Annual Reviews in Control 54 (2022) 327–351

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

Review article

Recent advances in the analysis and control of large populations of neural oscillators

Dan Wilson [a,*], Jeff Moehlis [b]

[a] Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, TN 37996, USA
[b] Department of Mechanical Engineering and Program in Dynamical Neuroscience, University of California, Santa Barbara, CA 93106, USA

ARTICLE INFO

Keywords:
Deep brain stimulation
Neural control
Phase
Phase reduction
Isochrons
Isostable coordinates
Optimal control
Synchronization
Entrainment
Phase resetting

ABSTRACT

Many challenging problems that consider the analysis and control of neural brain rhythms have been motivated by the advent of deep brain stimulation as a therapeutic treatment for a wide variety of neurological disorders. In a computational setting, neural rhythms are often modeled using large populations of coupled, conductance-based neurons. Control of such models comes with a long list of challenges: the underlying dynamics are nonnegligibly nonlinear, high dimensional, and subject to noise; hardware and biological limitations place restrictive constraints on allowable inputs; direct measurement of system observables is generally limited; and the resulting systems are typically highly underactuated. In this review article, we highlight a collection of recent analysis techniques and control frameworks that have been developed to contend with these difficulties. Particular emphasis is placed on the problem of desynchronization for a population of pathologically synchronized neural oscillators, a problem that is motivated by applications to Parkinson's disease where pathological synchronization is thought to contribute to the associated motor control symptoms. We also discuss other recent neural control applications that consider entrainment, phase randomization, synchronization, and clustering.

1. Introduction

Deep brain stimulation (DBS) is an FDA approved medical treatment used by over 100,000 patients worldwide (Lozano & Lipsman, 2013) whereby high frequency electrical current pulses are injected into an appropriate brain region. DBS was initially developed as a treatment for medication resistant Parkinson's disease (Benabid, Pollak, Louveau, Henry, & Rougemont, 1987; Koller, et al., 1997). More recently, DBS has shown promise as a potential treatment for a variety of other neurological disorders including depression (Mayberg, et al., 2005), Alzheimer's disease (Laxton, et al., 2010), Tourette syndrome (Schrock, et al., 2015), epilepsy (Li & Cook, 2018), and obsessive-compulsive disorder (Greenberg, et al., 2006). From a purely electrical perspective, the action of DBS is relatively well-understood, with DBS pulses inducing some combination of activation or inhibition in the neurons adjacent to the probe (Anderson, Farokhniaee, Gunalan, Howell, & McIntyre, 2018; McIntyre, Mori, Sherman, Thakor, & Vitek, 2004). By contrast, the dynamical mechanisms that influence the aggregate behavior of large populations of neurons are not well understood, leading to many open questions and control problems associated with DBS as a therapeutic treatment for a variety of neurological disorders.

In this review article, we highlight a collection of techniques that have been developed to handle the difficulties associated with analysis and control of neural brain rhythms. Specifically, this review is focused on control of aggregate oscillations that emerge in large populations of coupled neurons which are often modeled using conductance-based ordinary differential equations (ODEs) to describe the flow of current across a cell membrane (Ermentrout & Terman, 2010). From the perspective of classical feedback control, the control design problem for neural populations is particularly challenging: conductance-based ODE models of spiking neurons are complex and highly nonlinear, one wishes to control a large number of neurons using a single control stimulus, the magnitude of the stimulus is constrained by biological and hardware implementation limitations, voltage is typically the only directly measurable state, and such measurements are quite noisy. In these applications, techniques such as feedback linearization (Khalil, 2002) are typically not useful. Additionally, because of the high-dimensional nature of these control problems, the direct implementation of nonlinear control techniques is computationally prohibitive in many practical applications.

* Corresponding author.
E-mail addresses: dwilso81@utk.edu (D. Wilson), moehlis@ucsb.edu (J. Moehlis).

https://doi.org/10.1016/j.arcontrol.2022.05.002
Received 15 December 2021; Accepted 19 May 2022
Available online 11 June 2022

Because of these aforementioned difficulties, it is often useful to work in a coordinate framework that is more amenable to both mathematical analysis and implementation of nonlinear control algorithms. With respect to tonically firing neurons (i.e., those with a nominal stable limit cycle), Art Winfree made great strides in developing such a coordinate framework by introducing the notion of an asymptotic phase (Winfree, 1967, 2001). The key insight here was that the behavior of a general limit cycle oscillator in response to perturbations could be characterized in terms of the timing of oscillations rather than in reference to the underlying state variables. In subsequent years, these ideas were formalized using the notion of isochrons (Brown, Moehlis, & Holmes, 2004; Guckenheimer, 1975; Kuramoto, 1984; Winfree, 2001), i.e., level sets of initial conditions that have the same asymptotic convergence to the limit cycle. By restricting one's attention to a small neighborhood of the limit cycle, one can obtain a 1-dimensional phase reduced equation (Ermentrout & Terman, 2010; Izhikevich, 2007; Kuramoto, 1984) that can accurately capture the response to weak perturbations. Such phase reductions have been used for upwards of half a century to understand complex emergent patterns in weakly perturbed neurons and other oscillators (Brown, Holmes, & Moehlis, 2003; Ermentrout & Kopell, 1991; Kuramoto, 1997; Moehlis, Shea-Brown, & Rabitz, 2006; Wilson & Moehlis, 2014c; Winfree, 1974). More recent applications have highlighted the utility of additionally considering the amplitude dynamics in a phase reduced coordinate framework in order to extend applicability beyond the weakly perturbed paradigm (Castejón, Guillamon, & Huguet, 2013; Letson & Rubin, 2020; Shirasaka, Kurebayashi, & Nakao, 2017; Wedgwood, Lin, Thul, & Coombes, 2013; Wilson, 2020c, 2022; Wilson & Ermentrout, 2018b).

In the context of control of neural populations and neural rhythms, there are a variety of objectives to consider. A significant focus in this review article is on the desynchronization of a pathologically synchronized population of neural oscillators. These works are primarily motivated by applications to Parkinson's disease. Evidence suggests that pathological synchronization among neurons in the basal ganglia contributes to the motor control symptoms of Parkinson's disease, and that DBS helps to restore normal function by disrupting this synchronization (Chen, et al., 2007; Kuhn, et al., 2008; Levy, Hutchison, Lozano, & Dostrovsky, 2000). In pursuit of this control objective, a variety of model order reduction, analysis, and control techniques have been developed in recent years that can be applied to large populations of coupled, periodically firing neurons. Other objectives including entrainment, phase randomization, synchronization, and clustering also have relevance in neural control applications, and are discussed as part of the applications considered in this review article.

The organization of this article is as follows: Section 2 gives necessary background information on standard phase-based model order reduction techniques which provide a foundation for the implementation of many of the control techniques highlighted in this review. Appendix A expands on these methods with a discussion of isostable coordinate-based phase–amplitude reduction methods that are used in many of the applications considered in the latter portions of this article. Section 3 discusses reduction techniques for large populations of nonlinear oscillators. Model independent control techniques are discussed in Section 4. Such strategies can be particularly useful when the underlying model is unknown and limited information about the underlying system is available. Sections 5–8 consider problems where the goal is desynchronization of a pathologically synchronized neural rhythm: Section 5 discusses the method of chaotic desynchronization that can be used to manipulate finite time Lyapunov exponents so that nearby trajectories diverge, Section 6 considers approaches that target clustered states, e.g., that split the population into distinct clusters, and Section 7 discusses the direct control of phase distributions. Section 8 considers phase randomization techniques to drive a collection of oscillators to a phaseless set where they can subsequently be desynchronized by inherent noise. A collection of open-loop control strategies are considered in Section 9 for use when real-time state measurements are unavailable. Section 10 discusses a set of control problems where entrainment is the control objective, as is often the case when considering circadian oscillations. Finally, Section 12 provides concluding remarks.

2. Isochrons and phase reduction

The ordinary differential equations used to describe the dynamics of large populations of neurons are typically high-dimensional and nonnegligibly nonlinear. Both of these factors generally preclude direct analysis of the equations associated with these models. When considering oscillatory dynamics, phase reduction is one strategy that is often used to understand complex emergent patterns in weakly perturbed oscillators (Brown et al., 2003; Ermentrout & Kopell, 1991; Kuramoto, 1997; Moehlis et al., 2006; Wilson & Moehlis, 2014c; Winfree, 1974). Here, we provide a brief review of the notion of phase for oscillatory dynamical systems (such as a tonically firing neuron) and give a description of the phase reduction strategy that can be used to characterize the behavior of limit cycle oscillators in terms of the timing of oscillations (rather than their full state dynamics).

While phase reduction is a powerful tool, we note that it is only valid in the weakly perturbed regime, i.e., when considering inputs that are small enough so that the state remains close to its underlying limit cycle. In applications requiring larger magnitude inputs, information about amplitude-based effects must be incorporated. Appendix A provides a detailed discussion on various phase–amplitude reduction methods that use isostable coordinates to encode for these amplitude-based effects. These methods are employed in applications considered in the latter parts of this article.

2.1. Isochrons

To begin, consider an autonomous vector field

$$\frac{dx}{dt} = F(x), \qquad x \in \mathbb{R}^n, \qquad (n \geq 2) \tag{1}$$

having a stable hyperbolic periodic orbit $x^\gamma(t)$ with period T. We define the set of all points in the basin of attraction as $\mathcal{B}$. For each point x^* in $\mathcal{B}$ there exists a unique $\theta(x^*)$ such that (Coddington & Levinson, 1955; Guckenheimer, 1975; Josic, Shea-Brown, & Moehlis, 2006; Kuramoto, 1984; Malkin, 1949; Winfree, 1967, 1974, 2001)

$$\lim_{t\to\infty} \left| x(t) - x^\gamma\left(t + \frac{T}{2\pi}\,\theta(x^*)\right)\right| = 0, \tag{2}$$

where $x(t)$ is a trajectory starting with the initial point x^*. The function $\theta(x)$ is called the *asymptotic phase* of x, and takes values in $[0, 2\pi)$. Other conventions, related to this through a simple rescaling, define the asymptotic phase to take values in $[0, T)$ or in $[0, 1)$.

Let x_0^γ be the point on the periodic orbit where the phase is zero. A common convention is to choose x_0^γ as corresponding to the global maximum of the first coordinate on the periodic orbit, although, other conventions can be used (for instance having x_0^γ corresponding to the crossing of an arbitrarily chosen Poincaré section). An *isochron* is defined as a level set of $\theta(x)$, that is, the collection of all points in the basin of attraction of x^γ with the same asymptotic phase (Winfree, 1967, 2001). By convention, we will denote Γ_ζ as the set of all x for which $\theta(x) = \zeta$. We note that if $x(0)$ is a point on a periodic orbit, the isochron associated with that point is the set of all initial conditions $\tilde{x}(0)$ such that $\|x(t) - \tilde{x}(t)\| \to 0$ as $t \to \infty$. Isochrons extend the notion of phase of a stable periodic orbit to the basin of attraction of the periodic orbit. It is conventional to define isochrons so that the phase of a trajectory on the periodic orbit advances at a constant rate according to

$$\frac{d\theta}{dt} = \frac{2\pi}{T} \equiv \omega \tag{3}$$

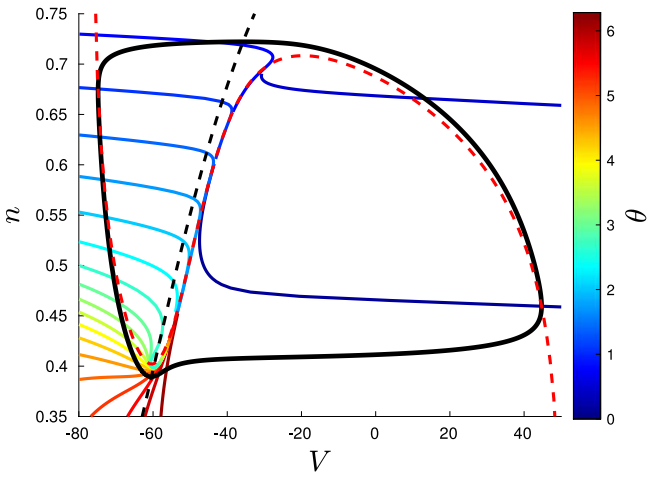

Fig. 1. Periodic orbit (black solid), V-nullcline (red dashed), and n-nullcline (black dashed), and twenty isochrons equally spaced in phase (as indicated by color) for the two-dimensional reduced Hodgkin–Huxley model in the absence of noise, coupling, and control. The location of the unstable fixed point (phaseless point) for this system is at the intersection of the nullclines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

both on and off the periodic orbit. Points at which isochrons of a periodic orbit cannot be defined form the *phaseless set* (Winfree, 1974), which will be important for the phase randomization control strategies discussed in Section 8. Isochrons for a single neuron from the two-dimensional reduced Hodgkin–Huxley model (Rinzel, 1985) given in Appendix B are shown in Fig. 1, where $x = [V, n]$.

