methodology yields a reduced order model that can provide accurate predictions for observed model behaviors in response to exogenous inputs. Additionally, this reduction framework is amenable for control applications requiring state feedback, allowing for the implementation of sophisticated control strategies that would otherwise be infeasible. It is envisioned that the data-driven phase and isostable reduction methodology will find use in experimental systems (where the underlying dynamics are unknown) and in particularly high-dimensional or complicated numerical systems for which

standard phase-amplitude reduction techniques are not computationally tractable.

I. INTRODUCTION

There are numerous difficulties that arise when working with high-dimensional dynamical systems. For example, model identification becomes difficult as the risk of overfitting increases with the dimension of the model considered, optimal control frameworks often become computationally intractable in high-dimensional settings, and dynamical mechanisms responsible for observed behaviors can become obscured by the size and complexity of high order systems. The aforementioned examples are facets of the "curse of dimensionality" coined by Richard Bellman,[5] a colorful interpretation of the fact that the amount of data needed to accurately

represent and characterize a dynamical system grows exponentially with its underlying dimensionality.

In an effort to mitigate some of these issues, model reduction has become an increasingly important first step in the mathematical analysis of high-dimensional dynamical systems and subsequent implementation of active control strategies. There are many different available strategies for model reduction. Proper orthogonal decomposition (POD) is a commonly used reduction technique that identifies an optimal set of modes (in an L_2 sense) to represent a given data set.[6,28,55,62] This method is particularly well suited for extraction of dynamical features from linear systems but is often inadequate in situations where nonlinear terms dominate. Koopman analysis has become increasingly popular in recent years, a strategy which can be used to represent the observables of a nonlinear dynamical system with a linear but infinite-dimensional operator.[10,42] In practice, the utility of Koopman based methods hinges on the ability to find an adequate set of observable functions in order to find a finite-dimensional approximation to the infinite-dimensional Koopman operator. Dynamic mode decomposition (DMD)[59] and extended DMD[65] are often used to fit a set of observable functions to data. Much like the POD strategy, the DMD algorithms can be difficult to implement in systems with dominant nonlinearities.

Other model reduction frameworks are better suited to nonlinear systems. In some cases, fast-slow analysis[7,28] can be used to partition dynamical systems into separate subsystems with different time scales. Averaging techniques are often used to characterize the influence of periodic perturbations that are sufficiently weak,[58] slow,[52] or rapid[53] and are often used to study bifurcations in nonlinear dynamical systems. Inertial manifolds,[12,19] can be used to identify lower dimensional manifolds to which solutions exponentially collapse but are often hard to identify. Unlike POD and Koopman analysis, these nonlinear reduction techniques require an underlying description of the model dynamics (i.e., they cannot be implemented solely from experimental data) and are difficult to employ in high-dimensional systems.

Few reduction frameworks exist that are both well suited to application in nonlinear systems and can be implemented using only recorded model output in the absence of an underlying model. Phase reduction[18,30,72] is one such framework that has shown promise for these applications. This methodology can be used to reduce a general oscillatory model of the form

$$\dot{x} = F(x) + U(t), \tag{1}$$

where $x \in \mathbb{R}^N$ is the state vector, F represents the nominal dynamics, and U is an external perturbation. Letting $x^\gamma(t)$ denote a T-periodic limit cycle solution of (1), solutions can be represented using a scalar equation of the form

$$\dot{\theta} = \omega + Z(\theta)^T U(t), \tag{2}$$

where $\theta \in [0, 2\pi)$, $\omega = 2\pi/T$, and $Z(\theta)$ is the phase response curve (PRC), which gives the gradient of the phase coordinate evaluated on the periodic orbit. Equation (2) is a linear approximation of the perturbed behavior near $x^\gamma(t)$, but the phase can readily be extended to the basin of attraction of the limit cycle using the notion of isochrons.[27,72] When $U = 0$, for any initial condition $a(0) \in x^\gamma(t)$,

the isochron associated with $a(0)$ is defined to be the set of all $b(0)$ such that

$$\lim_{t \to \infty} ||a(t) - b(t)|| = 0, \tag{3}$$

where $|| \cdot ||$ can be any vector norm. Using isochrons in conjunction with isostable coordinates, which represent level sets of initial conditions that approach the periodic orbit together,[40,68,71] a phase-amplitude reduction can be performed that explicitly accounts for dominant nonlinearities in the underlying model.[66,70]

Phase reduction using (2) to describe the reduced dynamics has been successfully employed in many experimental applications,[2,8,46,47,75–77] (i.e., where the full model equations are unknown). Phase reduction is particularly useful in situations where the Floquet exponents are large in magnitude relative to the magnitude of applied perturbations. As the magnitude of perturbations becomes larger, the dynamics of the amplitude coordinates may need to be considered. Application of the phase-amplitude reduction framework as suggested in Refs. 68 and 71 has been difficult to implement in situations where model equations are unavailable because of the challenge associated with finding a suitable basis of exponentially decaying mode shapes from model output. This problem was explored in Ref. 67, but general frameworks for finding a reduced set of phase-amplitude equations from model output have yet to be developed.

In this work, a new coordinate system is proposed for determining a phase-amplitude reduced model from system observables. The organization of this paper is as follows: Sec. II provides necessary background information on isostable coordinates that have been used previously in the development of phase-amplitude reduction strategies for nonlinear oscillators. Section III suggests and investigates a phase-amplitude reduced modeling framework using so-called data-driven phase and isostable coordinates. In contrast to previously suggested strategies for identifying phase and isostable reduced equations directly from model output,[69,70] the strategy presented in Sec. IV can be used to identify data-driven phase and isostable reduced equations with an arbitrary number of isostable coordinates and gives accurate predictions for the system observables. Section V provides an illustrative example of this framework implemented on a population of coupled neurons, and Sec. VI illustrates an optimal control strategy made possible by the data-driven phase and isostable reduction framework with applications to circadian physiology and recovery from jet lag. Section VII gives concluding remarks.

II. BACKGROUND ON ISOSTABLE COORDINATE SYSTEMS

Phase reduction (2) is a well-established and tremendously useful tool for the analysis of weakly perturbed oscillatory systems, but it ultimately requires perturbations $U(t)$ to be small relative to the Floquet exponents[37] so that the state remains close to the periodic orbit. In order to better understand behaviors in response to stronger perturbations, coordinates representing the transient decay in directions transverse to the periodic orbit must be considered.

Various coordinate frameworks have been used to represent the behavior transverse to the periodic orbit.[13,21,39,61,64,71] In this work,

the isostable coordinate framework will be used, which characterizes the infinite time decay of transient solutions. To define isostable coordinates, as in Refs. 61 and 68, consider the $\theta = 0$ isochron, Γ_0, and let $U = 0$. By definition, any initial condition that starts on Γ_0 returns to this surface at time T allowing for the construction of a Poincaré map,

$$Q : \Gamma_0 \to \Gamma_0;$$
$$x \to \nu(T, x), \tag{4}$$

where $\nu(t, x)$ represents the unperturbed flow of (1). The mapping (4) has a fixed point, x_0, located at the intersection of Γ_0 and the periodic orbit. Consequently, in a neighborhood of x_0, $\nu(T, x)$ can be approximated by local linearization as

$$\nu(T, x) = x_0 + J_\nu(x - x_0), \tag{5}$$

where J_ν is the Jacobian of $\nu(T, x)$ evaluated at x_0. Supposing that J_ν is diagonalizable with eigenvalues λ_j for $j = 1, \ldots, N$ and corresponding left eigenvectors w_j, one can define associated isostable coordinates for the components of the solution associated with some of the largest magnitude eigenvalues λ_j as

$$\psi_j(x) = \lim_{k \to \infty} \left[w_j^T (\nu(t_\Gamma^k, x) - x_0) \exp(-\kappa_k t_\Gamma^k) \right], \tag{6}$$

where t_Γ^k is the kth return time to Γ_0 under the action of the flow and $\kappa_k = \log(\lambda_k)/T$ is often referred to as a Floquet exponent.[32] Intuitively in the definition above, $\exp(-\kappa_k t_\Gamma^k)$ increases by a factor of $1/\lambda_j$ after every T time units. The left eigenvector w_j selects for a component with a matching decay rate in the limit that k approaches infinity. Related definitions can be used to define isostable coordinates in the neighborhood of the limit cycle.[71]

As illustrated in Ref. 71, isostable coordinates can be used alongside (2) to characterize the transient behavior of perturbations transverse to the limit cycle according to

$$\dot{\theta} = \omega + Z(\theta)^T U(t),$$
$$\dot{\psi}_j = \kappa_j \psi_j + I_j(\theta)^T U(t), \tag{7}$$
$$j = 1, \ldots, N - 1,$$

where $I_j(\theta)$ is the isostable response curve (analogous to the phase response curve) associated with the isostable coordinate ψ_j. Efficient numerical methods based on the adjoint method[9] have been developed for numerically computing the phase response curves[9] and isostable response curves.[71] To a first order approximation, the phase dynamics are uncoupled from the isostable dynamics. In general, this is not true for higher order accuracy phase-amplitude reductions; these considerations were explored in detail in Refs. 66 and 68.

III. DATA-DRIVEN PHASE-AMPLITUDE COORDINATES

While strategies for the computation of the phase-amplitude coordinate reduction (7) are fairly straightforward when the right hand side of the underlying model Eq. (1) is known, it is significantly more difficult to implement this reduction framework in experimental situations where the model equations are unknown. When the model equations are unknown, the standard phase reduction (2) can be found by using the "direct method,"[22,30,47] to estimate

$Z(\theta)$ by inferring the phase change in response to a set of exogenous perturbations given at many different phases. These methods have been applied successfully to gain understanding about the timing of oscillations in various experimental applications.[29,36,46] Conversely, there are no well-established frameworks that can be used to determine phase-amplitude reduced equations from experimental data. Previous authors have suggested that a model can be built by considering the timing of a distinct periodic event, for example, the upstroke of a neural action potential. This framework has led to the notion of higher order or residual PRCs,[13,21,48,69] which characterize the influence of a perturbation not only on the next cycle, but also on all subsequent cycles. In practice, it can be difficult to infer higher order phase response curves in the presence of noise, or when there are multiple slowly decaying components of the transient solution. Other methods based on the isostable framework have been suggested,[69,70] but are limited to situations where one representative isostable coordinate can be used to characterize the rate of decay of solutions toward the limit cycle.

In the derivation to follow, a procedure is developed to determine phase-amplitude reduced coordinate framework for a general oscillatory model with a single observable. Unlike other recently developed strategies, this method can be used to characterize the transient decay of multiple isostable coordinates. Additionally, the numerical examples to follow will illustrate that the resulting reduced equations are accurate enough so that they can be used to implement sophisticated control strategies, even when the reduced equations are obtained in the presence noise. The reduction strategy will be developed for use in dynamical systems of the form

$$\dot{x} = F(x, \epsilon U(t)) + \epsilon \eta(t),$$
$$y = G(x), \tag{8}$$

where $x \in \mathbb{R}^N$ is the state, $y \in \mathbb{R}$ is a single system observable, $\eta(t) \in \mathbb{R}^N$ is a vector with each component η_i being an independent Gaussian white noise process with intensity D_i, F gives the nominal dynamics, U is an external perturbation, G maps the state to the output, and $0 < \epsilon \ll 1$. Correlations between the noise processes, colored noise processes, and influences of measurement noise are not considered here.

Suppose in the absence of noise and when $U(t) = 0$, Eq. (8) admits a stable, T-periodic orbit $x^\gamma(t)$ with a corresponding output $y^\gamma(t) = G(x^\gamma(t))$. Compared with (1), Eq. (8) does not explicitly assume additive control and also assumes that individual state variables are not directly measurable. Near the periodic orbit to leading order ϵ, linearization of (8) yields

$$\Delta \dot{x} = A(t) \Delta x + \epsilon B(t) + \epsilon \eta(t) + \mathcal{O}(\epsilon^2),$$
$$\Delta y = C(t) \Delta x + \mathcal{O}(\epsilon^2), \tag{9}$$

where $\Delta x(t) \equiv x(t) - x^\gamma(t)$, $\Delta y \equiv y(t) - y^\gamma(t)$, $A(t) = F_x(x^\gamma(t), 0)$, $B(t) = F_U(x^\gamma(t), 0)U(t)$, $C(t) = G_x(x^\gamma(t))$. The notation $Q_R(x, U)$ is used to denote $\partial Q/\partial R$; all partial derivatives are evaluated at $x = x^\gamma(t)$ and $U = 0$. When both $U(t)$ and noise are absent, Eq. (9) is T-periodic. Supposing that the fundamental matrix associated with this linear time varying system is diagonalizable, solutions near the

periodic orbit can be determined from Floquet theory,[26,32]

$$\Delta x(t) = \sum_{j=1}^{N} c_j e^{\kappa_j t} p_j(t), \tag{10}$$

where $p_j(t)$ are T-periodic functions with constants c_j chosen so that initial conditions are satisfied, and κ_j are Floquet exponents. By convention, Floquet exponents will be sorted so that $\kappa_N = 0$ is the Floquet exponent corresponding to the unity Floquet multiplier of the periodic orbit. Using the isostable coordinate framework (7) along with the Floquet basis of solutions (10), as shown in Refs. 66 and 68, it is possible to represent solutions to (9) using a set of equations of the form,

$$\dot{\theta} = \omega + \epsilon Z^T(\theta)(B(t) + \eta(t)),$$
$$\dot{\psi}_j = \kappa_j \psi_j + \epsilon I_j^T(\theta)(B(t) + \eta(t)), \quad j = 1, \ldots N-1, \tag{11}$$
$$\Delta x = x(\theta + \Delta\theta, \psi_1, \ldots, \psi_{N-1}) - x^\gamma(\theta)$$
$$= \Delta\theta p_N^\theta(\theta) + \sum_{k=1}^{N-1} \left(\psi_k p_k^\theta(\theta)\right),$$

where $\theta + \Delta\theta$ is the phase of oscillation with $\Delta\theta = \mathcal{O}(\epsilon)$, $\omega = 2\pi/T$ is the natural frequency, ψ_j is the jth isostable coordinate, and $p_k^\theta(\omega t) = p_k(t)$. $x^\gamma(0)$ will be defined to correspond to $\theta = 0$ on the periodic orbit. In many situations, some of the Floquet exponents are negative and large in magnitude so that the corresponding isostable coordinate decays rapidly to zero in response to perturbation. In these instances, these isostable coordinates are often taken to be zero (e.g., as done in Refs. 45 and 70) providing a reduction to the overall dimensionality as compared to (8). As a final note, the noise intensity is assumed to be small enough so that the terms from the Ito correction[23] are negligible and can be ignored in (11).