Isochrons can be shown to exist for any stable hyperbolic periodic orbit. They are codimension one manifolds as smooth as the vector field, and transversal to the periodic orbit x^γ. Their union covers an open neighborhood of x^γ. This can be proved directly by using the Implicit Function Theorem (Coddington & Levinson, 1955; Guckenheimer, 1975), and is also a consequence of results on normally hyperbolic invariant manifolds (Wiggins, 1994). Methods for calculating isochrons numerically are summarized in Monga, Wilson, Matchen, and Moehlis (2019).

2.2. Phase reduction and phase-based control

Control theory seeks to design inputs to a dynamical system which change its behavior in a desired way. With this in mind, we consider the perturbed system

$$\frac{dx}{dt} = F(x) + U(x,t), \tag{4}$$

where $U(x,t)$ is a small control input. The evolution of this system in terms of isochrons is (Brown et al., 2004; Kuramoto, 1984)

$$\begin{aligned}\frac{d\theta}{dt} &= \frac{\partial\theta}{\partial x}\cdot\frac{dx}{dt}\\ &= \frac{\partial\theta}{\partial x}\cdot(F(x)+U(x,t))\\ &= \omega + \frac{\partial\theta}{\partial x}\cdot U(x,t).\end{aligned} \tag{5}$$

Evaluating on the periodic orbit x^γ for the unperturbed system gives, to leading order,

$$\frac{d\theta}{dt} = \omega + Z(\theta)\cdot U(x^\gamma,t), \tag{6}$$

where

$$Z(\theta) = \left.\frac{\partial\theta}{\partial x}\right|_{x^\gamma(\theta)} \equiv \nabla_{x^\gamma}\theta. \tag{7}$$

Here $Z(\theta)\in\mathbb{R}^n$ is the gradient of phase variable θ evaluated on the periodic orbit, and is referred to as the *(infinitesimal) phase response curve (PRC)* (Ermentrout, 1996; Ermentrout & Terman, 2010; Hansel, Mato, & Meunier, 1995; Netoff, Schwemmer, & Lewis, 2012; Winfree, 2001). It quantifies the effect of a small magnitude external perturbation on the phase of a periodic orbit. We call (6) the *standard phase reduction*. Further discussion of phase response curves is given in Monga et al. (2019).

Since phase-reduced models have lower dimension than the full models from which they came, optimal control problems for phase-reduced models give lower-dimensional boundary value problems, and thus are simpler to solve. Many applications in this article will consider control problems for which the control input is a rank-1 perturbation, i.e., $U(x^\gamma,t) = \delta u(t)$ where $\delta\in\mathbb{R}^n$ and $u(t)\in\mathbb{R}$ is the control stimulus. When δ has only one non-zero entry, this captures input that only affects a single state variable, for instance, control in the form of an injected electrical current which only affects the equation for the transmembrane voltage. Additionally, if the control input only depends on time, and not on the state variables, the standard phase reduction becomes (see, for example, Brown et al., 2004)

$$\frac{d\theta}{dt} = \omega + z(\theta)u(t), \tag{8}$$

where $z(\theta) = Z(\theta)\cdot\delta$.

3. Reduction techniques for coupled oscillator populations

3.1. Phase difference coupling

One of the most well-studied models describing coupled population oscillations is the Kuramoto model (Kuramoto, 1984), (cf., Strogatz, 2000):

$$\frac{d\theta_k}{dt} = \omega_k + \frac{K}{N}\sum_{j=1}^{N}\sin(\theta_j-\theta_k), \quad k = 1,\dots,N, \tag{9}$$

where θ_k is the phase of oscillator k with natural frequency ω_k with frequencies distributed according to some probability distribution $g(\omega)$ that is symmetric about its mean frequency, and K is the coupling strength. Many mean-field models, including the Kuramoto model (9), consider phase difference coupling despite the fact that this formulation rarely explicitly occurs in nature. Nonetheless, phase difference coupling can be used as a reasonable approximation for models of the form (4) provided coupling is sufficiently weak. To illustrate this, consider two coupled oscillators

$$\frac{dx_j}{dt} = F(x_j) + \epsilon G_j(x_j,x_k), \tag{10}$$

where $j = 1,2$ and $k = 3-j$ with $0<\epsilon\ll 1$. Supposing that $dx/dt = F(x)$ has a stable limit cycle $x^\gamma(t)$, changing to phase coordinates using the strategy described in Section 2.2 yields

$$\frac{d\theta_j}{dt} = \omega + \epsilon Z(\theta_j)\cdot G_j(x^\gamma(\theta_j),x^\gamma(\theta_k)). \tag{11}$$

Above, it is assumed that $x_j\approx x^\gamma(\theta_j)$ because the coupling is small. Defining $\phi_j = \theta_j - \omega t$ yields

$$\frac{d\phi_j}{dt} = \epsilon Z(\phi_j+\omega t)\cdot G_j(x^\gamma(\phi_j+\omega t),x^\gamma(\phi_k+\omega t)). \tag{12}$$

Because Z and x^γ are both T-periodic, we can apply the method of averaging (Guckenheimer & Holmes, 1983) to yield

$$\begin{aligned}\frac{d\phi_1}{dt} &= \epsilon H_1(\phi_2-\phi_1),\\ \frac{d\phi_2}{dt} &= \epsilon H_2(\phi_1-\phi_2),\end{aligned} \tag{13}$$

where Eq. (13) provides a good approximation to (12) taking $H_j(\phi) = \frac{1}{T}\int_0^T Z(\omega t)\cdot G_j(x^\gamma(t),x^\gamma(t+\phi))dt$. Intuitively, the method of averaging is used here to approximate the solution of the weakly perturbed, time periodic ordinary differential equation (12) by an average system (13). This general approach is valid in the limit that ϵ is small and yields a

time-invariant set of equations that usually simplifies the subsequent analysis. Noting that $\phi_2 - \phi_1 = \theta_2 - \theta_1$, the desired form of phase difference coupling can be recovered with an appropriate change of variables. Note that the derivation above is for two oscillators, but can be readily extended to accommodate an arbitrarily large population.

Transformations that consider phase difference coupling provide a useful starting point for assessing the emergence of spontaneous synchronization in a population of coupled neurons (Brown et al., 2003; Ermentrout & Terman, 2010; Hoppensteadt & Izhikevich, 1997; Vreeswijk, Abbott, & Ermentrout, 1994; Wilson, Faramarzi, Moehlis, Tinsley, & Showalter, 2018) and entrainment to an exogenous input (Pyragas, Fedaravičius, Pyragienė, & Tass, 2018; Wilson, Holt, Netoff, & Moehlis, 2015; Zlotnik & Li, 2012). Recent works have investigated strategies for obtaining coupling functions that are accurate beyond first order accuracy (Gengel, Teichmann, Rosenblum, & Pikovsky, 2020; Park & Wilson, 2021; Rosenblum & Pikovsky, 2019; Wilson & Ermentrout, 2019b). Such techniques can be used to understand how the shape of the coupling functions changes as the coupling strength increases, and subsequently how the coupling strength itself influences the network dynamics.

3.2. Ott-Antenson ansatz

Brain regions can contain on the order of billions of neurons. For large populations of neurons, rather than examining individual neurons, it may be more appropriate to consider information about population-level activity by considering a probability density of neurons at a given phase. More specifically, in the limit as $N \to \infty$, the population of oscillators can be analyzed in terms of probability distributions $\rho(\theta, \omega, t)$ rather than discrete oscillators. Note here that the natural frequency of each oscillator is often assumed to be drawn randomly from a distribution. In these cases the probability distribution ρ is considered a function of both θ and ω.

In some very specific circumstances, further simplifications are available allowing for a concise description of the population dynamics. Ott–Antonsen reduction (also known as the Ott–Antonsen ansatz) (Ott & Antonsen, 2008) can be used to accomplish this task. This seminal paper showed that under a specific set of conditions, probability distributions associated with a population of coupled oscillators evolve in time on an invariant manifold. Additionally, if the distribution of nominal frequencies is Lorentzian, the evolution on this manifold can be described by a single ordinary differential equation. Note that other frequency distributions can be used (Martens, Barreto, Strogatz, Ott, So, & Antonsen, 2009; Ott & Antonsen, 2008) resulting in different reduced order differential equations.

Application of the Ott–Antonsen ansatz can be used to greatly simplify the analysis of a large population of neurons. Indeed, Montbrió, Pazó, and Roxin (2015) uses this approach to characterize a biophysically relevant population of tonically firing neurons in terms of firing rate and mean membrane potential. Similar strategies have been used to investigate dynamical behaviors of macroscopic brain rhythms (Dumont & Gutkin, 2019; Schmidt, Avitabile, Montbrió, & Roxin, 2018). For a comprehensive review of other mean field reduction strategies with an emphasis on neural oscillators, the interested reader is referred to Bick, Goodfellow, Laing, and Martens (2020).

4. Model independent control strategies

Given the complexity of the larger brain circuit that gives rise to many pathological conditions, model independent control algorithms that do not require an accurate description of the underlying dynamics can be an attractive choice in many situations. Below, we highlight a few such strategies.

4.1. Adaptive deep brain stimulation

In a clinical setting, DBS as a treatment for Parkinson's disease is generally applied in an open loop manner by injecting pulses of current in a therapeutic range of 130–180 Hz (Kuncel & Grill, 2004; Volkmann, Herzog, Kopper, & Deuschl, 2002). There has been a growing interest in closing the loop, for example, by using an adaptive or demand controlled approach. Such a strategy could reduce the power required by an implantable device and limit unwanted side effects by turning off the device when it is not needed; effective adaptive DBS (aDBS) implementation could eliminate the time consuming process of tuning the parameters of the DBS probe (pulse width, frequency, etc.).

The identification of a suitable control variable is still an open problem for the development of general aDBS strategies. Initial investigations using local field potentials (LFPs) as a control variable have been promising (Little, et al., 2013; Meidahl, Tinkhauser, Herz, Cagnan, Debarros, & Brown, 2017; Priori, Foffani, Rossi, & Marceglia, 2013; Rosa, et al., 2015). When using LFP data as a control variable, DBS input is generally turned on when the amplitude of the pathological beta rhythms exceeds some prescribed threshold, and turned off once the beta rhythm has been sufficiently suppressed. Note, however, that more sophisticated implementations are possible. Other studies have investigated the possibility of timing the application of DBS pulses relative to a patient's tremor rhythm. Such strategies have also yielded promising results in both computational (Azodi-Avval & Gharabaghi, 2015; Duchet, Weerasinghe, Cagnan, Brown, Bick, & Bogacz, 2020) and experimental studies (Cagnan, et al., 2017; Holt, et al., 2019; Rosin, et al., 2011). Further investigation of these techniques will likely yield improvements over standard open-loop DBS methods.

4.2. BayesIan optimization for exploration of large parameter spaces

A fundamental problem in the implementation of DBS and aDBS strategies is the selection of stimulation parameters (e.g., pulse frequency, pulse amplitude, oscillation phase trigger, etc.). In clinical settings, parameter selection is a time consuming process that generally involves making a modification to the parameters, waiting long enough to assess the resulting effect, and making an informed guess about the next parameter to change. General procedures have been developed to perform this task (Kuncel & Grill, 2004; Volkmann, Moro, & Pahwa, 2006), but it still remains somewhat of an art.

Looking towards the development of methods to better implement this tuning process, the authors of Grado, Johnson, and Netoff (2018) consider a Bayesian optimization algorithm (Jones, Schonlau, & Welch, 1998; Shahriari, Swersky, Wang, Adams, & Freitas, 2015) in order to optimize stimulator parameters. The Bayesian optimization algorithm relies on the definition of an objective function in order to gauge the efficacy of a given parameter set. When the goal is the design of an algorithm to alleviate the motor control symptoms of Parkinson's disease, mean beta power is a natural choice. In contrast to frequentist approaches, Bayesian techniques incorporate a prior belief about the objective function, compute a posterior after each new sample, and iterate to gain better and better estimation about the objective function. During the course of the iterations, the Bayesian optimization approach also considers an acquisition function to guide the sampling process in the search for an optimal set of parameters. This approach was implemented successfully in a computational model of Parkinson's disease in Grado et al. (2018). A similar strategy was also effective at controlling seizures in a mouse model of temporal lobe epilepsy (Stieve, Richner, Krook-Magnuson, Netoff, & Krook-Magnuson, 2021). It would be interesting to see how this approach performs in human trials.

4.3. Delayed feedback

The delayed feedback control strategy was originally proposed in the context of chaos control (Pyragas, 1992, 2006). To implement such a control strategy, consider a general dynamical system

$$\frac{dx}{dt} = f(x, p), \tag{14}$$

$$y(t) = g(x(t)), \tag{15}$$

where $x \in \mathbb{R}^n$ is the state, p is a scalar parameter that can be used to apply a control input, and $y(t)$ is a scalar observable which is a function of the state. The delayed feedback approach applies the control signal

$$p(t) = K(y(t-T) - y(t)), \tag{16}$$

where K is a proportional gain and T is a time delay. This general approach was considered in Popovych, Hauptmann, and Tass (2006), Rosenblum and Pikovsky (2004a, 2004b) and Schöll, Hiller, Hövel, and Dahlem (2009) in the context of designing a feedback controller to eliminate pathological oscillations in a large population of neurons. In practice, such delayed feedback techniques work best when the observable gives a good representation of the aggregate population behavior, for instance, when using a mean field as the observable. While this control algorithm is relatively easy to implement, it typically does not come with performance guarantees and effective choices of the time delay and the proportional gain must generally be obtained by sweeping through a set of allowable values.

4.4. Control using machine learning

Machine learning has seen tremendous growth over the last several years. While there have been countless applications to problems such as computer vision and natural language processing, extensions into other areas such as neuroscience and control are still emerging. Two types of machine learning which have shown promise in this area are artificial neural networks and reinforcement learning.

An artificial neural network consists of a system that maps an input vector to an output vector via one or more layers; the system is a "deep" network if at least one of these layers is "hidden", meaning it does not map directly to the output. Each layer is made up of a number of artificial "neurons", each consisting of a matrix that linearly transforms the output from the previous layer followed by a nonlinear activation function. Training a neural network consists of repeated iterations of a forward propagation step followed by a backward propagation step. In the forward propagation step, the current parameters of the neural network are used to make a prediction of the output based on the input, which is compared to the actual output. This provides an accuracy metric based on a loss function which is then fed back into the neural network to compute appropriate derivatives for updating the weights in each layer of the neural network. An appealing aspect of artificial neural networks is that any smooth function can be accurately approximated by a neural network consisting of only a single hidden layer (Cybenko, 1989). In practice, however, it is typically more efficient to use a larger number of hidden layers. The deep neural network architecture allows nonlinearities to compound at each layer, allowing the approximation of highly nonlinear functions with less computational cost than a simpler, single-hidden layer network.

Matchen and Moehlis (2021) designed a deep artificial neural network for predicting the effects of square wave stimuli on a model of a neuron in the subthalamic nucleus. The goal was to utilize as little information as possible when training the neural network and designing the control, so it was assumed that the only measurement we could recover from the neuron was the timing of its spikes. The control signal used, however, was assumed to be known, and was characterized by three variables: signal amplitude u_0, signal width t_{width}, and signal delay time t_{delay}. The delay time is the length of time after the previous spike that the input stimulus is applied. The system's output was the expected value of the neuron's next firing time, t_{final}. This can be viewed as a regression problem, where the input $\vec{x}$ and output $\hat{y}$ of the neural network are written as:

$$\vec{x} = \left[u_0,\ t_{width},\ t_{delay}\ \right]^T, \qquad \hat{y} = E\left[t_{final}|\vec{x}\right]. \tag{17}$$

The neural network was trained via random trials of $\vec{x}$ selected uniformly from appropriate ranges, and was validated by comparing with trials from outside the training set (Matchen & Moehlis, 2021). This deep artificial neural network regression model was then used to find the best signal width for a constant-amplitude stimulus to desynchronize a pair of neurons. Moreover, it was incorporated into a dynamic programming formulation to desychronize a pair neurons with multiple inputs, and a cost function defined in terms of generalized order parameters (Daido, 1996) was used with the regression model to achieve clustered solutions (Matchen & Moehlis, 2021).

Another promising machine learning approach is reinforcement learning, which enables the learning of a control policy by trial and error using feedback from its own actions. The classical formulation of reinforcement learning is in terms of a Markov Decision Process in which $s \in S$ is the state of the system, a is an action drawn from a finite set A of possible actions, $R(s,a)$ is a reward function which depends on the state and action being implemented, and $\pi : S \to A$ is a policy which depends on the state of the system and is trained to give the optimal action. As a generalization, one can optimize based on the value function $V(s,a)$ which equals the reward function plus a constant multiplying the expected future reward. Q-learning is a model-free reinforcement learning algorithm. Recent results (Mitchell & Petzold, 2018; Nagaraj, Lamperski, & Netoff, 2017; Yu, Narayanan, Ching, & Li, 2020) have begun leveraging reinforcement learning to generate control strategies.

Finally, Monga and Moehlis (2020) shows that supervised learning applied to a reward function defined as the absolute value of spike time differences between neurons can be used to develop a control algorithm to desynchronize a population of neural oscillators.

5. Chaotic desynchronization

Wilson and Moehlis (2014c) presents a control strategy called *optimal chaotic desynchronization* for finding an energy-optimal stimulus which exponentially desynchronizes a population of neurons, based only on a neuron's phase response curve (PRC). Here we give a summary of this approach and discuss related control strategies.

5.1. Optimal chaotic desynchronization

As shown in Section 2.2, phase reduction gives the following reduced model for a single neuron:

$$\frac{d\theta}{dt} = \omega + z(\theta)u(t), \tag{18}$$

where θ is the phase of the neuron and is 2π-periodic on $[0, 2\pi)$. By convention, $\theta = 0$ corresponds to the spiking of the neuron. Here, ω gives the neuron's baseline dynamics which are determined from its natural period T as $\omega = 2\pi/T$, $z(\theta)$ is the PRC, and $u(t)$ is the control input.

The approach is based on Lyapunov exponents, which can be used to describe the rate at which nearby trajectories diverge. An expression for the Lyapunov exponent for (18) is obtained by considering two identical neurons subject to a common stimulus:

$$\frac{d\theta_i}{dt} = \omega + z(\theta_i)u(t), \quad i = 1, 2. \tag{19}$$

Here it is assumed that the neurons are nearly in-phase, so that $\theta_1 \approx \theta_2$. Letting $\phi \equiv |\theta_2 - \theta_1|$, one obtains

$$\frac{d\phi}{dt} = z'(\theta)u(t)\phi + \mathcal{O}(\phi^2). \tag{20}$$

Dropping the $\mathcal{O}(\phi^2)$ terms and assuming that solutions are of the form $\phi \sim e^{\Lambda t}$, the finite time Lyapunov exponent (cf., Abouzeid & Ermentrout, 2009) is

$$\Lambda(\tau) = \frac{\log(\phi(\tau))}{\tau} = \frac{\int_0^\tau z'(\theta(t))u(t)dt}{\tau}. \tag{21}$$

Now consider a population of neurons, each described by an equation of the form (18), and choose some time t_1. Suppose that for all stimuli $u(t)$ that advance θ from $\theta(0) = 0$ to $\theta(t_1) = \omega t_1$ (that is, stimuli

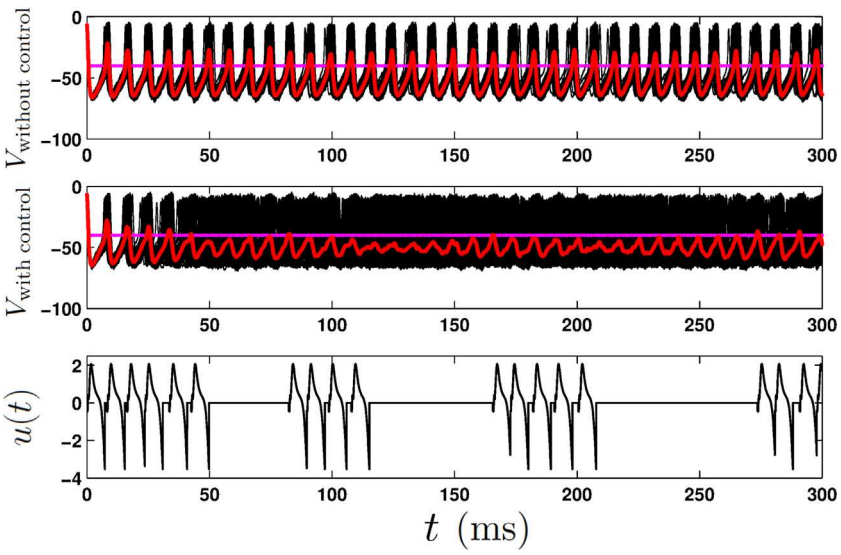

Fig. 2. Optimal chaotic synchronization for a population of $N = 100$ noisy, coupled neurons. The first panel shows the network in the absence of control. The second and third panels show results for the same network with the event-based control applied. In the first and second panels, the black traces are for individual neurons, the red trace shows the mean voltage for the system, and the horizontal dotted line shows the control activation threshold.
Source: Adapted from Figure 6 from Wilson and Moehlis (2014c).

that do not create any net change of phase), one wants to find the stimulus that minimizes the cost functional

$$G[u(t)] = \int_0^{t_1} [u(t)^2 - \beta z'(\theta(t))u(t)]dt. \tag{22}$$

Here, $\int_0^{t_1}[u(t)^2]dt$ corresponds to the power associated with the stimulus, and $\beta > 0$ determines the relative importance of minimizing this versus maximizing the Lyapunov exponent $\Lambda(t_1)$. Now, apply calculus of variations to minimize (Forger & Paydarfar, 2004)

$$C[u(t)] = \int_0^{t_1} \left\{ u(t)^2 - \beta z'(\theta)u(t) + \lambda \left(\frac{d\theta}{dt} - \omega - z(\theta)u(t) \right) \right\} dt, \tag{23}$$

where the Lagrange multiplier λ forces the neural dynamics to obey Eq. (18). The resulting Euler–Lagrange equations can be solved subject to the boundary conditions $\theta(0) = 0$, $\theta(t_1) = \omega t_1$ using a shooting method.

Because pathological neural synchronization in the thalamus is thought to play an important role in Parkinson's disease (Hammond, Bergman, & Brown, 2007), Wilson and Moehlis (2014c) illustrated optimal chaotic desynchronization by considering a three-dimensional model to describe thalamic neural activity augmented from Rubin and Terman (2004) and given in Appendix B. Fig. 2 shows results from event-based control applied to a network of coupled, initially synchronized, noisy thalamic neurons. Here the controller is only triggered when the mean voltage $\bar{V}$ crosses a certain threshold, in which case the pre-computed stimulus which maximizes the Lyapunov exponent is applied; see Wilson and Moehlis (2014c) for more details. We note that Wilson and Moehlis (2014c) shows that heterogeneity in the properties of the neurons and the coupling tends to increase the effectiveness of optimal chaotic desynchronization.

This control strategy can be generalized to obtain charge-balanced (CB) stimuli (Wilson & Moehlis, 2014c). The importance of CB stimuli in DBS has been known for many years: over time, non-charge-balanced (NCB) stimuli can create an accumulation of charge and cause harmful Faradaic reactions that may damage surrounding neural tissue or the DBS electrode (Merrill, Bikson, & Jefferys, 2005). The total charge q imparted to a neuron is the integral of the input current from the control, so $\dot{q}(t) \sim u(t)$. One can derive an optimal CB stimulus by optimizing the same cost function as in the NCB case, $G[u(t)]$, subject to the additional constraints $q(0) = 0$ and $q(t_1) = 0$, which ensures that the stimulus will be charge neutral at t_1. Calculus of variations can be applied to give new Euler–Lagrange equations which can be solved using a double bisection algorithm; see Wilson and Moehlis (2014c) for details. Optimal chaotic desynchronization has also been extended to the more realistic case of extracellular stimulation in Wilson and Moehlis (2014b).

5.2. Phasic burst stimulation

Considering the optimal chaotic desynchronization control strategy, the cost functional from Eq. (22) is imposed to balance the optimal trade-off between the energy consumed by a continuous input and the resulting finite time Lyapunov exponent. In experimental applications that are more relevant to clinical DBS, however, one is generally limited to the application of short duration electrical pulses. Given this constraint, one could conclude that in order to efficiently generate a positive finite time Lyapunov exponent, and hence desynchronization, pulses should be applied where the derivative of the phase response curve is the largest. These issues were investigated in depth in Holt, Wilson, Shinn, Moehlis, and Netoff (2016) in the context of providing a burst of stimulation at a specific phase of oscillation in order to disrupt pathological beta rhythms. In a detailed computational model of a population of conductance-based neurons (Hahn & McIntyre, 2010) that displays a pathological 34 Hz oscillation in a Parkinsonian state, stimulus pulses applied when the derivative of the effective population PRC is positive reduces pathological oscillations as gauged by the power spectrum of the signal obtained from the summation of phases of the aggregate neural population. Conversely, pulses applied when the derivative of the population PRC is negative tend to increase pathological oscillations. Both of these effects grow more pronounced when more pulses are applied in rapid succession.