In situations where the full equation (8) is known, it is straightforward to calculate the transformation from (8) to (11) through the numerical computation of $Z(\theta)$[9] and $I_j(\theta)$.[71] Currently, there are no reliable methods for inferring each $I_j(\theta)$ and κ_j when multiple isostable coordinates are necessary to adequately capture the transient behavior. Additionally, there are no existing strategies to identify the associated functions p_j^θ that relate the reduced coordinates to the model output. In the derivation to follow, proper orthogonal decomposition is used in concert with Floquet theory to achieve these goals from a noisy data set. This strategy does not require any knowledge of the underlying dynamical equations.

A. Preliminary identification of modes using proper orthogonal decomposition

To begin, consider the phase dynamics (11) in the absence of noise and taking $U(t) = 0$. For this simple system with an initial condition $\Psi(t_0) = [\psi_1(t_0) \ \ldots \ \psi_{N-1}(t_0) \ \theta(t_0)]^T$, it is immediately apparent that $\theta(t) = \theta(t_0) + \omega(t - t_0)$ and $\psi_j(t) = \psi_j(t_0) \exp(\kappa_j (t - t_0))$. Further assuming that $\theta(t_0)$ and each $\psi_j(t_0)$ are order ϵ terms and substituting Δx from (11) into the output equation of (9),

to leading order ϵ, one finds

$$y_i(t) - y^\gamma(t) = C(\theta(t))\Delta x$$
$$= \sum_{k=1}^{N-1} \left[C^\theta(\omega t) p_k^\theta(\omega t) \psi_k(t) \right] + C(\omega t) p_N^\theta(\omega t) \Delta\theta(t)$$
$$= \sum_{k=1}^{N-1} \left[C^\theta(\omega t) p_k^\theta(\omega t) \psi_k(0) \exp(\kappa_k t) \right]$$
$$+ C(\omega t) p_N^\theta(\omega t) \Delta\theta(0), \tag{12}$$

where $C^\theta(\omega t) = C(t)$ and t_0 is assumed to be 0 for simplicity of notation.

With an underlying description of the output dynamics from (12), it is possible to devise a strategy to determine a reduced order model for the observed output. To do so, consider a collection of α sets of data signals $y_1, y_2, \ldots, y_\alpha \in \mathbb{R}^\beta$ arranged as column vectors so that their first components correspond to an initial condition for which $\theta = \mathcal{O}(\epsilon)$ and the remaining $\beta - 1$ elements are measurements taken $\Delta t = T/(\beta - 1)$ time units apart. In other words, each y_i represents a recording over one period which starts and ends at $\theta \approx 0$. Define the matrix $Y \in \mathbb{R}^{\beta,\alpha}$

$$Y \equiv \begin{bmatrix} y_1 - \gamma & \ldots & y_\alpha - \gamma \end{bmatrix}, \tag{13}$$

where $\gamma \in \mathbb{R}^\beta$ is a discretized version of $y^\gamma(t)$ with the kth element of γ corresponding to $y^\gamma(\Delta t(k-1))$. It will be assumed that the columns of Y contain a good representation of the possible outputs that can be observed from the underlying system (8).

The goal of this analysis is to find a reduced order dynamical model similar to (11) to represent the information contained in the matrix Y. As a preliminary step, the method of proper orthogonal decomposition (POD) will be employed.[6,28,55,62] This is a well-established framework that can be used to decompose a set of data snapshots into a minimal number of representative modes. It will be shown that the dynamical behavior of the resulting reduced order model is closely related to the phase and isostable coordinate dynamics. To begin, POD can be used to identify representative mode shapes from (13) by solving an eigenvalue problem. Specifically, let v_j and ζ_j denote the eigenvectors and eigenvalues, respectively, of the matrix $Y^T Y$, with eigenvalues ordered such that $\zeta_j > \zeta_{j+1}$. Individual POD modes, ϕ_j, can be found using the method of snapshots according to[28]

$$\phi_j = \frac{1}{\sqrt{\zeta_j}} Y v_j, \tag{14}$$

where the resulting basis of POD modes is orthogonal. Additionally, the POD basis is arranged so that modes associated with larger values of ζ_j capture more of the temporal fluctuations in the data set. A total of m modes will be chosen such that $\sum_{j=1}^{m} \zeta_j / \sum_{j=1}^{\alpha} \zeta_j \approx 1$, resulting in a reduced set of modes that can be used to accurately reconstruct each element of Y. Define $\Phi \in \mathbb{R}^{\beta \times m}$ such that $\Phi = [\phi_1 \ \ldots \ \phi_m]$. Any column of the matrix Y can be projected onto this new POD basis according to

$$y_i - \gamma = \sum_{j=1}^{m} \left(\phi_j \xi_j^i \right) + q_i = \Phi \mu_i + q_i, \tag{15}$$

where q_i is some small residual that the POD basis does not capture and $\mu_i = \begin{bmatrix} \xi_1^i & \dots & \xi_m^i \end{bmatrix}^T$ is a vector of POD coefficients associated with y_i that can be found according to

$$\mu_i = \Phi^T(y_i - \gamma) \qquad (16)$$

as the result of orthogonality of the POD basis. It will be assumed that enough POD modes are included so that $q_i = \mathcal{O}(\epsilon_1)$ for all i with $0 < \epsilon_1 \ll 1$. Using (12), one can also write

$$y_i - \gamma = P\Psi_i, \qquad (17)$$

where $P \in \mathbb{R}^{\beta, N}$ with ith row and jth column equal to $C^\theta(\omega\Delta t (i-1))p_j^\theta(\omega\Delta t(i-1))\exp(\kappa_i\Delta t(i-1))$ and $\Psi_i = [\psi_1(t_i) \dots \psi_{N-1}(t_i) \ \Delta\theta(t_i)]^T$ are the initial conditions of the transformed phase and isostable coordinates with t_i being the time associated with the first entry of y_i. It will be assumed that the columns of P are linearly independent in the derivation to follow. For this condition to hold, it is necessary that $\beta \geq N$ and that for all j, $p_j^\theta(\theta)$ is not in the null space of $C^\theta(\theta)$ for at least some values of θ. Setting (15) and (17) equal to each other yields

$$\Phi\mu_i + q_i = P\Psi_i, \qquad (18)$$

which can be manipulated to provide the following relationships:

$$\mu_i = \Phi^T P\Psi_i - \Phi^T q_i, \qquad (19)$$

$$\Psi_i = P^\dagger \Phi\mu_i + P^\dagger q_i, \qquad (20)$$

where † denotes the pseudoinverse. Above, Eq. (19) follows because the POD basis is orthonormal while (20) follows from the assumption that the columns of P are linearly independent so that $P^\dagger P$ yields the identity matrix.

Now that relationships between μ_i and Ψ_i have been established, consider the temporal evolution of these reduced coordinates. Suppose two successive measurements y_i and y_i^+ are taken over two time intervals $[t_i, t_i + T]$ and $[t_i + T, t_i + 2T]$, respectively, chosen so that $\theta(t_i) \approx 0$. Let the set of reduced coordinates Ψ_i and Ψ_i^+ correspond to the phase and isostable coordinates associated with the outputs $y_i - \gamma$ and $y_i^+ - \gamma$, respectively. In the absence of perturbation and noise, the phase and isostable coordinates change over one period according to $\Psi_i^+ = \Lambda\Psi_i$ where $\Lambda \in \mathbb{R}^{N,N} = \text{diag}(\exp(\kappa_1 T), \dots, \exp(\kappa_{N-1} T), 0)$. Using this relationship along with (19) and (20), one finds

$$\begin{aligned}
\mu_i^+ &= \Phi^T P\Psi_i^+ + \Phi^T q_i^+ \\
&= \Phi^T P\Lambda\Psi_i + \Phi^T q_i^+ \\
&= \Phi^T P\Lambda \left[P^\dagger \Phi\mu_i + P^\dagger q_i \right] + \Phi^T q_i^+ \\
&= A_\mu \mu_i + \mathcal{O}(\epsilon_1), \qquad (21)
\end{aligned}$$

where $A_\mu \equiv \Phi^T P\Lambda P^\dagger \Phi$ and the final line is obtained by recalling that q_i is assumed to be an $\mathcal{O}(\epsilon_1)$ term. Thus, the model

$$\begin{aligned}
\mu_i^+ &= A_\mu \mu_i + \mathcal{O}(\epsilon_1), \\
y_i &= \gamma + \Phi\mu_i + \mathcal{O}(\epsilon_1),
\end{aligned} \qquad (22)$$

provides an approximation for the underlying dynamical behavior for the outputs of initial conditions that start near $\theta = 0$. Finally,

standard techniques can be used to write $A_\mu = VDV^{-1}$, where D is in the Jordan normal form. The coordinate transformation $\Sigma_i = V^{-1}\mu_i$ yields

$$\Sigma_i^+ = D\Sigma_i + \mathcal{O}(\epsilon_1), \qquad (23)$$

where the eigenvalues of A_μ appear on the diagonal of D. Noting the relationships $\Sigma_i^+ = V^{-1}\mu_i^+ = V^{-1}(\Phi^T P\Psi_i^+ - \Phi^T q_i^+)$, and the fact that $\Psi_i^+ = \Lambda\Psi_i$, it is also possible to write

$$\Sigma_i^+ = V^{-1}\Phi^T P\Lambda\Psi_i + \mathcal{O}(\epsilon_1). \qquad (24)$$

Using the notation Σ_i^{n+} and Ψ_i^{n+} to correspond to the coordinates Σ and Ψ resulting from the measurement y_i^{n+} taken over the time interval $[t_i + nT, t_i + (n+1)T]$, one can write

$$\Sigma_i^{n+} = D^n\Sigma_i + \mathcal{O}(\epsilon_1) = V^{-1}\Phi^T P\Lambda^n\Psi_i + \mathcal{O}(\epsilon_1). \qquad (25)$$

Multiplying (25) by e_j^T, where e_j^T is the jth element of the standard unit basis, one can write

$$\begin{aligned}
e_j^T D^n\Sigma_i &= \left(e_j^T V^{-1}\Phi^T P\right)\Lambda^n\Psi_i + \mathcal{O}(\epsilon_1) \\
&= Q_j\Lambda^n\Psi_i + \mathcal{O}(\epsilon_1), \qquad (26)
\end{aligned}$$

where $Q_j^T \in \mathbb{R}^{1\times n} = e_j^T V^{-1}\Phi^T P$. Note that Eq. (26) is valid for any initial conditions Ψ_i. For the moment, initial conditions with magnitudes that are sufficiently larger than ϵ_1 will be considered so that the $\mathcal{O}(\epsilon_1)$ terms from (26) can be neglected. Recall that Λ is diagonal so that the right hand side of (26) is a sum of exponentially decaying functions. This implies that D must also be diagonal so that, neglecting $\mathcal{O}(\epsilon_1)$ terms, (26) can be rewritten as

$$d_j^n e_j^T\Sigma_i = Q_j\Lambda^n\Psi_i, \qquad (27)$$

where d_j is the jth element on the diagonal of D. In order for (27) to hold, for any isostable coordinate ψ_k with associated Floquet multiplier $\lambda_k \neq d_j$, the corresponding kth entry of Q_j must be zero. This implies that $e_j^T\Sigma_i$ can be written as a linear combination of all isostable coordinates that have corresponding Floquet multipliers identical to d_j. This relationship will become important momentarily.

For a given POD basis, the matrix A_μ [as defined below (21)] determines the dynamics of the POD coefficients. This is relatively easy to calculate when the nominal dynamics $F(x, 0)$ are known (allowing Λ and P to be calculated straightforwardly using standard methods from Floquet theory). When the nominal dynamics are not known, however, the A_μ matrix must be inferred from data, a consideration that will be discussed in Sec. IV.

B. Definition of data-driven phase and isostable coordinates

The relationship (23), which is valid near the limit cycle, can be used to define a new set of coordinates that have similar properties to isostable coordinates in the basin of attraction of the limit cycle. These will be called data-driven isostable coordinates. To do so, first suppose that A_μ has already been found using methods described in Sec. III A. Consider some hypothetical signal $y = \begin{bmatrix} y(t_0) & y(t_0 + \Delta t) & \dots \end{bmatrix}^T$, which has been measured in the absence of noise with $\Delta t = T/(\beta - 1)$. Suppose that at time

$t_0 + a\Delta t$, it is detected that $\theta \approx 0$ (for example, by detecting the crossing of a voltage threshold during the upstroke of an action potential for a periodically firing neuron). Using these data, one can define the following vectors:

$$s_k \in \mathbb{R}^\beta \equiv \begin{bmatrix} y\left(t_0 + a\Delta t + (\beta - 1)k\Delta t\right) \\ y\left(t_0 + (a+1)\Delta t + (\beta - 1)k\Delta t\right) \\ \vdots \\ y\left(t_0 + (\beta - 1)\Delta t + (\beta - 1)k\Delta t\right) \end{bmatrix}, \quad (28)$$

for $k = 0, 1, \ldots$. Recall that A_μ can be written as $A_\mu = VDV^{-1}$, where D is a diagonal matrix and let d_j be the jth element on the diagonal of D (i.e., the jth eigenvalue of A_μ) with an associated left eigenvector w_j. In Sec. III A, it was shown that each d_j is a Floquet multiplier of the underlying periodic orbit. The unity Floquet multiplier of the periodic orbit will receive special attention and D will be sorted so that this eigenvalue appears in the mth row and mth column. For any $d_j \neq 1$, an associated data-driven isostable coordinate can be defined according to

$$\psi_j^D = w_j^T \Phi^T (s_k - \gamma) \exp\left(-\frac{\log(d_j)(a\Delta t + kT)}{T}\right), \quad (29)$$

for $k \in \mathbb{N}$ is chosen appropriately so that s_k is close enough to the limit cycle so that the linearization (9) is valid. The definition of data-driven isostable coordinates (29) is closely related to (6) and has a similar dynamical behavior. Intuitively, $w_j^T \Phi^T$ behaves like a left eigenvector that takes $(s_k - \gamma)$ and returns the jth element of Σ_k. Additionally, for the unity Floquet multiplier, $d_m = 1$, a corresponding a data-driven phase coordinate, $\theta^D \in [0, 2\pi)$, can be defined as

$$\theta^D = 2\pi \left(1 - \frac{a\Delta t}{T}\right) + \frac{1}{c} w_m^T \Phi^T (s_k - \gamma), \quad (30)$$

where k is chosen to be the same value used in the definition of (29) and c is a constant to be determined shortly.

In the absence of noise and external input, it will be shown that the dynamical behavior of data-driven isostable coordinates is well approximated by $\dot\psi_j^D = \log(d_i)/T$ when $|\psi_j^D| \gg \epsilon_1$. Likewise, by choosing c appropriately in (30), data-driven phase coordinates have dynamics that are well approximated by $\dot\theta^D = \omega$, which is identical to the standard definition of phase coordinates. To show this, recalling the relationships $\Sigma_i = V^{-1}\mu_i$ and $\mu_i = \Phi^T(s_i - \gamma)$, one can rewrite (29) as

$$\psi_j^D = e_j^T \Sigma_k \exp\left(-\frac{\log(d_j)(a\Delta t + kT)}{T}\right). \quad (31)$$

Taking the time derivative of (31) by noticing that $e_j^T \Sigma_k$ is the only nonconstant term, one finds

$$\dot\psi_j^D = \frac{d}{dt}\left(e_j^T \Sigma_k\right) \exp\left(-\frac{\log(d_j)(a\Delta t + kT)}{T}\right). \quad (32)$$

Recall that as a consequence of (27) [which neglects $\mathcal{O}(\epsilon_1)$ terms] near the periodic orbit, one can show that $e_j^T \Sigma_k$ can be written as a linear combination of the isostable coordinates that have associated Floquet multipliers d_j [and corresponding Floquet exponent $\log(d_j)/T$]. Because all isostable coordinates decay at a rate

governed by their Floquet exponents, this implies that $\frac{d}{dt}\left(e_j^T \Sigma_k\right)$ $= \frac{\log(d_j)}{T}\left(e_j^T \Sigma_k\right) + \mathcal{O}(\epsilon_1)$. Substituting this into (32) yields

$$\dot\psi_j^D = \frac{\log(d_j)}{T}\psi_j^D + \mathcal{O}(\epsilon_1). \quad (33)$$

A similar strategy can be used to find $\dot\theta^D$. Substituting $e_m^T \Sigma_k = w_m^T \Phi^T(s_k - \gamma)$ into (30) and taking the time derivative yields

$$\dot\theta^D = \frac{1}{c}\frac{d}{dt}\left(e_m^T \Sigma_k\right). \quad (34)$$

Again as a consequence of (27), near the periodic orbit, $e_m^T \Sigma_k$ in the above equation is a constant multiple of $\Delta\theta$ plus some $\mathcal{O}(\epsilon_1)$ residual, i.e., $e_m^T \Sigma_k = c_\theta \Delta\theta + \mathcal{O}(\epsilon_1)$, where c_θ is a constant. Choosing c from (30) to be identical to c_θ and noting that $\theta = \Delta\theta$ when $\theta \approx 0$, one finds

$$\dot\theta^D = \omega + \mathcal{O}(\epsilon_1) \quad (35)$$

in the absence of perturbation.

Finally, in the event that none of the values of d_j are equal to zero, the data-driven phase and isostable coordinates can be used to predict the future behavior of a system. To see this, consider an initial condition taken at $\theta \approx 0$ so that $a\Delta t$ from (29) and (30) can be chosen to be 0. Also, suppose that ψ_k are small so that the state is close to the periodic orbit. Because the state is close to the periodic orbit, the data-driven phase and isostable coordinates can be approximated as $\theta^D \approx \frac{1}{c} w_m^T \Phi^T(s_1 - \gamma)$ and $\psi_j^D \approx \frac{1}{d_j} w_j^T \Phi^T(s_1 - \gamma)$. Written in matrix form, this gives

$$\Psi^D = L^{-1}V^{-1}\Phi^T(s_1 - \gamma), \quad (36)$$

where $L = \mathrm{diag}(d_1, d_2, \ldots, d_{m-1}, 1/c)$ is an invertible matrix because each element on the diagonal is nonzero and $\Psi^D \equiv \begin{bmatrix} \psi_1^D & \ldots & \psi_{m-1}^D \\ \theta^D \end{bmatrix}^T$. Recalling the relationships from (15) and (16), Eq. (36) can be manipulated to yield

$$VL\Psi^D = \mu_1,$$
$$\Phi VL\Psi^D + \mathcal{O}(\epsilon_1) = s_1 - \gamma, \quad (37)$$

giving a direct relationship between the isostable coordinates and the expected model output.

IV. PRACTICAL IMPLEMENTATION OF THE DATA-DRIVEN PHASE AND ISOSTABLE REDUCTION STRATEGY

In Sec. III, it is assumed that the matrix A_μ from (22) (which determines the decay rates and left eigenvectors for the data-driven phase and isostable coordinates) can be calculated from direct computation of $\Phi^T P \Lambda P^\dagger \Phi$. In experimental applications, the matrix P is generally not directly accessible since it requires a complete description of the underlying system dynamics (8). In this section, a strategy for identification of A_μ directly from output data is illustrated. Consequently, it will be shown how this information can be used as part of a "direct method"[30] for computation of response curves associated with the data-driven phase and isostable coordinates in order to find a set of reduced equations similar to (7).

A. Inferring output dynamics from data

Consider the phase dynamics from (11) when noise is present and $U(t)$ is taken to be zero. The goal here is to determine a strategy to numerically compute A_μ as defined in (22) using only the output $y(t)$. Asymptotically expanding $\theta(t)$ in powers of ϵ yields $\theta(t) = \theta_0(t) + \epsilon\theta_1(t) + \cdots$. Substituting this into Eq. (11), one finds that $\dot\theta_0 = \omega$ so that $\theta_0(t) = \theta_0(t_0) + \omega(t - t_0)$. With this information, let $\Delta\theta(t) = \theta(t) - \theta_0(t_0) - \omega(t - t_0)$. The following analysis will suppose that both $\theta(t_0)$ and $\psi_k(t_0)$ for all k are order ϵ so that $\theta_0(t_0) = 0$.

Consider the behavior of the phase θ and isostable coordinate ψ_j in response to noise. Starting with ψ_j, letting $\psi_j(t) = R_j(t)\exp(\kappa_j t)$, substituting this into the isostable relations from (11) gives

$$\dot R_j\exp(\kappa_j t) + \kappa_j(t)R_j\exp(\kappa_j t) = \dot\psi_j = \kappa_j R_j\exp(\kappa_j t) + \sum_{i=1}^{N}\epsilon I_{j,i}(\theta)\eta_i(t),$$

$$\dot R_j = \exp(-\kappa_j t)\sum_{i=1}^{N}\epsilon I_{j,i}(\theta)\eta_i(t), \qquad (38)$$

where $I_{j,i}$ is the ith component of I_j. Integrating both sides yields

$$R_j(t) = R_j(t_0) + \epsilon\int_{t_0}^{t}\sum_{i=1}^{N}\left(\exp(-\kappa_j s)I_{j,i}(\theta)\eta_i(s)\right)ds$$

$$= R_j(t_0) + \epsilon\int_{t_0}^{t}\sum_{i=1}^{N}\left(\exp(-\kappa_j s)I_{j,i}(\omega s)\eta_i(s)\right)ds + \mathcal{O}(\epsilon^2). \quad (39)$$

Finally, substituting $R_j(t) = \psi_j(t)\exp(-\kappa_j t)$, one finds that to leading order ϵ,

$$\psi_j(t) = \psi_j(t_o)\exp(\kappa_j(t - t_0)) + \epsilon\int_{t_0}^{t}\sum_{i=1}^{N}\left(\exp(\kappa_j(t-s))I_{j,i}(\omega s)\eta_i(s)\right)ds. \quad (40)$$

The noise terms from the stochastic integral from (40) have a mean of zero and hence, neglecting terms of order ϵ^2 and higher,

$$E[\psi_j(t)] = \psi_j(t_0)\exp(\kappa_j(t - t_0)). \qquad (41)$$

The expected value of $\theta(t)$ is much more straightforward to find. Substituting the asymptotic expansion for θ into (11), again when $U(t) = 0$, one finds that $\dot\theta_1 = \sum_{k=1}^{N}Z_k(\theta(t))\eta_k(t)$. Direct integration yields $\theta_1(t) = \theta_1(t_0) + \int_{t_0}^{t}\sum_{k=1}^{N}Z_k(\theta(s))\eta_k(s)ds$. The terms of the stochastic integral have a mean of zero and hence $E[\theta_1(t)] = \theta_1(t_0)$. Using information about the expected values of the random variables θ and ψ_k, and again using $t_0 = 0$ for simplicity, to leading order ϵ, one finds that the expected value of (12) is

$$E[y_i(t) - y^\gamma(t)] = \sum_{k=1}^{N-1}\left[C^\theta(\omega t)p_k^\theta(\omega t)\psi_k(0)\exp(\kappa_k t)\right]$$

$$+ \Delta\theta(0)C^\theta(\omega t)p_N^\theta(\omega t) + \mathcal{O}(\epsilon^2). \qquad (42)$$

Notice that the right hand side of (42) is the same as the right hand side of (12), i.e., the solution when no noise is considered. Recalling

the definitions y_i, y_i^+, γ, P, Ψ, and Λ from (15), (17), and (21), using (42), one can show

$$E[y_i - \gamma] = P\Psi_i,$$
$$E[y_i^+ - \gamma] = P\Lambda\Psi_i. \qquad (43)$$

Let y_i^R and y_i^{+R} be random variables representing measurements taken starting at t_i and $t_i + T$, respectively, with $\theta(t_i) \approx \theta(t_i + T) \approx 0$. Noting that Φ and γ are both matrices with constant values, the expected value of the resulting measurements of the associated POD coefficients can be written as

$$E[\mu_i] = E[\Phi^T(y_i^R - \gamma)]$$
$$= \Phi^T P\Psi_i$$
$$= \mu_i + \mathcal{O}(\epsilon_1), \qquad (44)$$

$$E[\mu_i^+] = E[\Phi^T(y_i^{+R} - \gamma)]$$
$$= \Phi^T P\Lambda\Psi_i$$
$$= \Phi^T P\Lambda(P^\dagger\Phi\mu_i + \mathcal{O}(\epsilon_1))$$
$$= A_\mu\mu_i + \mathcal{O}(\epsilon_1),$$

where equality in the second to last line comes from (20). It should be emphasized that (44) is derived with the assumption that white noise enters additively in (8) and that the noise intensity is small enough so that terms from the Ito correction are negligible.[23] Multiplicative noise, colored noise, and larger intensity noise will influence the behavior of the system in ways that are not explicitly considered here. The relationships in (44) suggest a strategy to estimate the matrix A_μ from data using the following procedure provided that enough POD modes are taken so that ϵ_1 is small:

Step (1) For a time series of oscillatory data, select some distinct feature to correspond to $\theta \approx 0$. This feature should be easily detectable and occur robustly exactly once per cycle (e.g., the crossing of a voltage threshold during the upstroke of an action potential for periodically firing neuron). Note that the observation of this feature in the dataset does not generally indicate $\theta = 0$ exactly on a cycle-to-cycle basis because the transient decay of amplitude coordinates can influence the moment that the threshold is crossed.[69] Take the average return time to this feature over multiple cycles to correspond to the nominal period of oscillation T. From the data, extract vectors $y_j \in \mathbb{R}^\beta$ for $j = 1, \ldots, \alpha$ that correspond to discrete measurements taken every $\Delta t = T/(\beta - 1)$ time units with the first entry of each vector corresponding to a moment that $\theta \approx 0$ is detected. Take γ (the nominal periodic orbit) to be $\frac{1}{\alpha}\sum_{j=1}^{\alpha}y_j$ and define the matrix Y according to (13).

Step (2) Perform POD on the elements of the matrix Y to obtain the basis vectors defined in (14). Choose m of these vectors so that $\sum_{j=1}^{m}\zeta_j / \sum_{j=1}^{\alpha}\zeta_j \approx 1$ to incorporate in a reduced POD basis that defines the matrix Φ.

Step (3) For y_i from $k = 1, \ldots, \alpha - 1$ from step 1, extract a complementary set of vectors $y_i^+ \in \mathbb{R}^\beta$ defined as follows:

supposing that the elements of y_i correspond to measurements taken at $t_i, t_i + \Delta t, \ldots, t_i + (\beta - 1)\Delta t$, let the elements of y_i^+ correspond to measurements taken at $t_i + (\beta - 1)\Delta T, t_i + \beta \Delta T, \ldots, t_i + (2\beta - 2)\Delta t$. Note that with this definition, $\Phi^T(y_i - \gamma)$ and $\Phi^T(y_i^+ - \gamma)$ correspond to μ_i and μ_i^+, respectively.

Step (4) By defining the matrix $Y^+ = \begin{bmatrix} y_1^+ - \gamma & \ldots & y_\alpha^+ - \gamma \end{bmatrix}$, the relationship $A_\mu Y = Y^+$ can be used to estimate A_μ, which determines the dynamics of μ_i. As shown in Appendix A, $A_\mu = Y^+ Y^\dagger$ can be chosen to minimize $||A_\mu Y - Y^+||_F$.