The phasic burst stimulation framework fits within the larger context of developing demand-controlled, adaptive DBS strategies that can be used to improve both the efficiency and efficacy of clinical DBS by implementing a closed-loop algorithm. The critical challenge relies on identifying a suitable control variable to use for feedback.

6. Cluster control

An approach to achieve partial desynchronization is to split the oscillating neurons into clusters, in which only a subpopulation of the neurons are spike-synchronized. Such a strategy can be especially useful

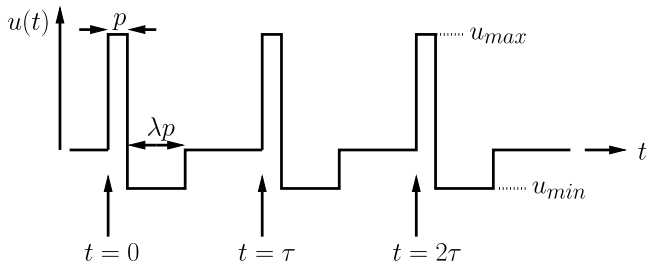

Fig. 3. Periodic sequence of identical pulses.

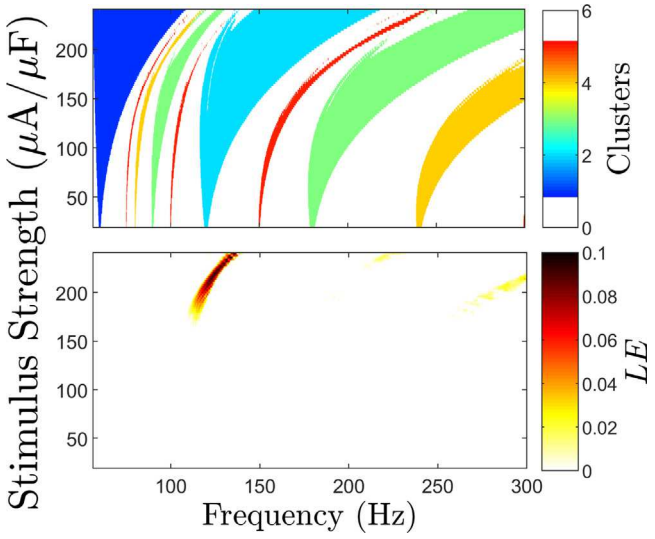

Fig. 4. Results for thalamic neurons driven by periodic pulsatile stimulation, showing regions of parameter space where (top) 1, 2, 3, 4, or 5 clusters occur, as indicated by color, and (bottom) positive Lyapunov exponents occur.
Source: Adapted from Figure 5 from Wilson and Moehlis (2015).

in an underactuated system where full desynchronization is not possible to achieve. In this section, we describe various control algorithms which lead to clustering.

6.1. Open-loop control through periodic pulses

Consider a population of uncoupled neural oscillators stimulated by periodic pulses of the type used for DBS, namely a charge-balanced τ-periodic sequence of pulses as shown in Fig. 3. By adjusting the magnitude and frequency of the stimulation, it is possible to cause the neural population to synchronize, desynchronize, or form clusters. The following describes how one can predict the effect of such periodic, pulsatile stimulation.

Fig. 4 shows numerical results for the effect of periodic pulsatile stimulation for a population of thalamic neurons, according to stimulus strength and frequency. The colored regions in the top panel correspond to clustering, in which neurons within a subpopulation are synchronized, but desynchronized with respect to the other subpopulations; only regions with 1, 2, 3, 4, or 5 clusters are shown, where a 1-cluster state corresponds to full synchronization. White regions represent parameters where clustering is not guaranteed, or where there are 6 or more clusters.

To understand the clustering behavior of a neural population in response to periodic, pulsatile stimuli, Wilson and Moehlis (2015) approximated the dynamics for this system as

$$\dot{\theta}_i = \omega + f(\theta_i)\delta(\mathrm{mod}(t,\tau)), \qquad i = 1,\dots,N, \tag{24}$$

where $\theta \in [0, 2\pi)$ is the phase of the neuron with $\theta = 0$ defined to be the phase at which the neuron fires an action potential, ω is the natural frequency, τ is the period of the stimuli, and $f(\theta)$ is a continuous and periodic function that describes the change in phase due to a single pulse (including the positive current for time p, and the negative current for time λp). Note that for weak inputs, $f(\theta)$ is directly related to the infinitesimal phase response curve defined in Eq. (6). If the pulse was a delta function with unit area, $f(\theta)$ would be equal to the infinitesimal PRC $z(\theta)$; for more general pulses, it can be calculated using a direct method in which a pulse is applied at a known phase, and the change in phase is deduced from the change in timing of the next action potential (Netoff et al., 2012). It is useful to think of the change in phase due to the pulse as occurring instantaneously, even though the pulse will typically have a finite duration; this will be a good approximation for pulses of short duration.

To understand the clustering behavior, consider the map which takes the phase of a neuron immediately after the start of an input pulse, $\theta(0^+) = \theta_0$, to the phase exactly n forcing cycles later (Wilson & Moehlis, 2015):

$$\theta(n\tau^+) = g^{(n)}(\theta_0), \tag{25}$$

where

$$g(s) = s + \omega\tau + f(s + \omega\tau), \tag{26}$$

and $g^{(n)}$ denotes the composition of g with itself n times. Fixed points of $g^{(n)}$, that is, solutions to $\theta^* = g^{(n)}(\theta^*)$, have the property that the phase has the same value after n pulses as where it started. For clustering, of particular interest are fixed points of $g^{(n)}$ which are not fixed points of $g^{(m)}$ for any positive integer m satisfying $m < n$; then there will be n fixed points of $g^{(n)}$ that correspond to points on a period-n orbit of g. If

$$\left| \frac{d}{d\theta}\bigg|_{\theta=\theta^*} (g^{(n)}(\theta)) \right| < 1, \tag{27}$$

then the fixed point θ^* of $g^{(n)}$ is stable, as is the corresponding period-n orbit of g. Neurons which start with initial phases within the basin of attraction of a given fixed point of $g^{(n)}$ will asymptotically approach that fixed point under iterations of $g^{(n)}$. The n different fixed points will each have a basin of attraction, so, for example, a uniform initial distribution of neurons will form n clusters, one for each of these fixed points of $g^{(n)}$, cf. Wilson and Moehlis (2015). Kuelbs, Dunefsky, Monga, and Moehlis (2020) shows how the number of clusters and their stability properties, bifurcations, and basins of attraction can be understood in terms of the iterates of the map g.

6.2. Clustering in noisy populations of neurons

Consider the behavior of a large population of neurons subject to both noise and periodic pulsatile inputs:

$$\dot{\theta}_i = \omega + f(\theta_i)\delta(\mathrm{mod}(t,\tau)) + \epsilon\eta_i(t)z(\theta_i), \qquad i = 1,\dots,N. \tag{28}$$

This equation is identical to (24) except for the addition of an independent and identically distributed Gaussian white noise process $\eta_i(t)$ to each neuron, where $0 < \epsilon \ll 1$. Because of noise, it is more convenient to think about the phase in terms of a probability distribution. For a large population of neurons, one can consider the dynamics of a probability distributions with a series of stochastic maps:

$$\rho(\theta, t + n\tau) = P_{n\tau}\rho(\theta, t), \tag{29}$$

where $P_{n\tau}$ is the linear Frobenius–Perron operator that maps the probability distribution $\rho(\theta, t)$ to $\rho(\theta, t + n\tau)$. The operator $P_{n\tau}$ was analyzed in Wilson and Moehlis (2015) in relation to fixed points of the map $g^{(n)}$. In particular, if $g^{(n)}$ is monotonic, has n stable fixed points, n unstable fixed points, and no center fixed points, then under repeated applications of the mapping from Eq. (29), $\rho(\theta, t)$ will approach a steady state with n distinct clusters, each containing approximately $1/n$ of

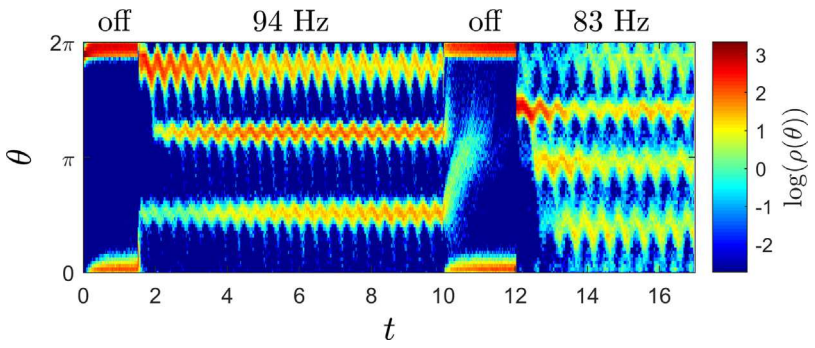

Fig. 5. A population of $N = 1000$ noisy thalamic neurons of the general form (B.1) is considered without electrotonic coupling. A background sinusoidal input common to each neuron serves to synchronize the population and is applied at all times. Pulsatile stimuli applied at 83 and 94 Hz eliminate synchronization in favor of 4- and 3-cluster solutions in steady state. In the figure above, snapshots of the probability distribution are taken immediately before every third and fourth pulse when applying pulsing at 83 and 94 Hz, respectively, and taken at the frequency of the sinusoidal forcing when no pulsatile input is applied.
Source: This figure is adapted from results originally presented in Wilson and Moehlis (2015).

the overall probability density. This result was derived for systems in the absence of coupling or other external inputs, nonetheless, in situations where periodic input is large and any additional inputs are weak, similar results can be observed. For instance, Fig. 5 shows a set of simulations of $N = 1000$ noisy neurons subject to a baseline sinusoidal input and pulsatile inputs applied at various frequencies. In the absence of pulsatile inputs, the sinusoidal input synchronizes the population. When pulsatile inputs are applied, synchronous behavior is disrupted and stable clustering emerges.

6.3. Chaotic desynchronization with pulsatile inputs

Periodic pulsatile inputs can also lead to chaotic desynchronization. The bottom panel of Fig. 4 shows the average Lyapunov exponent of the resulting steady state distributions $\rho^*(\theta)$ from periodic, pulsatile stimulation, calculated as (Wilson, Beverlin, & Netoff, 2011)

$$\mathrm{LE} = \int_0^{2\pi} (\rho^*(\theta) \log[1 + f'(\theta)]) d\theta. \tag{30}$$

For parameter regions with a positive Lyapunov exponent, the pulsatile stimulation will, on average, desynchronize neurons which are close in phase. It was proposed in Wilson et al. (2011) that the parameters which give therapeutic benefit to Parkinson's patients are those with a positive Lyapunov exponent, leading to chaotic desynchronization of the neural population. The bottom panel of Fig. 4 illustrates the Lyapunov exponent associated with a population of thalamic neurons for various stimulus strengths and frequencies. The colormap is scaled to highlight regions where the Lyapunov exponent is positive. Regions of parameter space that yield positive Lyapunov exponents are substantially narrower than the regions that yield clustering.

6.4. Coordinated reset

Coordinated reset uses multiple electrode implants delivering a series of identical impulses separated by a time delay between implants to achieve clustering. This has been studied extensively (Lücken, Yanchuk, Popovych, & Tass, 2013; Lysyansky, Popovych, & Tass, 2011, 2013; Tass, 2003) with preliminary clinical success (Adamchic, et al., 2014). Modeling and clinical results for coordinated reset suggest that relatively strongly clustered groups of neurons do not lead to pathological outcomes in the user and can be effective in treatment for Parkinson's disease.

Coordinated reset relies on a number of electrodes equal to the number of clusters desired, which may not always be physically feasible in practice. It also requires the powering of multiple electrodes simultaneously, which limits its energy efficiency. However, coordinated reset is clearly a promising approach to achieving clustering of neural populations.

6.5. Closed-loop cluster control

Matchen and Moehlis (2018) considers a population of N identical neurons, all receiving the identical input $u(t)$. If the neurons are uncoupled, then from (18) the dynamics of neuron j is given by

$$\frac{d\theta_j}{dt} = \omega + z(\theta_j)u(t), \qquad j = 1, 2, \ldots, N. \tag{31}$$

Now, label the neurons such that, at time $t = 0$, the neuron phases are ordered as $\theta_1(0) < \theta_2(0) < \theta_3(0) < \cdots < \theta_N(0)$. Since the neurons are identical, the response of a neuron is bounded by the neurons of phase initially less than the neuron and those greater than the neuron, so for $t > 0$, $\theta_1(t) < \theta_2(t) < \theta_3(t) < \cdots < \theta_N(t)$, that is, the ordering of the neurons is maintained; here θ_j is allowed to increase beyond 2π, rather than using the modulo 2π value (Li, Dasanayake, & Ruths, 2013; Matchen & Moehlis, 2018).