B. Reduced modeling and direct method for measurement of phase and isostable response curves

Data-driven isostable coordinates can be used to infer phase-amplitude reduced models from experimental systems using a strategy similar to the direct method,[22,30,47] which has been well established for the determination of phase models of the form (2). To do so, suppose one has a general system of the form (8) without noise. Taking $U(t) = qu(t)$ with $q \in \mathbb{R}^N$ and recalling that F_U denotes the partial of F with respect to U, using the set of data-driven phase and isostable coordinates near the periodic orbit, one finds

$$
\begin{aligned}
\frac{d\theta^D}{dt} &= \frac{\partial \theta^D}{\partial x}^T \frac{dx}{dt} = \frac{\partial \theta^D}{\partial x}^T \\
&\quad \times \left[F(x^\gamma(\theta^D(t)), 0) + \epsilon F_U(x^\gamma(\theta^D(t)), 0)qu(t) \right], \\
\frac{d\psi_j^D}{dt} &= \frac{\partial \psi_j^D}{\partial x}^T \frac{dx}{dt} = \frac{\partial \psi_j^D}{\partial x}^T \\
&\quad \times \left[F(x^\gamma(\theta^D(t)), 0) + \epsilon F_U(x^\gamma(\theta^D(t)), 0)qu(t) \right],
\end{aligned}
\tag{45}
$$

where $x^\gamma(\theta^D)$ gives the limit cycle solution parameterized by the data-driven phase. The above equation can be simplified by using (35) and (33), which mandate that in the absence of external perturbation $\frac{d\theta^D}{dt} = \left(\frac{\partial \theta^D}{\partial x}^T \right) F(x^\gamma(\theta^D(t)), 0) = \omega + \mathcal{O}(\epsilon_1)$ and $\frac{d\psi_j^D}{dt} = \left(\frac{\partial \psi_j^D}{\partial x}^T \right) F(x^\gamma(\theta^D(t)), 0) = \kappa_j^D \psi_j^D + \mathcal{O}(\epsilon_1)$. Substituting this result into (45) yields the simplification

$$
\begin{aligned}
\frac{d\theta^D}{dt} &= \omega + \mathcal{Z}^D(\theta^D)\epsilon u(t) + \mathcal{O}(\epsilon_1), \\
\frac{d\psi_j^D}{dt} &= \kappa_j^D \psi_j^D + \mathcal{I}_j^D(\theta^D)\epsilon u(t) + \mathcal{O}(\epsilon_1),
\end{aligned}
\tag{46}
$$

where $\mathcal{Z}^D(\theta^D) = \frac{\partial \theta^D}{\partial x}^T F_U(x^\gamma(\theta^D(t)), 0)q$, $\mathcal{I}_j^D(\theta^D) = \frac{\partial \psi_j^D}{\partial x}^T F_U(x^\gamma(\theta^D(t)), 0)q$, and all partial derivatives are taken at $x^\gamma(\theta(t))$ and $U = 0$. For practical purposes, it will be assumed that the $\mathcal{O}(\epsilon_1)$ terms are small enough to be neglected from (46).

The unknown functions $\mathcal{Z}(\theta^D)$ and $\mathcal{I}_j(\theta^D)$ from (46) can be found by using a strategy that mirrors the direct method for calculating phase response curves. Specifically, by first allowing transient behavior to die out so that each ψ_j^D is zero (corresponding to the state being on the periodic orbit), a perturbation with magnitude u_M lasting t_L time units can be given at some known θ^D. The resulting change to each of the isostable and phase coordinates ($\Delta \psi_j^D$

and $\Delta \theta^D$, respectively) can then be measured from the response using the relation given in (29) and (30), respectively. An approximation for the data-driven isostable and phase response curves can then be obtained according to $\mathcal{I}_j^D(\theta_0) = \Delta \psi_j^D/(u_M t_L)$ and $\mathcal{Z}^D(\theta_0) = \Delta \theta^D/(u_M t_L)$, respectively. This process can be repeated over multiple measurements by applying perturbations at different initial values of θ^D and resulting periodic curves can be fit to the data.

To summarize, once phase response curves have been identified, a reduced order model can be obtained according to

$$
\begin{aligned}
\dot{\theta}^D &= \omega + \epsilon \mathcal{Z}^D(\theta^D)u(t), \\
\dot{\psi}_j^D &= \kappa_j^D \psi_j^D + \epsilon \mathcal{I}_j^D(\theta^D)u(t), \quad j = 1, \ldots, m-1.
\end{aligned}
\tag{47}
$$

Additionally, from (37), if $\theta^D \approx 0$, the output over the next cycle can be predicted as $s_1 = \gamma + \Phi V L \Psi^D + \mathcal{O}(\epsilon_1)$.

The reduced framework (47) is particularly useful for systems with noise and other uncertainties because direct information about the phase and isostable coordinates can be obtained through real time measurement of the system output. For instance, by measuring the output over $[t_1, t_1 + T]$ with initial conditions $\theta^D(t_1) \approx 0$, one can use $k = 1$ in (30) and (29) to obtain an estimate of $\theta^D(t_1)$ and each $\psi_j^D(t_1)$, i.e., the reduced coordinates at time t_1. In order to obtain an estimate of these coordinates at $t_1 + T$, nominally $\dot{\theta}^D = \omega$ so that $\theta^D(T + t_1) \approx \theta^D(t_1)$. Likewise, due to the nominal exponential decay of isostable coordinates, $\psi_j^D(T + t_1) \approx \exp(\kappa_j^D T)\psi_j^D(t_1)$ can be used.

V. AN ILLUSTRATIVE EXAMPLE IN A COUPLED POPULATION OF NEURONS

As an illustration of the proposed data-driven phase and isostable reduction method, consider a population of coupled, noisy, periodically firing thalamic neurons taken from[56]

$$
\begin{aligned}
C\dot{V}_i &= -I_L(V_i) - I_{Na}(V_i, h_i) - I_K(V_i, h_i) - I_T(V_i, r_i) \\
&\quad + I_{SM} + \frac{\alpha_c}{N} \sum_{k=1}^{N} (V_k - V_i) + \sqrt{2D}\eta_i(t) + u_i(t), \\
\dot{h}_i &= (h_\infty(V_i) - h_i)/\tau_h(V_i), \\
\dot{r}_i &= (r_\infty(V_i) - r_i)/\tau_r(V_i), \quad i = 1, \ldots, N, \\
y(t) &= \frac{1}{N} \sum_{j=1}^{N} V_i(t).
\end{aligned}
\tag{48}
$$

Here, V_i is the transmembrane voltage of neuron i, h_i and r_i are corresponding gating variables, $u_i(t)$ represents direct current input applied to neuron i, $C = 1 \, \mu\text{F/cm}^2$, $I_{SM} = 5 \, \mu\text{A/cm}^2$ is a baseline current so that each neuron fires periodically in the absence of external input, noise, and coupling, $\alpha_c = 0.05 \, \text{mS/cm}^2$ is the magnitude of the electrotonic coupling strength,[31] $\sqrt{2D}\eta_i(t)$ is an independent and identically distributed zero-mean white noise process with intensity $D = 0.1$, and all other functions and parameters are identical to those discussed in greater detail in Ref. 56. In all simulations to follow, $N = 10$ neurons will be considered.

For the neural model (48), the location at which $\theta \approx 0$ (i.e., a spike) is defined to occur at the moment when the transmembrane
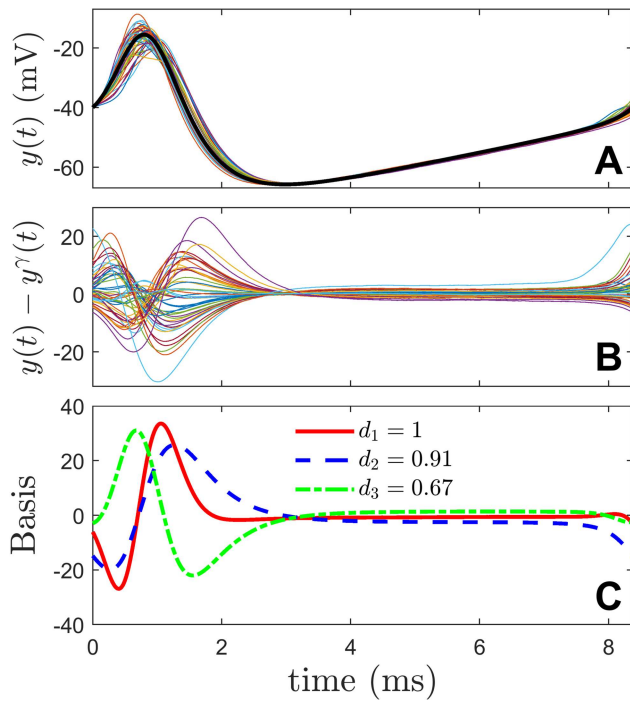
FIG. 1. Colored lines in panel (a) shows example outputs from (48) measured over the course of one period and the thick black line shows $y^\gamma(t)$. Panel (b) shows the same data, but with $y^\gamma(t)$ subtracted. The data-driven phase and isostable framework identifies three dominant basis elements and associated decay rates shown in panel (c).

voltage crosses $-40\,$mV during the upstroke of an action potential. Simulation of (48) is performed for 20 000 ms of simulated data resulting in $\alpha = 2385$ interspike intervals yielding an average period of $T = 8.38\,$ms. Colored lines in panel (a) of Fig. 1 show examples of the model output plotted over one period. The mean of these outputs is taken to be $y^\gamma(t)$. The same data from panel (a) are shown

in panel (b) with $y^\gamma(t)$ subtracted. Steps 1–4 from Sec. IV A are performed to identify a POD basis and associated dynamics A_μ. Here, choosing $m = 3$ POD modes yields an adequate representation of the deviation from the mean, with $\sum_{j=1}^{3} \zeta_j / \sum_{j=1}^{\alpha} \zeta_j = 0.987$. Data-driven phase and isostable coordinates are defined as in Sec. III B and the resulting mode shapes [i.e., the columns of ΦVL from (37)] are shown in panel (c) along with their associated decay rates d_j.

In the examples to follow which use the neural model (48), external stimulation is chosen so that $u_i(t) = 0$ for $i > N/2$ at all times. In other words, perturbation is only applied to half of the neurons. For other choices and distributions of $u_i(t)$, the measured phase and isostable response curves would be different, but the strategy for measuring response curves would not change.

Figure 2 gives an illustration of the strategy used to measure a PRC in the data-driven isostable reduced model (47) detailed in Sec. IV B. Over multiple trials, a short perturbation with $u_M = 250\,\mu\text{A}/\mu\text{F}$ lasting $t_L = 0.02\,$ms (resulting in a 5 mV increase in transmembrane voltage) is applied to half of the neurons at a known phase, θ^D. The third full interspike interval occurring after the perturbation is considered to calculate the resulting changes to the phase and isostable coordinates [panels (a) and (b)]. This delay corresponds to taking $k = 3$ in (29) and (30) and is chosen so that rapidly decaying transient behavior is not considered in the algorithm. The inner product of the resulting value of $y_i - y^\gamma$ is taken with $w_j^T \Phi^T$ [panels (c)–(e)] yielding associated datapoints [panels (f)–(g)]. This process is repeated for 500 trials and phase and isostable response curves are fit to functions of the form $\sum_{n=0}^{3} \left[a_n \sin(n\theta) + b_n \cos(n\theta) \right]$. The resulting fits are shown as black lines in panels (f)–(h).

The expected output of the reduced model (47) is compared to the true output from full model simulations of (48). In results shown in panel (a) of Fig. 3, the full model (48) is integrated forward in time allowing initial transients to decay so that both ψ_1^D and ψ_2^D start near zero. At phase $\theta^D = 4\,$rad, a perturbation of $u_M = 400\,\mu\text{A}/\mu\text{F}$ lasting $t_L = 0.02\,$ms is applied to half of the neurons and the full model output is compared to the predicted output from the relation (37). Good agreement is observed between the full and reduced model, particularly in the replication of the dual peaks in model output
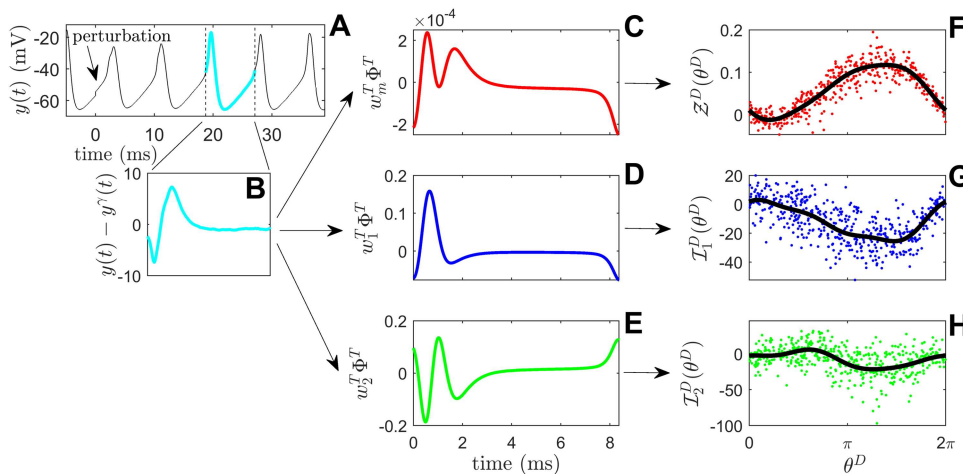

FIG. 2. An illustration of the direct method for determining the response curves from the reduced model (47). After allowing transient dynamics to decay, perturbations are applied at known values of θ^D [panel (a)]. For each trial, the third full interspike interval after the perturbation is extracted [panel (b)] and its inner product with appropriately scaled vectors $w_j^T \Phi^T$ [shown in panels (c)–(e)] is used to identify the resulting changes to the data-driven phase and isostable coordinates. Individual dots in panels (f)–(h) represent calculated values of $\Delta \theta^D / u_M t_L$ and $\Delta \psi_j^D / u_M t_L$ on individual trials and solid lines are fit to the datapoints.
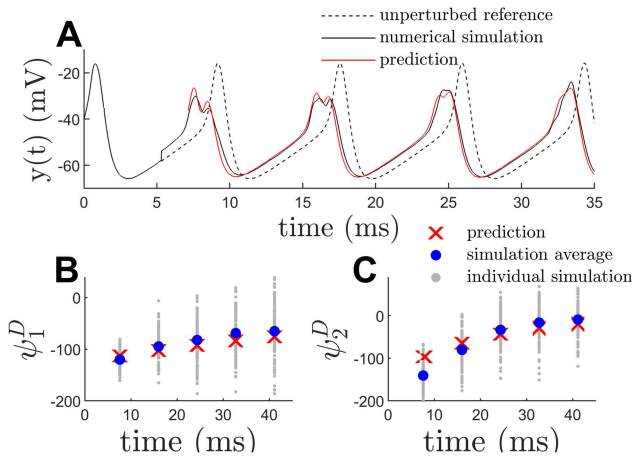
FIG. 3. Comparisons between outputs of the full and reduced model. Panel (a) shows model output for perturbation of $u_M = 400\,\mu A/\mu F$ lasting $t_L = 0.02$ applied to half of the neurons and given at a phase of $\theta^D = 4\,\mathrm{rad}$. Good agreement is observed between the full and predicted model outputs shown as black and red lines, respectively. The dashed line illustrates the predicted behavior had the perturbation not been given. In the bottom panels, the noisy full model (48) is simulated until initial transients die out, after which a perturbation of $u_M = 250\,\mu A/\mu F$ lasting $t_L = 0.02\,\mathrm{ms}$ is applied to half of the neurons when the phase reaches $\theta^D = 4\,\mathrm{rad}$. Coordinates ψ_1^D and ψ_2^D are calculated according to (29) for individual trials (grey dots) using model output from the previous spike (i.e., with $k = 1$) and the process is repeated for 100 trials. Red crosses show the predicted values from the reduced model.

observed immediately after perturbation is applied that gradually disappear during recovery toward steady state. Panels (b) and (c) illustrate predicted and measured values of ψ_1^D and ψ_2^D after perturbation is applied indicating good agreement between the reduced and full models. Note that because the full model includes additive white noise, there is a spread in the values of the reduced coordinates over multiple simulations.

VI. OPTIMAL CONTROL OF A LARGE MODEL OF COUPLED CIRCADIAN OSCILLATORS WITH APPLICATIONS TO RECOVERY FROM JET LAG

In this section, a more involved application with relevance to circadian physiology will be considered to showcase the potential for the proposed phase-amplitude reduction to be used in feedback control applications. Circadian rhythms are evolutionarily advantageous for humans and other organisms to anticipate environmental changes to external lighting and temperature and respond accordingly.[51,74] In humans, circadian rhythms are governed by the suprachiasmatic nucleus (SCN),[54] a master pacemaker comprised of approximately 10 000 coupled neurons that maintains a free-running period of the circadian rhythm close to 24 h so that the circadian cycle can be stably entrained to a 24-h light–dark cycle.[14,24,73] Circadian misalignment (most commonly due to jet lag) represents a disruption to this steady entrainment and is caused by a mismatch between the environmental time and one's internal circadian clock.[1,57] Without treatment, recovery times from jet lag have

been measured to be approximately 1 and 1.5 time zones per day for eastward and westward travel, respectively. Recent years have seen a growing interest in the development of specific treatments to accelerate reentrainment to limit the negative effects of jet lag.

One useful approach for developing treatments relies on the circadian pacemaker's phase response curve (PRC) to perturbations such as bright light.[33,43] For instance, bright light given in the evening before bedtime has the tendency to delay the onset of sleep while the same bright light tends to advance one's circadian phase when applied in the morning. Information about human PRCs has proven invaluable in the development of simple and practical jet-lag recovery treatments that incorporate carefully timed combinations of light exposure and light avoidance.[11,63]

In addition to information gathered directly from human subjects, quantitative and qualitative numerical models detailing circadian oscillations have also been useful in the development and understanding of new jet-lag treatment strategies.[20,25,34,38] These and related models can be used to find mathematically optimal schedules of light avoidance and exposure for rapidly acclimating to a new time zone,[3,15,60] some of which have been implemented successfully in real-time using smartphones.[49] While these computational models can be used to identify general principles that result in rapid recovery from circadian misalignment, they cannot be tailored to an individual without a personalized model. Additionally, it is not obvious how to incorporate state feedback into these control algorithms, necessitating a trade-off between optimality of the control strategy and robustness to noise and other uncertainties. As illustrated in the example to follow, provided a single output signal can be measured from a circadian model, using the data-driven phase and isostable reduction strategy detailed in Sec. IV, it is possible to infer a model of the form (47), which can subsequently be used to determine time-optimal strategies for recovery from circadian misalignment. Additionally, in contrast to other recently proposed circadian misalignment recovery strategies, model output can be used to provide a real-time estimate of the data-driven phase and isostable coordinates resulting in a feedback control strategy that is robust to noise and modeling uncertainty.

To begin, a model describing the entrained behavior of N coupled SCN cells[25] will be considered

$$
\begin{aligned}
\dot{a}_i &= h_1 \frac{K_1^n}{K_1^n + c_i^n} - h_2 \frac{a_i}{K_2 + a_i} + h_c \frac{KF(t)}{K_c + KF(t)} \\
&\quad + S_i\left[L(t_s) + \Delta L(t)\right] + \sqrt{2D}\eta_i, \\
\dot{b}_i &= h_3 a_i - h_4 \frac{b_i}{K_4 + b_i}, \\
\dot{c}_i &= h_5 b_i - h_6 \frac{c_i}{K_6 + c_i}, \\
\dot{d}_i &= h_7 a_i - h_8 \frac{d_i}{K_8 + d_i}, \quad i = 1, \ldots, N, \\
\dot{t}_s &= 1, \\
y(t) &= F(t) \equiv (1/N) \sum_{j=1}^{N} d_i(t).
\end{aligned}
\tag{49}
$$

Above, for the ith cell, the variable a_i represents concentrations of an mRNA clock gene, b_i and c_i are the associated protein and nuclear form of the protein, respectively, d_i is a neurotransmitter with $F(t)$ being its population average value, $L(t_s)$ is the external light input, $t_s \in \mathbb{S}^1$ is a time-like variable that takes values in the range of $[0, 24)$ and determines the nominal external light magnitude, $N = 3000$, and $S_i = \max(0.5 + 0.4\mathcal{N}(0, 1), 0)$ is an individual cell's sensitivity to light where $\mathcal{N}(0, 1)$ is a standard normal distribution. A light exposure or avoidance strategy can be implemented with $\Delta L(t)$ in order to promote reentrainment. Noise with intensity $D = 5 \times 10^{-4}$ is added to the a_i variable of each oscillator. Nominal parameters are taken to be identical to those given in Fig. 1 of Ref. 25 with the exception of $h_2 = 0.503$ and $h_c = 1.2$. Heterogeneity is incorporated by drawing $h_1, h_2, h_3, h_4, h_5,$ and h_6 from a normal distribution with a standard deviation of 0.03 with the mean value equal to the nominal parameter value.

In simulations of (49), the external light is taken to be

$$L(t_s) = 0.015 \left[\frac{1}{1 + \exp(-4(t_s - 6))} - \frac{1}{1 + \exp(-4(t_s - 18))} \right],$$

(50)

where $\Delta L(t)$ can be used to implement a light exposure or avoidance strategy in order to promote reentrainment. Additionally, the constraint $0 \leq L(t_s) + \Delta L(t) \leq 0.015$ is required to limit the magnitude of total light input and to prevent the application of a negative value of light. When $\Delta L(t) = 0$, the model (49) stably entrains to the 24-h light-dark cycle as shown in panels (a) and (b) of Fig. 4. Note that because t_s is explicitly included as a state, the model (49) is autonomous. From this perspective, using the definition of isochrons from (6), one can show that the asymptotic phase $\theta = 2\pi t_s/24$. This point was illustrated for a different periodically forced model in Ref. 66. Because the asymptotic phase, θ, is known exactly, it is not necessary to use a data-driven phase, θ^D, in a reduced order model. However, data-driven isostable coordinates will still be used to determine a reduced order model for the amplitude dynamics.

While the coupled oscillator model (49) contains thousands of states, it is possible to use the output from (49) to determine a reduced order model based on data-driven isostable coordinates using the strategy detailed in Sec. III. This is done by taking $\alpha = 968$ measurements over t_s ranging from 0 to 24. Individual trials are plotted in panel (c) of Fig. 4 as colored lines; the periodic orbit $y^\gamma(t)$ is taken to be the average value of these outputs and plotted as a black line. Next POD analysis is implemented on the data set where it is found that $m = 5$ modes provides a good representation of the data with $\sum_{j=1}^{5} \zeta_j / \sum_{j=1}^{\alpha} \zeta_j = 0.987$. Subsequently, A_μ is fit to the data using the strategy detailed in Sec. IV A and the resulting matrix is used to determine the data-driven isostable coordinates for this model. The resulting numerically determined decay rates are $d_1 = 0.23 + 0.77i$, $d_2 = 0.23 - 0.77i$, $d_3 = 0.07$, $d_4 = -0.06$, and $d_5 = 0.02$. The corresponding modes for the complex conjugate eigenvalues are shown in panel (d) of Fig. 4, and the other modes are ignored in the reduction to follow because the decay rates are so rapid.

A reduced set of three differential equations will be used to characterize the amplitude dynamics of perturbations to the limit
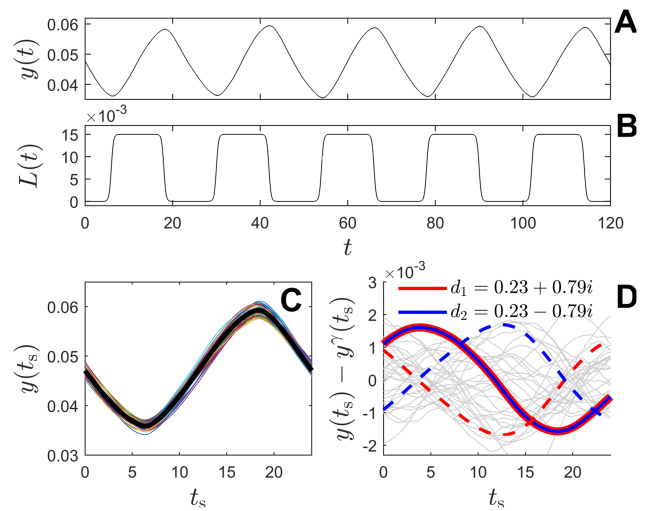

FIG. 4. Panel (a) shows a representative output from the SCN model (49) with the 24-h light-dark cycle shown in panel (b). In panel (c), representative measurements are shown as colored lines over t_s ranging from 0 to 24 and the average of these outputs is plotted as a black line. In panel (d), grey lines show individual measurements corresponding to $y(t_s) - y^\gamma(t_s)$ with real and imaginary components of the modes corresponding to data-driven isostable coordinates shown as solid and dashed lines, respectively.

cycle of (49),

$$\dot{t}_s = 1,$$
$$\dot{\psi}_1^D = \kappa_1^D \psi_1^D + \mathcal{I}_1^D(2\pi t_s/24)\Delta L(t),$$
$$\dot{\psi}_2^D = \kappa_2^D \psi_2^D + \mathcal{I}_2^D(2\pi t_s/24)\Delta L(t).$$

(51)

Here, t_s is a phase-like variable and is required to implement the external forcing from the 24-h light-dark cycle. It is assumed that there is no direct control over t_s and hence no associated response curve for the variable t_s; it can only increase at a constant rate. In this example, because d_1 and d_2 are complex conjugate, the data-driven isostable coordinates ψ_1^D and ψ_2^D and the response curves $\mathcal{I}_1^D$ and $\mathcal{I}_2^D$ must also be complex conjugates. This point is discussed in greater detail in Ref. 66.

The direct method as described in Sec. IV B is used to determine isostable response curves. For each trial, any initial transients in the model (49) are allowed sufficient time to decay so that ψ_1^D and ψ_2^D are close to zero. In each trial, for a given initial value of t_s the external light is set to either its maximum or minimum value (0.015 or 0, respectively) for a duration of $t_L = 1$ h. The value ΔL is determined and the resulting shifts in the isostable coordinates are determined directly from (29) using $k = 2$. Resulting estimates for the isostable response curves are shown in panels (a) and (b) of Fig. 5. Dots and open circles indicate trials for which ΔL is positive and negative, respectively. A basis of the form $\sum_{n=0}^{2} \left[a_n \sin(n\theta) + b_n \cos(n\theta) \right]$ is fit to the datapoints and taken to be the response curves.
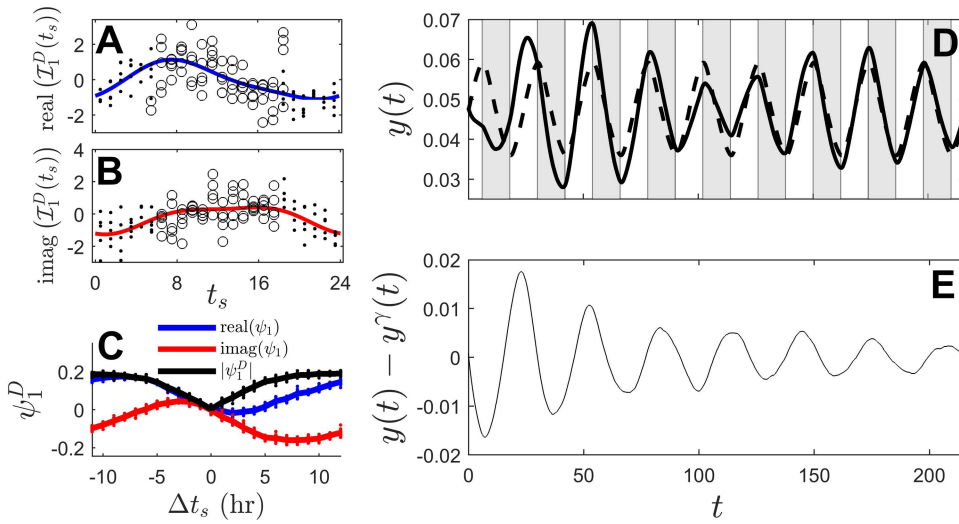
FIG. 5. Panels (a) and (b) illustrate results from using the direct method to infer the necessary parameters of the response curves from (47). Dots and open circles indicate values obtained using positive and negative values of ΔL, respectively, and colored lines show fits to the resulting data. Panel (c) illustrates the effect on ψ_1^D of suddenly shifting t_s by Δt_s hours. Panels (d) and (e) illustrate the recovery from a time shift of $\Delta t_s = 12$ h. Solid and dashed lines in panel (d) show the perturbed output $y(t)$ and the nominal entrained output $y^\gamma(t)$, respectively. White and grey bands correspond the nominal light-dark cycle indicating environmental day and night, respectively. Panel (e) shows the difference between the solid and dashed lines from panel (d), illustrating the expected exponential decay during reentrainment.

In the context of the reduced model (51), full entrainment corresponds to a state for which $\psi_1^D = \psi_2^D = 0$. Circadian misalignment can be caused by sudden shifts in the time t_s, which reflect changes in the environmental time occurring after rapid travel across multiple time zones. Panel (c) shows the changes in ψ_1 resulting from advancing or delaying t_s by Δt_s hours starting from an initial time of $t_s = 0$. The shifts in isostable coordinates are calculated from (29) taking $k = 1$, i.e., using output from the first full period occurring after the shift. Dots show the results from individual trials and solid lines show the average for a given value of Δt_s. In general, choosing Δt_s closer to ± 12 yields larger values of $|\psi_1|$ and will take longer to recover to the nominal entrained solution. Note that starting the simulation at values of t_s other than zero does not qualitatively change the results. Panels (d) and (e) show output resulting from taking $\Delta t_s = 12$ h illustrating the expected exponential decay toward the nominal entrained solution.