Suppose that the goal is to split the N neurons into K clusters. The approach from Matchen and Moehlis (2018) is illustrated for $K = 2$ in Fig. 6, where each neuron is represented as a dot with angle corresponding to its phase. In the absence of control input, the dots will move with constant angular velocity ω around the circle. Choose two pairs of neurons – here, one pair is the red and blue neuron, and the other pair is the green and yellow neuron – with initial phases as shown in Fig. 6(a). The control described below is applied so that the neurons in each of these pairs go a state in which they are $2\pi/K = \pi$ out of phase with each other; because the ordering of the neurons is maintained, this will lead to a configuration such as that shown in Fig. 6(b), and in the absence of additional control input they will maintain such separation. Now, suppose that there are additional neurons as shown as black dots in Fig. 6(c). Because the ordering of the neurons is maintained even in the presence of a control input, when the control which makes the pairs of neurons $2\pi/K = \pi$ out phase with each other is applied, the additional neurons' phases get "squeezed" as shown in Fig. 6(d). This results in a two-cluster state. A similar argument applies for obtaining a K-cluster state for $K > 2$.

To find the desired control input, Matchen and Moehlis (2018) defines a state function $r(t)$ which is zero when the appropriate pairs of neurons are $2\pi/K$ out of phase with each other, and tends towards infinity when their phase difference approaches 0 or 2π. It also defines a cost function

$$C(t) = \int_0^t \left[u(\tau)^2 + \alpha r(\tau)\right] \mathrm{d}\tau, \tag{32}$$

where the first term is proportional to the power associated with the control input, and the second term is a penalty associated with the state function, both of which we want to minimize. The value of α determines the relative importance of these two terms. Matchen and Moehlis (2018) rewrites C as a function of r, and chooses the input

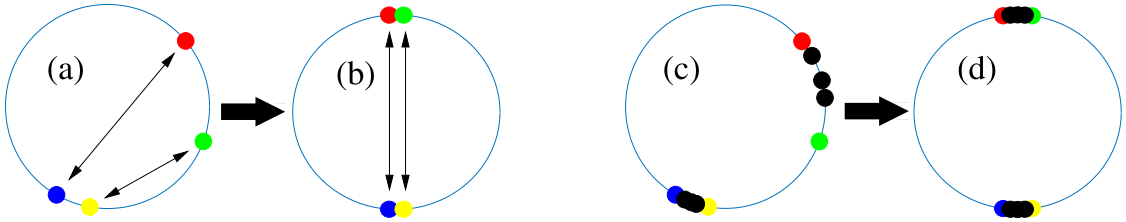

Fig. 6. Illustration of cluster control as proposed in Matchen and Moehlis (2018), here for $K = 2$ clusters. Each neuron is represented as a dot, with angle corresponding to its phase. The (red,blue) and (green,yellow) pairs of neurons are controlled from configuration (a) to configuration (b). When additional neurons are present, shown as black dots, the same control takes the system from configuration (c) to configuration (d), giving two clusters. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$u(t)$ such that the instantaneous magnitude of $\frac{dC}{dr}$ is instantaneously minimized, which can be interpreted as the input that is most efficient in terms of cost relative to change in r. Such control is called an input of maximal instantaneous efficiency (IMIE), and the effectiveness of this algorithm has been shown numerically (Matchen & Moehlis, 2018) and experimentally (Faramarzi & Netoff, 2021).

7. Direct control of phase distributions

Rather than attempting to identify a set of reduced ordinary differential equations, in many control applications that consider a large population of identical neurons it can be useful to consider the dynamical evolution of the probability distribution itself. To this end, consider a population of noise-free, identical, uncoupled neural oscillators all receiving the same control input, with the dynamics of a single neuron given by (18). It is convenient to represent the population dynamics in terms of its probability distribution $\rho(\theta, t)$ which evolves according to the advection equation (Brown et al., 2004; Tass, 2007)

$$\frac{\partial \rho(\theta,t)}{\partial t} = -\frac{\partial}{\partial \theta}\left[\left(\omega + z(\theta)u(t)\right)\rho(\theta,t)\right]. \tag{33}$$

Wilson and Moehlis (2014c) considers the problem of minimizing the value of the peak of such a phase distribution, which has a desynchronizing effect because fewer neurons have nearby phases. It accomplishes this by finding ordinary differential equations for the location of the distribution peak (ρ_M, θ_M), which are then used to formulate an optimal control problem for the energy-optimal minimization of ρ_M.

As an alternative approach, (Kuritz, Zeng, & Allgöwer, 2019; Monga, Froyland, & Moehlis, 2018; Monga & Moehlis, 2019b) formulated an algorithm to control the full distribution. Here we follow the treatment in Monga et al. (2018). The desired final probability distribution $\rho_f(\theta, t)$ was taken to be a traveling wave which evolves according to

$$\frac{\partial \rho_f(\theta,t)}{\partial t} = -\omega \frac{\partial \rho_f(\theta,t)}{\partial \theta}. \tag{34}$$

Note that (34) is of the same form as (33) with $u = 0$. The approach is to select the control $u(t)$ at each time instant so that the L^2 difference between the current density $\rho(\theta, t)$ and the final target density $\rho_f(\theta, t)$ is instantaneously decreased as much as possible subject to the control limits. In particular, Monga et al. (2018) defines the L^2 norm of the probability distribution difference as

$$V(t) = \int_0^{2\pi} \left(\rho(\theta,t) - \rho_f(\theta,t)\right)^2 d\theta. \tag{35}$$

This is positive over the range of all probability distributions, except being zero when $\rho(\theta, t) = \rho_f(\theta, t)$. Its time derivative is given as

$$\dot{V}(t) = I(t)u(t), \tag{36}$$

where

$$I(t) = 2\int_0^{2\pi} \left(\frac{\partial \rho(\theta,t)}{\partial \theta} - \frac{\partial \rho_f(\theta,t)}{\partial \theta}\right) z(\theta)\rho(\theta,t)d\theta. \tag{37}$$

Taking the control input $u(t) = -KI(t)$ with $K > 0$ gives the time derivative of the L^2 norm, $\dot{V}(\rho, \rho_f, t) = -K[I(t)]^2$, as negative definite. Thus, the control law $u(t) = -KI(t)$ will decrease the L^2 norm until the current probability distribution becomes equal to the desired distribution. For both experimental and numerical reasons, it is more practical to have a bounded control input, so Monga et al. (2018) also considered a "clipped" proportional control law and a "bang–bang" control law. These ideas were generalized in Monga and Moehlis (2019b) to formulate a control law when the neural population experiences white noise, and to formulate an optimal control algorithm for phase distributions.

We note that, apart from degenerate cases as explained in Monga and Moehlis (2019b), this algorithm is able to take any initial distribution to a broad range of possible target distributions, including a flat distribution corresponding to desynchronization, a single-peaked distribution corresponding to synchronization, or a multi-peaked distribution corresponding to a cluster state. Fig. 7 shows an example where the initial distribution is flat, and the final distribution corresponds to a 2-cluster state.

8. Phase randomization

The most reliable state measurements for individual neurons are typically the timing of voltage spikes. For a population of neurons, often all that is available is the local field potential, which represents a filtered sum of current traveling across the cell membranes of a population of cells close to a measurement electrode. The availability of such limited measurements suggests the use of *event-based control*, which was developed as an improvement to fixed sample-rate feedback control for digital systems (Åström & Bernhardsson, 2003). In particular, whenever an event is detected, such as a voltage spike or the local field potential crossing some threshold, a pre-computed control stimulus is applied.

Danzl, Hespanha, and Moehlis (2009) and Nabi, Mirzadeh, Gibou, and Moehlis (2013) design controllers that wait for a voltage spike event, then stimulate the neuron with a pre-computed waveform designed to drive the neuron's state close to the phaseless set; for example, in Fig. 1 the phaseless set would be the unstable fixed point "inside" the periodic orbit which is surrounded with closely packed isochrons. When the state is near the phaseless set, background noise will perturb the system onto a random isochron and thus randomize its asymptotic phase and its next spike time as described in Winfree (2001). Here we focus on control algorithms which are either optimal in the sense of minimizing the time or minimizing the amount of energy needed to reach the phaseless set. When such control is applied to a population of neurons, each neuron is randomized differently, and the whole population is desynchronized.

First consider the system corresponding to a single neuron with state $x \in \mathbb{R}^n$ and control $u(t)$, but no noise:

$$\frac{dx}{dt} = F(x) + Bu(t), \tag{38}$$

where $B = [1, 0, \ldots, 0]^T$, and the first component of x corresponds to the membrane voltage of the neuron. From a theoretical standpoint, the objective of the control strategy will be to drive the system's state to a pre-defined target location $\mathcal{T}$ corresponding to the phaseless set. For simplicity, here we assume that $\mathcal{T}$ is an isolated fixed point, and

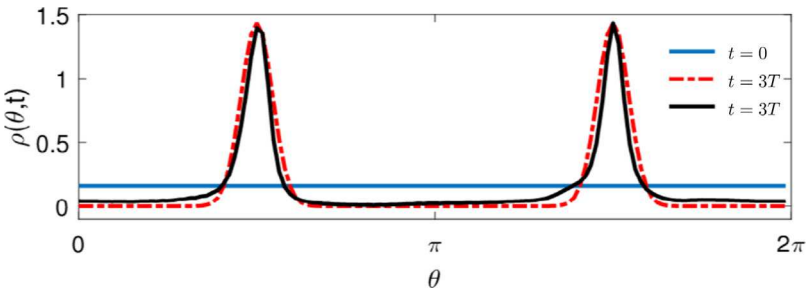

Fig. 7. Phase distribution control, where the initial uniform distribution in shown in blue, and the true distribution and target distribution at time $t = 3T$ are shown in black and red, respectively, where T is the period of the neural oscillators. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: Adapted from Figure 3 from Monga and Moehlis (2019b).

we assume that the system has been translated so $\mathcal{T}$ is at the origin. (In practice, $\mathcal{T}$ is typically taken to be a small neighborhood of the phaseless set.) The constraint on the input due to tissue sensitivity or the available hardware for experimental implementation is taken to be $u(t) \in \mathcal{U} \equiv \{u(t) | u_{min} \leq u(t) \leq u_{max}\}$.

8.1. Time-optimal phase randomization

First consider the *minimum-time-to-reach* control objective, which is to drive the system from its initial state x_0 to the target set $\mathcal{T}$ in minimum time. The optimal control signal is found by first computing an approximation of the minimum-time-to-reach value function (sometimes called the cost-to-go function), which is a viscosity solution of a Hamilton–Jacobi–Bellman (HJB) PDE. The numerical approximation of the value function is then used to generate the optimal state trajectory and the optimal control signal by forward simulation.

This approach begins by defining the terminal time, $t_{end} \in [0, \infty]$, which is the minimum time at which the state reaches the target set when starting from $x(0) = x$ under the control signal $u(t)$:

$$t_{end}(x, u(t)) = \min\{t : x(t) \in \mathcal{T} | x(0) = x\}. \tag{39}$$

The terminal time is taken to be infinite when the trajectory $x(t)$ never reaches the target set, which can occur if the constraint on $u(t)$ does not allow enough control authority to over-ride the system's natural autonomous dynamics. The terminal time is not known at the outset, and is only found through calculating the optimal stimulus and optimal state trajectories.