In the control examples to follow, the transformation $\rho_1 \equiv \psi_1^D + \psi_2^D = 2\,\text{real}(\psi_1^D)$ and $\rho_2 \equiv -i(\psi_1^D - \psi_2^D) = 2\,\text{imag}(\psi_1^D)$ will be used in order to avoid the use of complex numbers. This change of coordinates transforms (49) to

$$\dot{t}_s = 1,$$

$$\dot{\rho}_1 = \text{real}(\kappa_1^D)\rho_1 - \text{imag}(\kappa_1^D)\rho_2 + 2\,\text{real}(\mathcal{I}_1^D(2\pi t_s/24))\Delta L(t), \quad (52)$$

$$\dot{\rho}_2 = \text{imag}(\kappa_1^D)\rho_1 + \text{real}(\kappa_1^D)\rho_2 + 2\,\text{imag}(\mathcal{I}_1^D(2\pi t_s/24))\Delta L(t),$$

Equation (52) is in the general form

$$\dot{x} = R(x) + B(x)u(t), \quad (53)$$

where $x = [t_s \ \rho_1 \ \rho_2]^T$, $B(x) = [0 \ \ 2\,\text{real}(\mathcal{I}_1^D(2\pi t_s/24)) \ \ 2\text{imag}(\mathcal{I}_1^D(2\pi t_s/24))]^T$, and $u(t) = \Delta L(t)$. Note here that $t_s \in \mathbb{S}^1$ dictates the nominal input from the external light source and $t \in \mathbb{R}$ is the overall time variable. Recalling that $\rho_1 = \rho_2 = 0$ corresponds to the entrained solution, an optimal jet-lag mitigation control strategy can be formulated as a minimum time to reach problem and solved using a Hamilton-Jacobi-Bellman approach.[35] This is described in detail in Appendix B, but the general goal is to find an allowable control $u^*(t)$

that takes an initial state x to a prespecified target set $\mathcal{T}$ (with a small value of ρ_1 and ρ_2) in the minimum time possible. Practically, this can be done by first defining a cost functional

$$J(x, u(t)) = \int_0^{t_{\text{targ}}} 1\,dt = t_{\text{targ}}(x, u(t)), \quad (54)$$

where t_{targ} is the time required to reach $\mathcal{T}$ starting from x under the application of $u(t)$. Next, one can define a value function (sometimes referred to a cost-to-go function) as the minimum possible value of t_{targ} for a given x

$$\mathcal{V}(x) = \inf_{u_{\min} \le u \le u_{\max}} J(x, u(t)) = \inf_{u_{\min} \le u \le u_{\max}} t_{\text{targ}}(x, u(t)). \quad (55)$$

One can numerically solve for $\mathcal{V}(x)$ by defining an auxiliary function $\Upsilon(x, t)$, which evolves according to

$$0 = \Upsilon_t(x, t) + \nabla\Upsilon(x, t)^T R(x) + \min_{u_{\min} \le u \le u_{\max}} \left[\nabla\Upsilon(x, t)^T B(x)u(t)\right]. \quad (56)$$

For appropriately chosen initial conditions of $\Upsilon(x, 0)$ (as explained in Appendix B), the value function is determined by $\mathcal{V}(x) = \{t | \Upsilon(x, t) = 0\}$. Once $\mathcal{V}(x)$ is known, the control policy

$$L(t) + \Delta L(t) = \begin{cases} 0.015 & \text{if } \nabla\mathcal{V}^T(x(t))B(x(t)) < 0, \\ 0 & \text{otherwise}, \end{cases} \quad (57)$$

provides the optimal control input that reaches $\mathcal{T}$ in the minimum time possible. Control policies of the form (57) are often called bang-bang control[35] because the minimum or maximum possible control is always applied. Finally, as a matter of practical implementation, switching from maximum to minimum allowable control will only be allowed once per hour of simulated time in order to avoid solutions that switch rapidly back and forth. A target set that is close to $\rho_1 = \rho_2 = 0$ will be chosen. In simulations of (49) when the model is fully entrained, the value of $||\rho|| \equiv \sqrt{\rho_1^2 + \rho_2^2}$ stays near 0.015 due to noise in the system and occasionally is driven close to 0.05. Because of this nominal behavior, the target set $\mathcal{T}$ is chosen to be the set of all
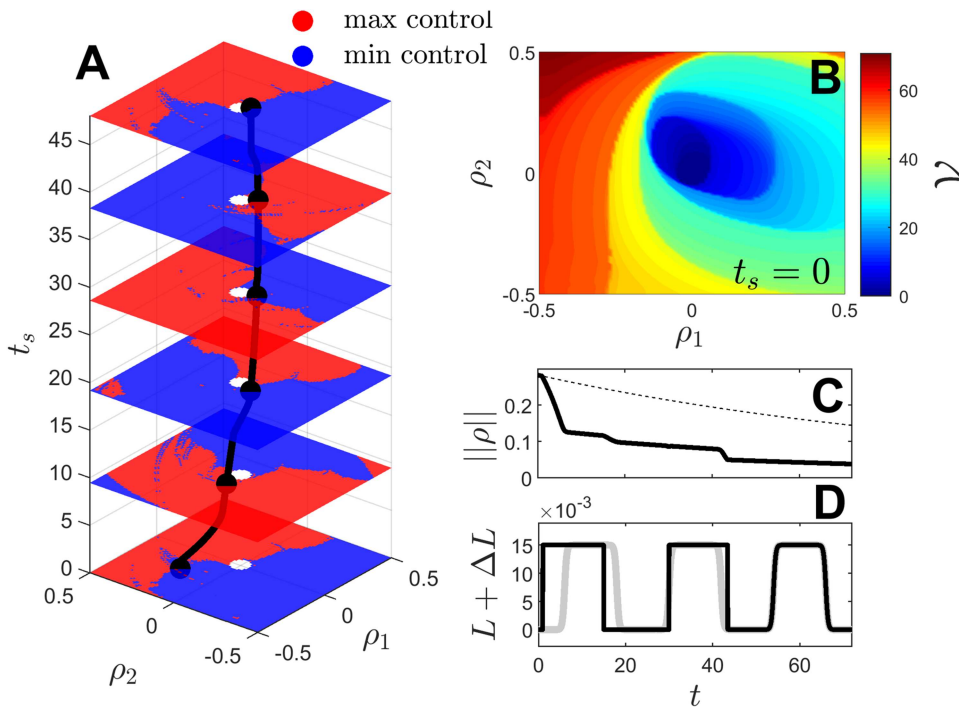
FIG. 6. The optimal control policy (57) is shown in panel (a) for various values of t_s. Red and blue regions indicate that the control policy calls for the maximum and minimum light, respectively. The white regions show locations inside the target set for which the controller is turned off. The black line shows an example of an optimally controlled trajectory of (52); black dots correspond to the intersection of the trajectory with each cross section. A cross section of the value function $\mathcal{V}(x)$ at $t_s = 0$ is shown in panel (b). The black line in panel (c) shows $||\rho||$ for the trajectory from panel (a) with the dashed line showing the decay in the uncontrolled case. Panel (d) shows the associated optimal light input as a black line along with the nominal light-dark cycle in grey.
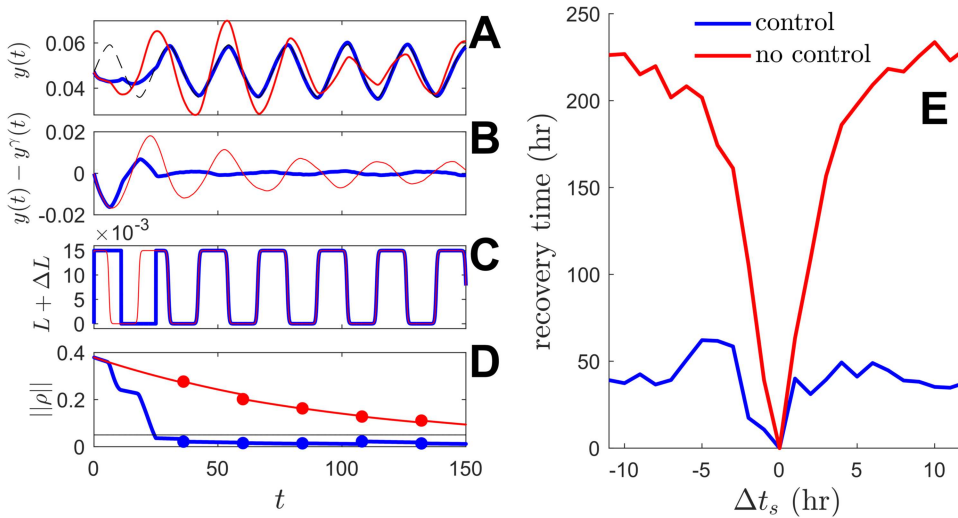

FIG. 7. Results from implementing the optimal reentrainment strategy (57) on the full, noisy model (49). Panel (a) compares the uncontrolled recovery from a $\Delta t_s = 12$ h shift (red line) to the recovery under the application of control (blue line). For reference, the black dashed line shows the fully entrained solution output $y^\gamma(t)$. This solution becomes indistinguishable from the blue line at $t \approx 30$ h. Panel (b) shows $y(t) - y^\gamma(t)$ for the same simulation. In panel (c), the uncontrolled light input is shown in red and the controlled light input is given in blue. Panel (d) gives the value of $||\rho||$; solid lines show the value used as an estimate by the controller and dots indicate states updated every 24 h directly from the model output in order to account for the influence of noise and uncertainty. The black line shows the boundary of the target set. After the state reaches the target set, the control is turned off. Simulations are repeated for 50 trials for each value of Δt_s to measure the recovery time with and without control. The averaged values are shown in panel (e).

states for which $||\rho|| \leq 0.05$. The partial differential equation (56) is computed using a freely available Matlab toolbox[44] and used to determine the value function $\mathcal{V}(x)$. Panel (a) of Fig. 6 shows cross sections of the resulting control policy with red and blue colors prescribing the maximum and minimum possible control application, i.e., full brightness or total darkness, respectively. Inside the target set, ΔL is taken to be zero. Panel (b) shows the value function for $t_s = 0$. In general, values that are a large distance from $\rho_1 = \rho_2 = 0$ have larger corresponding values of $\mathcal{V}(x)$. Panel (a) also shows a controlled trajectory of (52) starting at $(\rho_1, \rho_2, t_s) = (-0.2, 0.2, 0)$ as a thick black line with black dots to help identify the intersections with the individual planes. Panel (c) shows values of $||\rho||$ plotted over time for this optimally controlled trajectory and panel (d) shows the optimal value of $L(t) + \Delta L(t)$.

The optimal control strategy obtained for the reduced model is applied to the full model Eq. (49). In order to apply the optimal control policy (57), the values of ρ_1 and ρ_2 are estimated from the model output. This can be done by periodically obtaining an estimate for the isostable coordinates (and hence ρ_1 and ρ_2) according to (29) using $k = 1$ and using the reduced model equations (52) to estimate the state in between measurements. In the illustrations to follow, state updates obtained directly from model output are taken every 24 h. In the simulations shown in Fig. 7, t_s is shifted by an amount Δt_s and the control strategy is implemented to recover from the misalignment as quickly as possible. In each simulation, initial values of ρ_1 and ρ_2 are calculated from the average values of ψ_1^D shown for the applied Δt_s. Panels (a)–(d) show an example of the observed recovery using $\Delta t_s = 12$ h. Results using the optimal control strategy are shown in blue, and red lines illustrate the recovery in the absence of control. For each value of Δt_s, the simulation is run for 50 trials and the recovery times, defined to be the time it takes to reach the target set, are computed. The results are then averaged for each Δt_s and shown in panel (e) of Fig. 7. When implementing the optimal control strategy, recovery is sped up by a factor of 3–6 with the largest differences seen for the largest values of Δt_s.

As a final note, the magnitude of the external forcing is large enough that standard phase-only reduction methods using equations of form (2) cannot be used to understand the dynamical behavior of this model. As explained in Appendix C, the entrained periodic orbit is far enough away from the unperturbed periodic orbit so that the phase dynamics are not well-captured by a single phase response curve. This issue is also considered in applications motivated by circadian physiology in Refs. 16 and 17. For model (49), attempts at using the information from Fig. 8 lead to incorrect light avoidance and exposure strategies that inadvertently delay reentrainment. It is imperative to incorporate information about the amplitude coordinates for this example in order to find a viable control strategy.

VII. DISCUSSION AND CONCLUSION

This work suggests and investigates a framework for phase-amplitude reduction using so-called data-driven phase and isostable coordinates. This strategy uses well-established POD algorithms to identify a modal basis from output data from a limit cycle oscillator and subsequently draws on the recently developed theory of phase and isostable coordinates to define a new data-driven phase and

isostable coordinate framework. Similar to standard phase coordinates, the data-driven phase coordinates (30) explored here increase at a constant rate in the absence of perturbation. Likewise, data-driven isostable coordinates (29) decay exponentially in the absence of perturbation. Techniques are developed to infer the phase and isostable response curves in a reduced order model (47) from output data in order to characterize the influence of external perturbations. The resulting reduction is particularly well suited for application in dynamical systems with dominant nonlinearities.

While the proposed reduction strategy is developed with experimental systems in mind, it may be a useful alternative to standard phase and isostable reduction strategies as illustrated in Refs. 66, 68, and 71 for particularly high-dimensional numerical systems. Isostable response curves, $I_j(\theta)$ from (7), generally require the identification of periodic, dynamically unstable solutions of a generalized adjoint equation, a nontrivial computation in high-dimensional systems. Even when working with numerical simulations where the full dynamical equations are explicitly known, reduction using data-driven phase and isostable coordinates may be more feasible in particularly high-dimensional or complicated systems.

The efficacy of the data-driven phase and isostable reduction framework is illustrated on a 30 dimensional model of coupled, periodically firing neurons[56] and on a 12 000 dimensional model describing circadian oscillations (49). In the former application, the reduction framework is used to accurately predict model output in response to exogenous perturbation over the span of multiple periods. In the circadian example, an optimal control strategy is implemented that would be computationally prohibitive without first implementing the proposed reduction framework. The resulting control strategy decreases recovery times from circadian misalignment by a factor of 3–6 depending on the initial time shift with the largest improvements observed for the largest time shifts. Despite the high-dimension of the underlying equation in both applications considered here, the output behaviors are well-characterized by significantly lower order reduced models.

Much like existing POD and DMD algorithms, the reduction technique presented here can be implemented using only measured model output, making it amenable for the analysis of experimental systems. While standard implementations of POD and DMD do not make assumptions about the underlying system behavior, the data-driven phase and isostable reduction explicitly assumes that an underlying periodic orbit exists and that Floquet theory can be applied to characterize its dynamics for small perturbations. As a preliminary step in the data-driven reduction strategy, a POD algorithm is used to identify features of the transient decay of solutions toward the limit cycle. A central requirement of the reduction strategy is that enough POD modes are kept so that any residuals from (16) are small enough so that the underlying dynamics can be adequately captured. In the examples shown here, this was accomplished by choosing m so that $\sum_{j=1}^{m} \zeta_j / \sum_j \zeta_j \approx 0.99$. In principle, using more POD modes will increase this ratio and provide a better reduced model, but this will also require a larger matrix A_μ to be used when fitting the data to the model behavior. In practice, a good balance needs to be achieved between the number of POD modes kept and the accuracy of the resulting basis.

There are numerous limitations and opportunities for extension of the proposed data-driven phase and isostable coordinate

reduction strategy. First, while the data-driven phase and isostable coordinates can be defined in the basin of attraction of a limit cycle, the reduced equations used here (47) only use a linearized approximation of the phase and isostable coordinates with respect to the limit cycle. Previous work has focused on obtaining higher order accuracy reduced models using standard phase and isostable coordinates.[66,68] It is likely that the associated reduction framework can be extended in a similar manner. Additionally, while the data-driven phase and isostable reduction strategy can be used to accurately predict outputs in response to perturbation, specific estimates for the errors caused by the truncated residual terms from (15) have not been obtained. Other frameworks such as POD come with an explicit energy estimate that gives a sense of the energy contained by each mode and hence its relative importance. It would be of interest to develop similar estimates for accuracy of the proposed reduction technique in relation to the number of modes used. Also, the current reduced framework is designed for replicating the model output when only one temporal measurement is taken, but it may be possible to extend the current framework for applicability in situations where data are measured from multiple channels simultaneously. All examples presented in this study are implemented in numerical systems. An experimental implementation would likely need to extend the proposed framework to account for the possibility of drifting parameters that would cause the phase and amplitude response properties of a given system to change over time. Furthermore, the illustrative examples given in this manuscript use a large number of trials to fit the underlying reduced model. Quantitative investigations into how many trials are necessary to obtain an accurate reduced model would be useful, especially in experimental scenarios where data may be limited. These and other experimental concerns would likely need to be addressed on an application specific basis.

ACKNOWLEDGMENTS

APPENDIX A: ESTIMATION OF THE DYNAMICS OF POD COEFFICIENTS FROM DATA

Consider the problem of finding an estimate for the matrix $X \in \mathbb{R}^{n \times n}$ as part of the relationship $XA = B$ where the matrices A and $B \in \mathbb{R}^{n \times m}$ are known. Letting $|| \cdot ||_F$ denote the Frobenius norm and $\dagger$ denote the pseudoinverse, here it will be shown that $X = BA^{\dagger}$ provides an estimate that minimizes the Frobenius norm of the residual, i.e., $||BA^{\dagger}A - B||_F \leq ||YA - B||_F$ for any choice of Y. To begin, the relationship $XA = B$ can be rewritten as a matrix equation of the form

$$\begin{bmatrix} A^T & & \\ & \ddots & \\ & & A^T \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}, \quad (A1)$$

where x_j and b_j are the jth columns of X^T and B^T, respectively. As explained in Ref. 41, a least squares estimate for the unknown vector (comprised of elements of X) can be obtained by taking the pseudoinverse of the large matrix. Since the only nonzero entries of the

large matrix are on the diagonal, this can be written as

$$\begin{bmatrix} x_1^{\mathrm{LS}} \\ \vdots \\ x_n^{\mathrm{LS}} \end{bmatrix} = \begin{bmatrix} A^{T\dagger} & & \\ & \ddots & \\ & & A^{T\dagger} \end{bmatrix} \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}, \quad (A2)$$

where $\left[\left(x_1^{\mathrm{LS}} \right)^T \ \ldots \ \left(x_n^{\mathrm{LS}} \right)^T \right]^T$ is the solution to (A1) which minimizes the error of the squared residuals. The relationship (A2) can be rewritten in terms of the original matrices as

$$(X^{\mathrm{LS}})^T = A^{T\dagger} B^T$$
$$= A^{\dagger T} B^T, \quad (A3)$$

and finally

$$X^{\mathrm{LS}} = BA^{\dagger}, \quad (A4)$$

where X^{LS} is the matrix that minimizes the Frobenius norm of the residual as desired.

APPENDIX B: MINIMUM TIME CONTROL OF PHASE-AMPLITUDE REDUCED EQUATIONS

Consider a nonlinear model of the form

$$\dot{x} = R(x) + B(x)u(t), \quad (B1)$$

with $x \in \mathbb{R}^n$, $u(t) \in \mathbb{R}$ is a control input with constraints $u_{\min}(x) \leq u(t) \leq u_{\max}(x)$, $B(x) \in \mathbb{R}^n$ determines the influence of the control stimulus, and $R(x) \in \mathbb{R}^n$ gives the unperturbed dynamics. Suppose the control objective is to find a control input that will bring the state to some target set $\mathcal{T}$ as quickly as possible. To formulate this as an optimal control problem, one can define $t_{\min}(x, u(t))$ to be the minimum time required for an initial state to reach $\mathcal{T}$ under the application of the control signal $u(t)$. For any initial state, x, a minimum-time stimulus, $u^*(t)$, must minimize the cost functional

$$J(x, u(t)) = \int_0^{t_{\mathrm{targ}}} 1 \, dt = t_{\mathrm{targ}}(x, u(t)), \quad (B2)$$

subject to the constraints on $u(t)$. Above, t_{targ} is the time required to reach the target set under the application of $u(t)$. The minimum value of t_{targ} can be found by identifying the control policy that yields $u^*(t)$. This can be accomplished by first defining a value function (sometimes referred to a cost-to-go function)

$$\mathcal{V}(x) = \inf_{u_{\min} \leq u \leq u_{\max}} J(x, u(t)) = \inf_{u_{\min} \leq u \leq u_{\max}} t_{\mathrm{targ}}(x, u(t)). \quad (B3)$$

Using a Hamilton-Jacobi-Bellman framework, the value function $\mathcal{V}(x)$ is the solution to[4]

$$0 = \min_{u_{\min} \leq u \leq u_{\max}} \left[1 + \nabla \mathcal{V}^T(x)(R(x) + B(x)u(t)) \right]$$
$$= 1 + \nabla \mathcal{V}^T(x)R(x) + \min_{u_{\min} \leq u \leq u_{\max}} \left[\nabla \mathcal{V}^T(x)B(x)u(t) \right], \quad (B4)$$

where ∇ denotes the gradient. By solving for $\mathcal{V}(x)$, the optimal control $u^*(t)$ can be found according to the control policy

$$u^*(t) = \arg \min_{u_{\min} \leq u \leq u_{\max}} \left(\nabla \mathcal{V}^T(x)B(x)u \right), \quad (B5)$$

which simplifies to

$$u^*(t) = \begin{cases} u_{\max} & \text{if } \nabla \mathcal{V}^T(x(t))B(x(t)) < 0, \\ u_{\min} & \text{otherwise.} \end{cases} \tag{B6}$$

Numerical solutions of (B4) can be obtained using a numerical package for solving time dependent partial differential equations.[44] Even though (B4) itself is not time dependent, it can be transformed to a suitable form using the following strategy detailed in Ref. 50 and summarized in Ref. 44. First, one can define an intermediate function

$$\xi(x, \nabla \mathcal{V}(x)) = 1 + \nabla \mathcal{V}^T(x)R(x) + \min_{u_{\min} \leq u \leq u_{\max}} \left[\mathcal{V}^T(x)B(x)u(t) \right], \tag{B7}$$

so that (B4) can be rewritten as

$$\xi(x, \nabla \mathcal{V}(x)) = 0 \quad \text{on } \mathcal{D}, $$
$$\mathcal{V}(x) = 0 \quad \text{on } \mathcal{D} \backslash \delta \mathcal{T}, \tag{B8}$$

where $\delta \mathcal{T}$ is the boundary of the target set and $\mathcal{D}$ is the domain of interest. Provided the boundary conditions are noncharacteristic, the value function can be found by introducing a new function $\Upsilon(x, t)$ and solving the equation

$$0 = \Upsilon_t(x, t) + \nabla \Upsilon(x, t)^T R(x) + \min_{u_{\min} \leq u \leq u_{\max}} \left[\nabla \Upsilon(x, t)^T B(x)u(t) \right], \tag{B9}$$

using an initial condition

$$\Upsilon(x, 0) = 0 \in \delta \mathcal{T},$$
$$\Upsilon(x, 0) < 0 \in \mathcal{T} \backslash \delta \mathcal{T}, \tag{B10}$$
$$\Upsilon(x, 0) > 0 \in \mathcal{D} \backslash \mathcal{T}.$$

The initial condition $\Upsilon(x, 0)$ is usually taken to be a signed function of distance to the boundary of $\delta \mathcal{T}$. Upon solving for $\Upsilon(x, t)$ the value function can be found by taking

$$\mathcal{V}(x) = \{t | \Upsilon(x, t) = 0\}. \tag{B11}$$

Once $\mathcal{V}(x)$ is known, it is straightforward to determine the optimal control according to (B6).

APPENDIX C: FAILURE OF STANDARD PHASE-ONLY REDUCTION METHODS FOR LARGE MAGNITUDE PERTURBATIONS

Standard phase reduction using (2) fails for the circadian model (49) used in Sec. VI. When attempting to use the reduction $\dot{\theta} = \omega + Z(\theta)^T U(t)$ to represent the dynamics of (49), recall that $Z(\theta) \equiv \frac{\partial \theta}{\partial x}$, where x is the overall state of the system and partial derivatives are evaluated on the limit cycle $x^\gamma(\theta)$. Intuitively, $Z(\theta)$ provides a linear approximation of the gradient of the phase coordinate evaluated on the periodic orbit.

Toward a standard implementation of a phase-only reduction on (49), one can nominally take $L_{\text{tot}} \equiv L(t_s) + \Delta L(t) = 0$ and treat L_{tot} as a parameter perturbation. From this perspective, one can write a phase reduced version of this model as

$$\dot{\theta} = \omega + Z_L^{x^\gamma}(\theta)(L(t_s) + \Delta L(t)), \tag{C1}$$

where

$$Z_L^{x^\gamma}(\theta) \equiv \sum_{i=1}^N \left(\frac{\partial \theta}{\partial a_i} S_i \right) \bigg|_{x^\gamma(\theta)}, \tag{C2}$$

where $x^\gamma(\theta)$ is the limit cycle that results when taking $L(t_s) + \Delta L(t) = 0$. Recall that S_i in (C2) is the sensitivity of oscillator i to light. As illustrated in Fig. 8, the phase reduction (C1) breaks down for the magnitudes of light perturbations used in this model.

Panel (a) of Fig. 8 shows $Z_L^{x^\gamma}(\theta)$ estimated using the direct method.[22,30,47] Briefly, (49) is simulated in the absence of light (i.e., taking $L(t_s) + \Delta L(t) = 0$) until initial transients die out so that the state is close to the limit cycle. At a known phase, $\Delta L(t)$ is set to 0.015 for a duration of 1 h, the resulting shift in phase $\Delta \theta$ is measured, and the phase response curve $Z_L^{x^\gamma}(\theta)$ is taken to be $\Delta \theta / 0.015$. This process is repeated 450 times with each red dot in panel (a) representing the result of a single trial. The red line shows a least square fit to a basis of the form $\sum_{n=0}^2 \left[a_n \sin(n\theta) + b_n \cos(n\theta) \right]$. This fit represents an approximation of the phase response curve that would be used in (2). A second set of simulations is also performed to measure $Z_L^{x^{\text{ent}}}(\theta)$, that is, the effective phase response curve for an initial condition on the orbit that results when (49) is fully entrained to $L(t_s)$ as defined by (50). Equation (49) is once again simulated with this
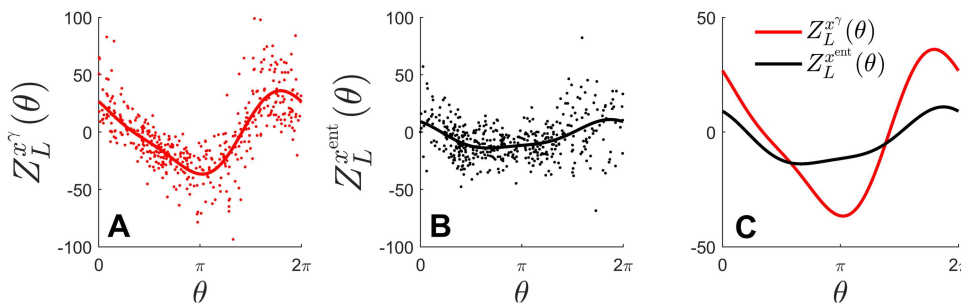

FIG. 8. Panels (a) and (b) show results when using the direct method to estimate the influence of light perturbations on the asymptotic phase. Panel (a) shows the phase response curve as would be used in the phase reduction from (C1). Panel (b) shows the effective response for initial conditions taken on the entrained orbit. Panel (c) gives a direct comparison between the two curves. The phase reduction (C1) does not accurately reflect the phase dynamics on the entrained orbit and hence does not provide a viable reduction strategy to study entrainment for this model.

light-dark cycle taking $\Delta L = 0$. Once all transients have died out so that the system is fully entrained to the light-dark cycle, the state at a randomly chosen value of t_s is saved in memory. The applied light $L(t_s) + \Delta L(t)$ is set to zero and the system is simulated long enough so that it can approach the limit cycle. The phase θ at the moment the light-dark cycle was eliminated can then be inferred. Performing an associated simulation using the state that was stored in memory as an initial condition, $L(t_s)$ is set to zero and $\Delta L(t)$ is set to 0.015 for 1 h and then set to zero for the remainder of the simulation. These data are used to infer the change in phase $\Delta\theta$ caused by the 1 h perturbation and $Z_L^{x^{ent}}(\theta)$ is taken to be $\Delta\theta/0.015$. This process is repeated 450 times with each black dot representing the result from a single trial. The black line is fit to the same basis as before. This curve gives a local approximation of the gradient of the phase on the periodic orbit that is entrained to the light-dark cycle. Panel (c) shows the two curves plotted on the same axis; the difference is stark with the phase response curve severely overestimating the effect of light perturbation for a state on the entrained periodic orbit.