Next, define the *cost* functional J of an $(x(t), u(t))$ trajectory starting at $x(0) = x$ (over the time interval $t \in [0, t_{end}]$) as

$$J(x, u(t)) = \int_0^{t_{end}} 1 \, dt = t_{end}(x, u(t)). \tag{40}$$

Minimizing J is equivalent to our objective of reaching the target set in minimum time. The minimum-time-to-reach *value* function, $\mathcal{V}(x)$, from state x at time $t = 0$ to the target set is computed as

$$\mathcal{V}(x) = \inf_{u \in \mathcal{U}} J(x, u) = \inf_{u \in \mathcal{U}} t_{end}(x, u(t)). \tag{41}$$

The value function is infinite at points in state space from which the controller cannot drive the state to the target set. From classical optimal control theory (Bardi & Capuzzo-Dolcetta, 1997; Kirk, 1970), in the minimum-time-to-reach framework the value function $\mathcal{V}(x)$ is a viscosity solution of the following Hamilton–Jacobi–Bellman equation:

$$0 = \min_{u \in \mathcal{U}} \{1 + \nabla \mathcal{V}(x) \cdot (F(x) + Bu(t))\} \tag{42}$$

with the boundary condition

$$\mathcal{V}(x) = 0 \quad \forall x \in \mathcal{T}. \tag{43}$$

The solution $\mathcal{V}(x)$, in the viscosity sense, enables the computation of the optimal state-feedback (and ultimately event-based open-loop) policies u^* and the corresponding state trajectories x^* for any given starting state x. Given $\mathcal{V}(x)$, the state-feedback form of the optimal control policy $u^*(x)$ is called the H-minimal control (Athans & Falb, 1966) evaluated at x:

$$\begin{aligned} u^*(x) &= \arg\min\{1 + \nabla \mathcal{V}(x) \cdot (F(x) + Bu)\} \\ &= -\text{sgn}\left(\frac{\partial \mathcal{V}}{\partial x_1}\Big|_x\right). \end{aligned} \tag{44}$$

The optimal state trajectory $x^*(t)$ satisfies the system dynamics driven by the optimal state-feedback control law:

$$\begin{aligned} \dot{x}^*(t) &= F(x^*(t)) + Bu^*(x^*(t)) \\ &= F(x^*(t)) - B \, \text{sgn}\left(\frac{\partial \mathcal{V}}{\partial x_1}\Big|_{x^*(t)}\right). \end{aligned} \tag{45}$$

One computes the open-loop optimal control signal $u^*(x_0, t)$ for all $t \in [0, t_{end}(x_0)]$ by simulating Eq. (45) starting from initial position $x(0) = x_0$ until x reaches the target set $\mathcal{T}$ at time $t = t_{end}(x_0) = \mathcal{V}(x_0)$. The simulation provides the optimal x trajectory, which is used to calculate the optimal control through Eq. (44). Thus, given any initial condition x_0 one has all the necessary tools to calculate a variable-time-length open-loop control signal $u^*(x_0, t)$. Danzl et al. (2009) describes how the HJB equation can be solved numerically for such problems using ToolboxLS (Mitchell, 2007, 2008). Note that the optimal control for minimum-time-to-reach control is "bang–bang", this is, it switches between u_{min} and u_{max}, which follows from Pontryagin's Minimum Principle (Pontryagin, Trirogoff, & Neustadt, 1962).

As an approximation to the local field potential, Danzl et al. (2009) considered the average voltage of the neural population:

$$\bar{V}(t) = \frac{1}{N} \sum_{i=1}^{N} V_i(t), \tag{46}$$

where V_i is the voltage associated with the ith neuron. The event for the population is taken to be when $\bar{V}$ crosses some threshold value; this only occurs when the neurons are all firing at approximately the same time, representing synchronization. When such an event is detected, the control (44) is applied, driving all neurons to a neighborhood of the phaseless set, thereby randomizing their phases, each differently because each feels a different realization of the noise. Fig. 8 shows results from applying such control to desynchronize a population of 100 coupled (2D reduced) Hodgkin–Huxley neurons. Here two cycles of control were needed to sufficiently desynchronize the population.

8.2. Energy-optimal phase randomization

Now, consider the system (38) where the objective is instead to find the optimal control law that would take the system to its phaseless set $\mathcal{T}$ in some prespecified length of time $[0, t_{end}]$, while minimizing the cost function

$$J(x, u(t)) = \int_0^{t_{end}} u^2 dt + \gamma q(x(t_{end})). \tag{47}$$

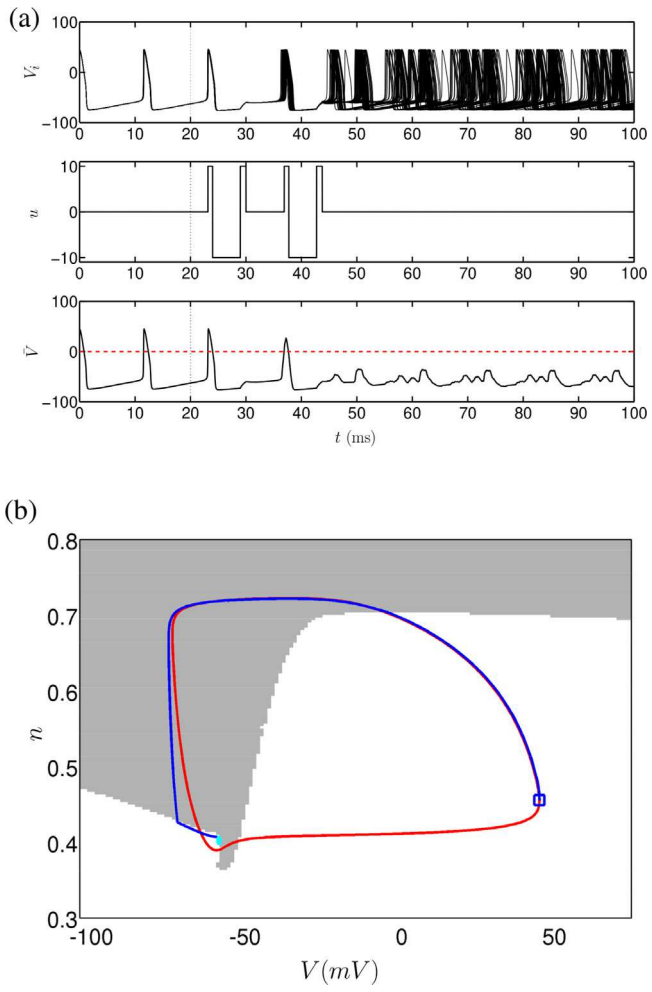

Fig. 8. (a) Time-optimal phase randomization for a population of $N = 100$ reduced Hodgkin–Huxley neurons. The top panel shows voltage traces for the individual neurons, the middle panel shows the optimal input $u(t)$, and the bottom panel shows the population average voltage $\bar{V}$. Control is activated at $t = 20$ ms, as indicated by the vertical dotted lines, and the optimal precomputed stimulus is applied when $\bar{V}$ crosses the threshold $V_{thresh} = 0$, shown as a dashed red line. Here $u_{min} = -10$ and $u_{max} = 10$. (b) Phase space illustration, showing the periodic orbit (red line), optimal state trajectory (blue line), and the target set (cyan circle). Gray regions signify using control u_{min}, and white regions signify using control u_{max}. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: Adapted from Figures 3 and 10 of Danzl et al. (2009).

This cost function has a time-additive portion, $\int_0^{t_{end}} u^2 dt$, that characterizes the total input power being used, and an end-point cost, $q(x(t_{end}))$, that depends on the end state. Here, γ is a penalizing scalar, and again consider bounded inputs $u(t) \in \mathcal{U}$.

Nabi et al. (2013) solves this by employing an HJB approach. One first defines the value function $\mathcal{V}(x, \tau)$ from state x and time $\tau \in [0, t_{end}]$ which gives the minimum cost to go from (x, τ) to a neighborhood of $\mathcal{T}$ subject to the constraints on the optimal control:

$$\mathcal{V}(x,\tau) = \min_{\substack{u(t)\in\mathcal{U} \\ \forall t\in[\tau, t_{end}]}} J = \min_{\substack{u\in\mathcal{U} \\ \forall t\in[\tau, t_{end}]}} \left[\int_\tau^{t_{end}} u^2 dt + \gamma q(x(t_{end}))\right]. \tag{48}$$

Notice that the minimum-energy value function (47) is a function of both time and state, whereas the minimum-time-to-reach value function (41) is only a function of the state. From classical optimal control theory (Bardi & Capuzzo-Dolcetta, 1997; Kirk, 1970), in the energy-optimal framework the value function $\mathcal{V}(x)$ is a viscosity solution of the following HJB equation:

$$0 = \frac{\partial \mathcal{V}}{\partial t}(x(\tau),\tau) + \min_{u\in\mathcal{U}}\left[u(\tau)^2 + \frac{\partial \mathcal{V}}{\partial x}(x(\tau),\tau)\Big(F(x(\tau)) + Bu(\tau)\Big)\right], \tag{49}$$

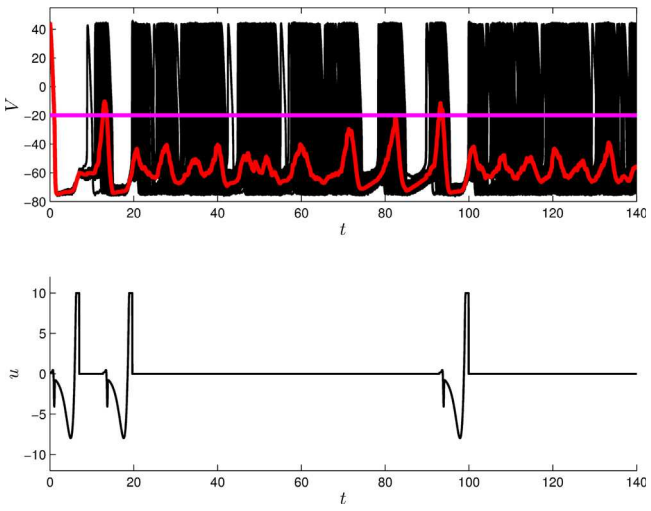

Fig. 9. Energy-optimal phase randomization for a population of $N = 100$ noisy, coupled reduced Hodgkin–Huxley neurons. The top panel shows voltage traces for the individual neurons in black and the average voltage $\bar{V}$ in red. The bottom panel shows the input $u(t)$; here the optimal precomputed stimulus is applied when $\bar{V}$ crosses the threshold shown as purple line, and $u_{min} = -10$ and $u_{max} = 10$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: Adapted from Figure 5 from Nabi et al. (2013).

with the boundary condition

$$\mathcal{V}(x(t_{end}), t_{end}) = \gamma q(x(t_{end})). \tag{50}$$

By defining

$$\mathcal{H}(z, \nabla\mathcal{V}, u) = u^2 + \nabla\mathcal{V}(z(t), t)(F(z(t)) + Bu(t)) \tag{51}$$

as the Hamiltonian for the system, one can rewrite (49) more succinctly as

$$\frac{\partial \mathcal{V}}{\partial t} + \min_{u\in\mathcal{U}} \mathcal{H}(x, \nabla\mathcal{V}, u) = 0, \tag{52}$$

where $\nabla\mathcal{V}$ is the gradient of the value function with respect to x. The optimal control that globally minimizes $\mathcal{H}$ is obtained as

$$u^*(t) = \arg\min_{u\in\mathcal{U}}[u^2 + \nabla\mathcal{V}(x^*(t), t)(F(x^*(t)) + Bu(t))],$$

where $x^*(t)$ is the optimal trajectory. In order to find the optimal control, one sets the derivative of the Hamiltonian (51) with respect to u equal to zero and solves for the extremal u. This is true as long as the magnitude of the control remains smaller than the predetermined bound u_{max}. When the magnitude of the optimal control reaches the bound u_{min} or u_{max}, it saturates in accordance to Pontryagin's minimum principle (Kirk, 1970; Pontryagin et al., 1962). See Nabi et al. (2013) for more detail on this and for a discussion of solving the HJB equation numerically. Fig. 9 shows results from applying event-based energy-optimal control to desynchronize a population of 100 coupled (2D reduced) Hodgkin–Huxley neurons. Notice that two cycles of control are needed to initially sufficiently desynchronize the population, and another cycle is needed when the coupling has caused the neurons to synchronize enough to trigger another event.

9. Open loop control strategies

Deep brain stimulation is generally applied in an open loop manner without feedback. Nonetheless, it is often able to yield robust changes in the population behaviors that result in desirable modifications in the dynamics of a larger brain circuit. In Wilson (2020b) it was suggested that open loop stimulation could be used to strategically destabilize pathologically synchronized solutions and replace them with clustered states, specifically stable rotating block or splay state solutions.

It has long been known that periodic inputs can be used to stabilize unstable periodic orbits (Braiman & Goldhirsch, 1991; Chacon & Bejarano, 1993; Lima & Pettini, 1990; Ramesh & Narayanan, 1999). However, theoretical understanding has generally been limited to low dimensional systems, often precluding the use of nonlinear control algorithms to achieve this objective in practical applications. More recently, Wilson (2019) developed a control framework for modifying the stability of a given periodic orbit using the phase-isostable coordinate paradigm. This strategy is briefly summarized below in Section 9.1, and subsequently considered in the context of neural control in Sections 9.2 and 9.3.