As explained in the text, the control strategy presented in Sec. VI avoids the above limitation by treating t_s as a state variable so that the system is viewed as an autonomous differential equation. From this perspective, the entrained periodic orbit itself is the limit cycle and subsequent perturbations $\Delta L(t)$ are small enough so that the dynamics stay close enough to the limit cycle. Note that using this formulation, the asymptotic phase is identically equal to $2\pi t_s/24$. Subsequently, $\Delta L(t)$ cannot be used to influence the phase. Instead, the magnitude of the isostable coordinates corresponds to the degree of circadian misalignment. These coordinates can be influenced by $\Delta L(t)$.

REFERENCES

[1]J. Arendt, "Managing jet lag: Some of the problems and possible new solutions," Sleep Med. Rev. **13**(4), 249–256 (2009).
[2]R. Azodi-Avval and A. Gharabaghi, "Phase-dependent modulation as a novel approach for therapeutic brain stimulation," Front. Comput. Neurosci. **9**, 26 (2015).
[3]N. Bagheri, J. Stelling, and F. J. Doyle III, "Circadian phase resetting via single and multiple control targets," PLoS Comput. Biol. **4**(7), e1000104 (2008).
[4]M. Bardi and I. Capuzzo-Dolcetta, Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations (Birkhauser, Boston, 2008).
[5]R. E. Bellman, Adaptive Control Processes: A Guided Tour (Princeton University Press, 1961), Vol. 2045.
[6]P. Benner, S. Gugercin, and K. Willcox, "A survey of projection-based model reduction methods for parametric dynamical systems," SIAM Rev. **57**(4), 483–531 (2015).
[7]R. E. Bertram and J. E. Rubin, "Multi-timescale systems and fast-slow analysis," Math. Biosci. **287**, 105–121 (2017).
[8]W. B. Bomela, I. S. Dasanayake, J. S. Li, Y. Chen, and I. Z. Kiss, "Optimal phase-to-phase control of chemical oscillations," Ind. Eng. Chem. Res. **57**(23), 7764–7770 (2018).
[9]E. Brown, J. Moehlis, and P. Holmes, "On the phase reduction and response dynamics of neural oscillator populations," Neural Comput. **16**(4), 673–715 (2004).
[10]M. Budišić, R. Mohr, and I. Mezić, "Applied Koopmanism," Chaos **22**(4), 047510 (2012).
[11]A. L. Chesson, M. Littner, D. Davila, W. M. Anderson, M. Grigg-Damberger, K. Hartse, S. Johnson, and M. Wise, "Practice parameters for the use of light therapy in the treatment of sleep disorders," Sleep **22**(5), 641–660 (1999).
[12]P. Constantin, C. Foias, B. Nicolaenko, and R. Temam, Integral Manifolds and Inertial Manifolds for Dissipative Partial Differential Equations (Springer, New York, 1989), Vol. 70.
[13]J. Cui, C. C. Canavier, and R. J. Butera, "Functional phase response curves: A method for understanding synchronization of adapting neurons," J. Neurophysiol. **102**(1), 387–398 (2009).
[14]C. A. Czeisler, J. F. Duffy, T. L. Shanahan, E. N. Brown, J. F. Mitchell, D. W. Rimmer, J. M. Ronda, E. J. Silva, J. S. Allan, J. S. Emens, D. J. Dijk, and R. E. Kronauer, "Stability, precision, and near-24-hour period of the human circadian pacemaker," Science **284**(5423), 2177–2181 (1999).
[15]D. A. Dean, D. B. Forger, and E. B. Klerman, "Taking the lag out of jet lag through model-based schedule design," PLoS Comput. Biol. **5**(6), e1000418 (2009).
[16]C. O. Diekman and A. Bose, "Entrainment maps: A new tool for understanding properties of circadian oscillator models," J. Biol. Rhythms **31**(6), 598–616 (2016).
[17]C. O. Diekman and A. Bose, "Reentrainment of the circadian pacemaker during jet lag: East-west asymmetry and the effects of north-south travel," J. Theor. Biol. **437**, 261–285 (2018).
[18]G. B. Ermentrout and D. H. Terman, Mathematical Foundations of Neuroscience (Springer, New York, 2010), Vol. 35.
[19]C. Foias, G. R. Sell, and R. Temam, "Inertial manifolds for nonlinear evolutionary equations," J. Differ. Equ. **73**(2), 309–353 (1988).
[20]D. B. Forger and C. S. Peskin, "A detailed predictive model of the mammalian circadian clock," Proc. Natl. Acad. Sci. U.S.A. **100**(25), 14806–14811 (2003).
[21]J. Foss and J. Milton, "Multistability in recurrent neural loops arising from delay," J. Neurophysiol. **84**(2), 975–985 (2000).
[22]R. F. Galán, G. B. Ermentrout, and N. N. Urban, "Efficient estimation of phase-resetting curves in real neurons and its significance for neural-network modeling," Phys. Rev. Lett. **94**(15), 158101 (2005).
[23]C. W. Gardiner, Handbook of Stochastic Methods: For Physics, Chemistry and the Natural Sciences (Springer, Berlin, 2004).
[24]D. A. Golombek and R. E. Rosenstein, "Physiology of circadian entrainment," Physiol. Rev. **90**(3), 1063–1102 (2010).
[25]D. Gonze, S. Bernard, C. Waltermann, A. Kramer, and H. Herzel, "Spontaneous synchronization of coupled circadian oscillators," Biophys. J. **89**(1), 120–129 (2005).
[26]R. Grimshaw, Nonlinear Ordinary Differential Equations (CRC Press, Baca Raton, FL, 1993), Vol. 2.
[27]J. Guckenheimer, "Isochrons and phaseless sets," J. Math. Biol. **1**(3), 259–273 (1975).
[28]P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, Turbulence, Coherent Structures, Dynamical Systems and Symmetry (Cambridge University Press, New York, 1996).
[29]A. B. Holt, D. Wilson, M. Shinn, J. Moehlis, and T. I. Netoff, "Phasic burst stimulation: A closed-loop approach to tuning deep brain stimulation parameters for Parkinson's disease," Comput. Biol. **12**(7), e1005011 (2016).
[30]E. M. Izhikevich, Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting (MIT Press, London, 2007).
[31]D. Johnston and S. M.-S. Wu, Foundations of Cellular Neurophysiology (MIT Press, Cambridge, MA, 1995).
[32]D. Jordan and P. Smith, Nonlinear Ordinary Differential Equations: An Introduction for Scientists and Engineers (Oxford University Press, Oxford, 2007), Vol. 10.
[33]S. B. S. Khalsa, M. E. Jewett, C. Cajochen, and C. A. Czeisler, "A phase response curve to single bright light pulses in human subjects," J. Physiol. (Lond.) **549**(3), 945–952 (2003).
[34]J. K. Kim and D. B. Forger, "A mechanism for robust circadian timekeeping via stoichiometric balance," Mol. Syst. Biol. **8**(1), 630 (2012).
[35]D. Kirk, Optimal Control Theory (Dover Publications, New York, 1998).
[36]I. Z. Kiss, C. G. Rusin, H. Kori, and J. L. Hudson, "Engineering complex dynamical structures: Sequential patterns and desynchronization," Science **316**, 1886–1889 (2007).
[37]Y. Kuramoto, Chemical Oscillations, Waves, and Turbulence (Springer-Verlag, Berlin, 1984).
[38]J. C. Leloup and A. Goldbeter, "Toward a detailed computational model for the mammalian circadian clock," Proc. Natl. Acad. Sci. U.S.A. **100**(12), 7051–7056 (2003).

[39] B. Letson and J. E. Rubin, "A new frame for an old (phase) portrait: Finding rivers and other flow features in the plane," J. Appl. Dyn. Syst. **17**(4), 2414–2445 (2018).

[40] A. Mauroy, I. Mezić, and J. Moehlis, "Isostables, isochrons, and Koopman spectrum for the action–angle representation of stable fixed point dynamics," Physica D **261**, 19–30 (2013).

[41] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra* (SIAM, Philadelphia, 2000), Vol. 2.

[42] I. Mezić, "Analysis of fluid flows via spectral properties of the Koopman operator," Annu. Rev. Fluid Mech. **45**, 357–378 (2013).

[43] D. S. Minors, J. M. Waterhouse, and A. Wirz-Justice, "A human phase-response curve to light," Neurosci. Lett. **133**(1), 36–40 (1991).

[44] I. Mitchell, "A toolbox of level set methods," Technical Report TR-2007-11, University of British Columbia, Vancouver BC, 2007, see http://www.cs.ubs.ca/˜mitchell/ToolboxLS/toolboxLS.pdf.

[45] B. Monga, D. Wilson, T. Matchen, and J. Moehlis, "Phase reduction and phase-based optimal control for biological systems: A tutorial," Biol. Cybern. **113**(1–2), 11–46 (2019).

[46] A. Nabi, T. Stigen, J. Moehlis, and T. Netoff, "Minimum energy control for *in vitro* neurons," J. Neural Eng. **10**(3), 036005 (2013).

[47] T. Netoff, M. A. Schwemmer, and T. J. Lewis, "Experimentally estimating phase response curves of neurons: Theoretical and practical issues," in *Phase Response Curves in Neuroscience* (Springer, 2012), pp. 95–129.

[48] T. I. Netoff, C. D. Acker, J. C. Bettencourt, and J. A. White, "Beyond two-cell networks: Experimental measurement of neuronal responses to multiple synaptic inputs," J. Comput. Neurosci. **18**(3), 287–295 (2005).

[49] University of Michigan, Entrain Yourself, see http://entrain.math.lsa.umich.edu (last accessed September 19, 2018).

[50] S. Osher, "A level set formulation for the solution of the Dirichlet problem for Hamilton-Jacobi equations," SIAM J. Math. Anal. **24**(5), 1145–1152 (1993).

[51] S. Panda, J. B. Hogenesch, and S. A. Kay, "Circadian rhythms from flies to human," Nature **417**(6886), 329 (2002).

[52] Y. Park and B. Ermentrout, "Weakly coupled oscillators in a slowly varying world," J. Comput. Neurosci. **40**(3), 269–281 (2016).

[53] K. Pyragas and V. Novičenko, "Phase reduction of a limit cycle oscillator perturbed by a strong amplitude-modulated high-frequency force," Phys. Rev. E **92**(1), 012910 (2015).

[54] S. M. Reppert and D. R. Weaver, "Coordination of circadian timing in mammals," Nature **418**(6901), 935 (2002).

[55] C. W. Rowley and S. T. M. Dawson, "Model reduction for flow analysis and control," Annu. Rev. Fluid Mech. **49**, 387–417 (2017).

[56] J. Rubin and D. Terman, "High frequency stimulation of the subthalamic nucleus eliminates pathological thalamic rhythmicity in a computational model," J. Comput. Neurosci. **16**, 211–235 (2004).

[57] R. L. Sack, "Jet lag," New Eng. J. Med. **362**(5), 440–447 (2010).

[58] J. A. Sanders, F. Verhulst, and J. Murdock, *Averaging Methods in Nonlinear Dynamical Systems*, 2nd ed. (Springer-Verlag, New York, 2007).

[59] P. J. Schmid, "Dynamic mode decomposition of numerical and experimental data," J. Fluid Mech. **656**, 5–28 (2010).

[60] K. Serkh and D. B. Forger, "Optimal schedules of light exposure for rapidly correcting circadian misalignment," PLoS Comput. Biol. **10**(4), e1003523 (2014).

[61] S. Shirasaka, W. Kurebayashi, and H. Nakao, "Phase-amplitude reduction of transient dynamics far from attractors for limit-cycling systems," Chaos **27**(2), 023119 (2017).

[62] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley, "Modal analysis of fluid flows: An overview," AIAA J. **55**(12), 4013–4041 (2017).

[63] J. Waterhouse, T. Reilly, G. Atkinson, and B. Edwards, "Jet lag: Trends and coping strategies," Lancet **369**(9567), 1117–1129 (2007).

[64] K. C. A. Wedgwood, K. K. Lin, R. Thul, and S. Coombes, "Phase-amplitude descriptions of neural oscillator models," J. Math. Neurosci. **3**(1), 2 (2013).

[65] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, "A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition," J. Nonlinear Sci. **25**(6), 1307–1346 (2015).

[66] D. Wilson, "Isostable reduction of oscillators with piecewise smooth dynamics and complex Floquet multipliers," Phys. Rev. E **99**(2), 022210 (2019).

[67] D. Wilson and S. Djouadi, "Isostable reduction and boundary feedback control for nonlinear convective flow," in *Proceedings of the 58th IEEE Conference on Decision and Control* (IEEE, 2019).

[68] D. Wilson and B. Ermentrout, "Greater accuracy and broadened applicability of phase reduction using isostable coordinates," J. Math. Biol. **76**(1–2), 37–66 (2018).

[69] D. Wilson and B. Ermentrout, "An operational definition of phase characterizes the transient response of perturbed limit cycle oscillators," J. Appl. Dyn. Syst. **17**(4), 2516–2543 (2018).

[70] D. Wilson and B. Ermentrout, "Augmented phase reduction of (not so) weakly perturbed coupled oscillators," SIAM Rev. **61**(2), 277–315 (2019).

[71] D. Wilson and J. Moehlis, "Isostable reduction of periodic orbits," Phys. Rev. E **94**(5), 052213 (2016).

[72] A. Winfree, *The Geometry of Biological Time*, 2nd ed. (Springer Verlag, New York, 2001).

[73] K. P. Wright, R. J. Hughes, R. E. Kronauer, D. J. Dijk, and C. A. Czeisler, "Intrinsic near-24-h pacemaker period determines limits of circadian entrainment to a weak synchronizer in humans," Proc. Natl. Acad. Sci. U.S.A. **98**(24), 14027–14032 (2001).

[74] M. W. Young and S. A. Kay, "Time zones: A comparative genetics of circadian clocks," Nat. Rev. Genet. **2**(9), 702 (2001).

[75] J. M. Zeitzer, D. J. Dijk, R. E. Kronauer, E. N. Brown, and C. A. Czeisler, "Sensitivity of the human circadian pacemaker to nocturnal light: Melatonin phase resetting and suppression," J. Physiol. (Lond.) **526**(3), 695–702 (2000).

[76] A. Zlotnik, Y. Chen, I. Z. Kiss, H. A. Tanaka, and J. S. Li, "Optimal waveform for fast entrainment of weakly forced nonlinear oscillators," Phys. Rev. Lett. **111**(2), 024102 (2013).

[77] A. Zlotnik, R. Nagao, I. Z. Kiss, and J. S. Li, "Phase-selective entrainment of nonlinear oscillator ensembles," Nat. Commun. **7**, 10788 (2016).