9.1. Modifying the stability of weakly stable and unstable periodic orbits with open-loop control

Consider a general dynamical model of the form (4). As discussed in Appendix A, such a model can be analyzed using a second order accurate phase-isostable reduction of the general form (A.6)

$$\begin{aligned}\frac{d\theta}{dt} &= \omega + z(\theta)u(t) + \sum_{k=1}^{\beta} \psi_k z^k(\theta)u(t),\\ \frac{d\psi_j}{dt} &= \kappa_j\psi_j + i_j(\theta)u(t) + \sum_{k=1}^{\beta} \psi_k i_j^k(\theta)u(t), \qquad j = 1,\dots,\beta.\end{aligned} \tag{53}$$

Above, θ is the phase with z and ω representing the PRC and natural frequency as defined in Eq. (8). Each $\psi_1,\dots,\psi_\beta$ is an isostable coordinate which, as discussed in Appendix A, give a sense of distance from the periodic orbit. The term i_j is the infinitesimal isostable response curve (IRC) of the ψ_j isostable coordinate, which quantifies the effect of a small amplitude input on the isostable coordinate. Additional terms $z^k(\theta)$ and $i_j^k(\theta)$ result when considering the gradient of the phase and isostable coordinates, respectively, using a Taylor expansion in a basis of isostable coordinates centered at the periodic orbit. Compared to the standard phase reduction (8), these extra terms yield a reduced order model that is more accurate when larger magnitude inputs are considered. Here, for simplicity of presentation the input $U(x,t)$ is assumed to be a rank-1 perturbation, i.e., $U(x,t) = \delta u(t)$, where $\delta \in \mathbb{R}^n$ is a constant. With this in mind, Eq. (53) can be derived from Eq. (A.6) by taking the expansions from Eq. (A.8) to first order accuracy in each isostable coordinate and simplifying appropriately (e.g., $z^k(\theta) = Z^k(\theta)\cdot\delta$). As discussed in Wilson (2019), the reduction (53) is valid both for unstable and stable periodic orbits; for an unstable orbit, the phase can no longer be determined according to isochrons and instead must be defined using either an analogous finite-time definition or implicitly using level sets of Koopman eigenfunctions. Taking $u(t)$ to be a T_1-periodic input, one can consider a rotating reference frame by letting $\phi = \theta - \omega_1 t$, where $\omega_1 = 2\pi/T_1$. The dynamics in this rotating reference frame become

$$\begin{aligned}\dot\phi &= \Delta\omega + \left[z(\phi+\omega_1 t) + \sum_{k=1}^{\beta}\left(\psi_k z^k(\phi+\omega_1 t)\right)\right]u(t),\\ \dot\psi_j &= \kappa_j\psi_j + \left[i_j(\phi+\omega_1 t) + \sum_{k=1}^{\beta}\left(\psi_k i_j^k(\phi+\omega_1 t)\right)\right]u(t), \qquad j = 1,\dots,\beta,\end{aligned} \tag{54}$$

where $\Delta\omega = \omega - \omega_1$. Provided $\Delta\omega$ and each ψ_j are $O(\epsilon)$ terms, formal averaging techniques (Guckenheimer & Holmes, 1983; Sanders, Verhulst, & Murdock, 2007) can be used to approximate (54) as

$$\begin{aligned}\dot\Phi &= \Delta\omega + \sigma(\Phi) + \sum_{k=1}^{\beta}\left(\Psi_k Y_k(\Phi)\right),\\ \dot y &= (A + E(\Phi))y + p(\Phi),\end{aligned} \tag{55}$$

where

$$y = \begin{bmatrix}\Psi_1\\ \Psi_2\\ \vdots\\ \Psi_\beta\end{bmatrix}, \quad A = \begin{bmatrix}\kappa_1 & 0 & \dots & 0\\ 0 & \kappa_2 & & 0\\ \vdots & & \ddots & \vdots\\ 0 & 0 & \dots & \kappa_\beta\end{bmatrix},$$

$$E(\Phi) = \begin{bmatrix}\alpha_{1,1}(\Phi) & \alpha_{1,2}(\Phi) & \dots & \alpha_{1,\beta}(\Phi)\\ \alpha_{2,1}(\Phi) & \alpha_{2,2}(\Phi) & \dots & \alpha_{2,\beta}(\Phi)\\ \vdots & \vdots & \ddots & \vdots\\ \alpha_{\beta,1}(\Phi) & \alpha_{\beta,1}(\Phi) & \dots & \alpha_{\beta,\beta}(\Phi)\end{bmatrix}, \quad p(\Phi) = \begin{bmatrix}\mu_1(\Phi)\\ \mu_2(\Phi)\\ \vdots\\ \mu_\beta(\Phi)\end{bmatrix},$$

with $\sigma(\Phi) = \frac{1}{T_1}\int_0^{T_1} z(\Phi+\omega_1 t)u(t)dt$, $Y_k(\Phi) = \frac{1}{T_1}\int_0^{T_1} z^k(\Phi+\omega_1 t)u(t)dt$, $\mu_k(\Phi) = \frac{1}{T_1}\int_0^{T_1} i_k(\Phi+\omega_1 t)u(t)dt$, $\alpha_{j,k}(\Phi) = \frac{1}{T_1}\int_0^{T_1} i_j^k(\Phi+\omega_1 t)u(t)dt$. Here, Φ and Ψ_i provide close approximations of ϕ and ψ_i, respectively, when considering the averaged equations. As detailed in Sanders et al. (2007), fixed points of (55) correspond to periodic orbits of (54) with the same stability.

The reduced order, averaged Eqs. (55) ultimately provide a strategy for modifying Floquet exponents associated with a given periodic orbit, and hence, its stability. Ultimately, the Floquet exponents are determined by the eigenvalues of $A + E(\Phi)$. Noting that $E(\Phi)$ contains $O(\epsilon)$ terms, shifts to a simple (i.e., unique) eigenvalue can be determined using perturbation methods (Demmel, 1997):

$$\begin{aligned}\text{eig}(A+E) &= \kappa_i + w_i^T E v_i/(w_i^T v_i) + \mathcal{O}(\epsilon^2),\\ &= \kappa_i + e_i^T E e_i + \mathcal{O}(\epsilon^2),\\ &= \kappa_i + \alpha_{i,i}(\Phi) + \mathcal{O}(\epsilon^2), \text{ for } i = 1,\dots,\beta,\end{aligned} \tag{56}$$

where $w_i^T = e_i^T$ and $v_i = e_i$ (elements of the standard basis) are left and right eigenvectors of A corresponding to eigenvalue κ_i. As such, the terms $\alpha_{i,i}$ from the matrix $E(\Phi)$ determine the shift in a unique Floquet exponent of the underlying periodic orbit in response to periodic input. If the control objective is to stabilize a nominally unstable periodic orbit with a single, real Floquet exponent $\kappa_i > 0$, this can be posed as an optimal control problem where the goal is to minimize the cost functional

$$C[u(t)] = \int_0^{T_1} u(t)^2 dt \tag{57}$$

subject to constraints

$$\begin{aligned}\sigma(\Phi_0) &= -\Delta\omega,\\ \mu_i(\Phi_0) &= 0 \text{ for } i = 1,\dots,\beta,\\ d\sigma/d\Phi|_{\Phi_0} &= \rho < 0,\\ \kappa_i + \alpha_{i,i}(\Phi_0) &= \eta_i < 0, \text{ for } i = 1,\dots,\beta.\end{aligned} \tag{58}$$

Here the first two conditions guarantee that Eq. (55) has a fixed point, and the second two guarantee stability of the associated periodic orbit under the application of the periodic input. This problem can be solved using a calculus of variations formulation as discussed in Wilson (2019).

9.2. Stabilization of unstable rotating block solutions in neural populations using open-loop control

Here, we consider a phase model with N identical coupled neurons

$$\dot\theta_i = \omega + \frac{1}{N}z(\theta_i)\sum_{j=1}^{N} f(\theta_i,\theta_j) + z(\theta_i)u(t), \tag{59}$$

for $i = 1,\dots,N$, where θ_i represents the phase of the ith neuron, $f(\theta_i,\theta_j)$ represents the influence of coupling from neuron j on neuron i, $u(t)$ is an external input applied identically to all neurons, and $z(\theta)$ is the phase response curve. For a population of identical, weakly coupled neurons, many configurations of the individual neural oscillators are guaranteed to exist (Ashwin & Swift, 1992; Brown et al., 2003). For example, synchronous solutions for which $\theta_i = \theta_j$ for all i and j are one common configuration. Additionally, splay states, with all phases of neurons spread equally, and rotating block states having N/G blocks, each with G oscillators, are also possible configurations (see Fig. 10). While these solutions usually exist as solutions of (59), not all of these solutions will be stable. With this in mind, in Wilson (2020b) the open-loop stability

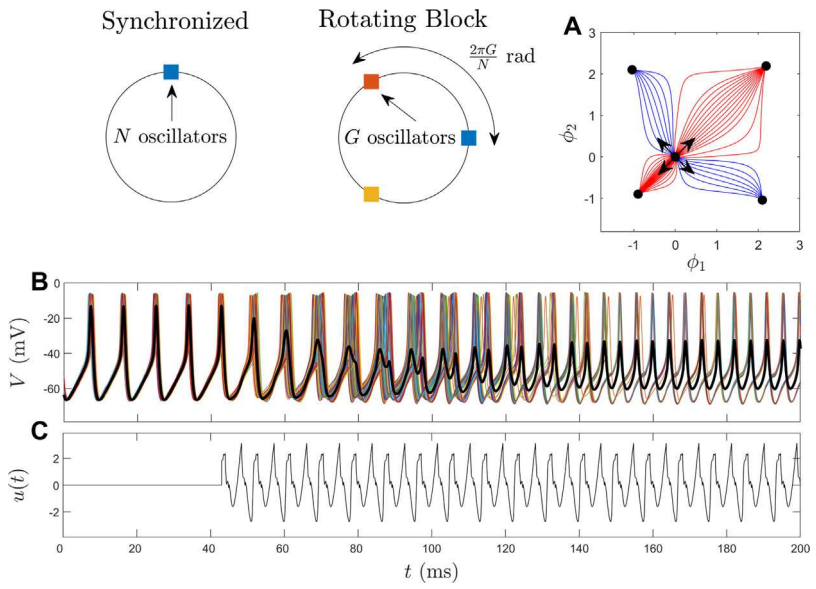

Fig. 10. Top-left panels illustrate the important features of synchronized and rotating block solutions for population of N identical oscillators. When considering the aggregate oscillation of two limit cycle oscillators, an energy optimal input can be designed in order to simultaneously destabilize a synchronous solution in favor of a splay solution. Here ϕ_1 and ϕ_2 are the phases of each oscillator in a rotating reference frame. The synchronous and splay states are indeed destabilized and stabilized, respectively. As shown in Panel A, the solution ends up at a new stable fixed point depending on the basin of attraction associated with the initial condition. Blue (resp., red) lines highlight trajectories that end at a stable splay (resp., synchronous) state. A similar control objective of destabilizing synchronous solutions in favor of a two block rotating solution is illustrated in panels B and C for a population of conductance-based neurons. In panel B, colored lines represent voltage traces for individual neurons and the black line gives the average value. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: This figure is adapted from results originally presented in Wilson (2020b).

modification framework described in Section 9.1 was investigated in the context of destabilizing stable, synchronous solutions in favor of less synchronized rotating block solutions or splay state solutions.

For an appropriate choice of coupling, Eq. (59) has a stable synchronized solution. Considering the dynamics of the entire collection of periodic orbits, these aggregate oscillations can be analyzed with the phase-isostable reduction strategy yielding equations of the form

$$\dot{\Theta} = \Omega + Z(\Theta)\cdot P(t) + \sum_{k=1}^{\beta}\left[Z^k(\Theta)\psi_k\right]\cdot P(t),$$
$$\dot{\psi}_j = \kappa_j\psi_j + I_j(\Theta)\cdot P(t) + \sum_{k=1}^{\beta}\left[I_j^k(\Theta)\psi_k\right]\cdot P(t), \qquad j=1,\dots,\beta, \tag{60}$$

where Θ is the population phase, Ω is the population frequency, each κ_j is a Floquet exponent of the population orbit, Z, Z^k, I_j, and I_j^k each in $\mathbb{R}^N$ are terms of the associated second order accurate phase-isostable reduction, and $P(t) = u(t)\begin{bmatrix}z(\theta_1) & \dots & z(\theta_N)\end{bmatrix}^T$ is the effective perturbation. This general framework can also be used to characterize oscillatory solutions of rotating block and splay state solutions, keeping in mind that the specific terms of the reduced order equations will be different when considering different periodic orbits.

Investigating aggregate oscillations in this manner yields general heuristics for designing effective stimuli to stabilize or destabilize a given periodic orbit. Recalling Eq. (56) and defining phase appropriately so that $\Phi = 0$, to leading order the term $\alpha_{j,J} = \int_0^{T_1} i_j^j(\omega_1 t)u(t)dt$ is solely responsible for changing the Floquet exponents. In Eq. (60), this term of interest takes the form $\int_0^{T_1} I_j^j(\Omega t)P(t)dt$. As shown in Wilson (2020b), for synchronous solutions, each $I_j^j(\Omega t)P(t) = z'(\Omega t)u(t)$ where $'$ denotes the derivative with respect to θ. Note the similarity between this result and the definition of the finite time Lyapunov exponent from Eq. (21); even when coupling is explicitly considered, applying inputs in proportion to the derivative of the phase response curve of the individual neurons is an effective strategy for destabilizing synchronous solutions. In addition to investigating relationships between the single neuron phase response curves and the terms of the phase-isostable reduction for population oscillations, Wilson (2020b) also considers the design of stimuli to simultaneously destabilize a synchronous solution in favor of stabilizing a rotating block solution. These results are summarized in Fig. 10.

9.3. Stabilization of unstable fixed points

Section 9.2 highlights design strategies for open-loop stimuli that can modify the stability of a given periodic orbit. Similar strategies can also be employed in order to modify the stability of fixed point solutions, as considered in Wilson (2020d). In this example, consider a large population of N identical periodically firing neurons that obeys

$$\dot{\theta}_j = \omega + z(\theta)\left(u(t) + \sqrt{2D}\eta_j + \frac{1}{N}\sum_{i=1}^{N}\sigma_c(V(\theta_i) - V(\theta_j))\right), \tag{61}$$

for $j = 1,\dots,N$. Here $\sqrt{2D}\eta_j$ is an independent and identically distributed zero-mean white noise process with intensity D, and the neurons each receive an identical input $u(t)$. The neurons are coupled through electrotonic (voltage difference) coupling with relative strength σ_c. Here, $V(\theta)$ denotes the transmembrane voltage at phase θ.

In the context of controlling the phase distribution of the coupled neural oscillators, also considered in Section 7 of this review article, the population dynamics can be understood in terms of a probability distribution $\rho(\theta,t)$ that evolves according to a nonlinear Fokker–Planck equation (Gardiner, 2004). Using an Ito interpretation for the noise yields the partial differential equation

$$\begin{aligned}\frac{\partial\rho(\theta,t)}{\partial t} &= -\frac{\partial}{\partial\theta}[(\omega + z(\theta)(u(t) + \sigma_c(\overline{V} - V(\theta))))\rho(\theta,t)] \\ &\quad + \frac{\partial^2}{\partial\theta^2}[Dz^2(\theta)\rho(\theta,t)], \\ &= \underbrace{-\omega\rho_\theta + \frac{\partial^2}{\partial\theta^2}[Dz^2(\theta)\rho] - \frac{\partial}{\partial\theta}[z(\theta)\sigma_c(\overline{V} - V(\theta))\rho]}_{\text{unperturbed dynamics}} \\ &\quad \times \underbrace{-[z(\theta)\rho_\theta + z'(\theta)\rho]u(t)}_{\text{external forcing}},\end{aligned} \tag{62}$$

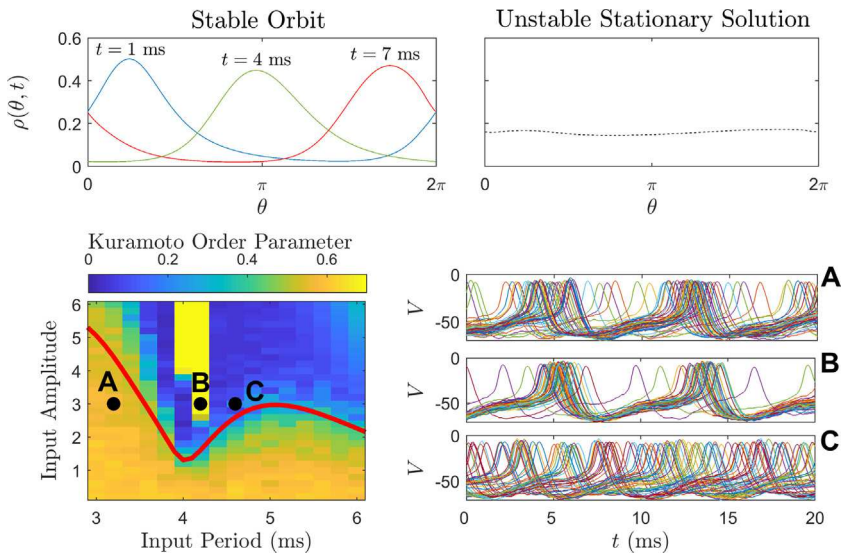

Fig. 11. For a large population of electrotonically coupled thalamic neurons, the top-left panel highlights snapshots of a stable periodic orbit ($T = 8.2$ ms) of Eq. (62) that governs the dynamics of the probability distribution. The top-right panels shows a less synchronized, unstable stationary solution of the same equation. In response to open-loop sinusoidal forcing $u(t) = A\sin(2\pi t/T_1)$ where T_1 is the input period, the bottom-left panel shows the Kuramoto order parameter Kuramoto (1984) in steady state which gives a sense of the phase cohesiveness within the population. An order parameter of 1 indicates a completely synchronized solution; an order parameter closer to zero generally correlates with a less synchronized state. Panels A, B, and C, show traces of representative neurons in steady state when using the input parameters indicated in the bottom-left panel. The red line denotes a theoretical prediction of the boundary over which the unstable stationary solution is stabilized (obtained through analysis (62)) and agrees well with simulations. Note that this prediction is valid in the limit of weak forcing. Additional bifurcations can occur for larger magnitude forcing, for instance, resulting in highly synchronized solutions when using an input period of approximately 4.2 ms.
Source: This figure is adapted from results originally presented in Wilson (2020d).

where $\overline{V} = \int_0^{2\pi} V(\theta)\rho(\theta)d\theta$ is the population averaged voltage, ρ_θ denotes the first partial derivative with respect to θ, and $'$ denotes the derivative with respect to θ. Comparing to Eq. (33), the added diffusive term accounts for the influence of noise. Eq. (62) highlights the difference between the terms that contribute to the nominal dynamics and terms that result due to external input. Considering a population of conductance-based thalamic neurons from Rubin and Terman (2004) to characterize the neurons from (62) (with equations given in Appendix B), in the absence of input, Eq. (62) has a stable periodic orbit highlighted in Fig. 11 for which the phases are strongly correlated. Eq. (62) also has an unstable stationary solution where the phases are spread more diffusely.

Wilson (2020d) considers strategies for stabilization of this unstable stationary solution to engender desynchronization in the network using open-loop feedback. In contrast to the work described in Section 9.2, the goal is to stabilize a stationary solution rather than a periodic solution. A recent study (Duchet, Weerasinghe, Bick, & Bogacz, 2021) found that phenomenological models fitted to patient tremor data often yielded systems with weakly stable fixed points containing complex-conjugate eigenvalues. Theoretical analysis from Wilson (2020d) reveals that to leading order, purely sinusoidal input applied at an appropriate frequency is energy-optimal for achieving this task when fixed points are weakly unstable, that is, when the unstable solutions have a single unstable eigenvalue or a single pair of complex-conjugate unstable eigenvalues with real components of order ϵ, where $0 < \epsilon \ll 1$. Theoretical analysis performed on the Fokker–Planck distribution from (62) is numerically validated using a model of $N = 1000$ conductance-based neurons. These results are summarized in Fig. 11. When applying a purely sinusoidal input, the threshold over which the nominally unstable, less synchronized fixed point is stabilized agrees well with theoretical predictions. In this application, the steady state distributions did not depend on initial conditions, e.g., the initial level of synchronization among neurons. Note that these predictions are valid in the weak forcing regime, and other solutions can result when the amplitude of the forcing becomes larger.

We note that Wilson and Moehlis (2016b) also considers (61) and a partial differential equation very similar to (62). Here the isostable reduction of this partial differential equation about the uniform distribution leads to a bang–bang control strategy to desynchronize a neural population.

10. Control of entrained oscillations

While the majority of the applications considered in this review article are on control of synchronization in populations of tonically firing neurons, there are a number of other control objectives that can be formulated in the context of neural control. In this section, we highlight problems associated with control of entrained oscillations, that is, limit cycle oscillators that are nominally entrained by an external periodic input. In terms of practical applications, we focus here on the control of circadian rhythms that are governed by the suprachaismatic nucleus (SCN) (Reppert & Weaver, 2002), the master pacemaker in the mammalian brain comprised of around 20,000 coupled neurons. The coupled oscillations of these neurons are responsible for maintaining a near 24-hour circadian cycle which can be entrained, for instance, by a 24-hour light-dark cycle (Golombek & Rosenstein, 2010; Wright, Hughes, Kronauer, Dijk, & Czeisler, 2001). When considering control of entrained oscillations, elimination of jet-lag is a natural choice. Mathematical models are often used in conjunction with optimal control techniques to identify optimal schedules of light exposure and avoidance to rapidly acclimate to a new time zone and recover from jet-lag (Bagheri, Stelling, & III, 2008; Dean, Forger, & Klerman, 2009; Serkh & Forger, 2014; University of Michigan, 2018). Alternatively, phase-based reduction strategies have also been developed, which first characterize the phase response to bright pulses of light and subsequently identify strategies that can be used to reentrain rapidly to a new time zone (Chesson, et al., 1999; Waterhouse, Reilly, Atkinson, & Edwards, 2007).

Two fundamental difficulties arise in the use of optimal control techniques in the context of jet-lag amelioration: First, full scale modeling of the coupled oscillations of SCN neurons result in high-order nonlinear systems for which it is computationally difficult to implement nonlinear control techniques. Second, collective rhythms are notoriously difficult to analyze in a reduced order setting, and standard phase

reduction methods often break down when considering control inputs with physiologically relevant magnitudes. In the following sections, we highlight recent efforts to obtain reduced order models that capture the coupled oscillations of SCN neurons for which nonlinear optimal control problems can be formulated and solved.

10.1. Phase-isostable reduction and control of entrained oscillations

Wilson (2020a) considers entrained oscillations in general models of the form

$$\frac{dx}{dt} = F(x) + \delta(u_{\text{nom}}(t_s) + u(t)),$$
$$\frac{dt_s}{dt} = 1, \tag{63}$$

where $x \in \mathbb{R}^n$, $F(x)$ gives the nominal dynamics that yield a stable T-periodic limit cycle $x^\gamma(t)$ in the absence of additional input, $\delta \in \mathbb{R}^n$, $t_s \in \mathbb{S} = [0, T_1)$ is a T_1-periodic time-like variable that governs the timing of the entraining stimulus $u_{\text{nom}}(t_s)$, and $u(t)$ is a control input. Standard phase reduction techniques (as described in Section 2.2) yield a phase model of the form

$$\frac{d\theta}{dt} = \omega + z(\theta)(u_{\text{nom}}(t_s) + u(t)), \tag{64}$$

where $\omega = 2\pi/T$ and $z(\theta)$ is defined appropriately. A common issue when studying entrained oscillations is that the nominal input is large enough to drive the state of the full system far from the unperturbed periodic orbit, thereby degrading the accuracy of the reduction from Eq. (64) and rendering it useless (cf., Diekman & Bose, 2016). To circumvent this issue, one can instead consider the entrained solution $x^\gamma_{\text{ent}}(t)$ that results under the application of the entraining stimulus, yielding reduced order phase and isostable equations of the form

$$\frac{d\theta}{dt} = \omega_1,$$
$$\frac{d\psi_j}{dt} = \kappa_j \psi_j + i_j(\theta)u(t), \tag{65}$$

for $j = 1, \dots, \beta$, where $\omega_1 = 2\pi/T_1$ and $i_j(\theta)$ is an effective isostable response curve. Note that when the reduction is taken relative to the entrained periodic orbit, the state x has no influence on the infinite time behavior. As such, per the definition of isochrons discussed in Section 2.1, the asymptotic phase can be written as $\theta = 2\pi t_s/T_1$ and no longer depends on $u(t)$. Hence, the effective phase response curve is equal to zero. From this perspective, in the context of the identification of jet-lag amelioration strategies, the relative degree of circadian misalignment following a sudden shift in environmental time is directly related to the magnitude of the isostable coordinates, with $\psi_1 = \dots = \psi_\beta = 0$ corresponding to the fully entrained solution. Therefore, following a sudden shift in the environmental time (for instance, caused by rapid travel across multiple time zones) an optimal jet-lag mitigation control strategy can be formulated as a minimum-time-to-reach problem and subsequently solved using a Hamilton–Jacobi–Bellman approach (Kirk, 1998). Specifically, the goal is to find an allowable control $u^*(t)$ that drives the system (65) with initial isostable coordinates $\psi_1(0), \dots, \psi_\beta(0)$ to some target set $\mathcal{T}$ for which the isostable coordinates are small in the minimum possible time. This can be accomplished by defining the cost functional

$$J(\Psi, u(t)) = \int_0^{t_{\text{targ}}} 1 dt = t_{\text{targ}}(\Psi, u(t)), \tag{66}$$

where $\Psi = \begin{bmatrix}\psi_1 & \dots & \psi_\beta\end{bmatrix}^T$ and t_{targ} is the time required to reach $\mathcal{T}$ starting from Ψ under the application of $u(t)$. The input $u^*(t)$ can be found using the strategy described in Section 8.1, whereby a value function $\mathcal{V}(\Psi)$ which represents the minimum possible value of t_{targ} for a given x is obtained and $u^*(t)$ is determined with knowledge of the value function. In Wilson (2020a), this control strategy is applied to a model of $N = 3000$ coupled SCN oscillators with model equations given by (B.3) and (B.4) from Appendix B. Here, the input $u(t)$ corresponds to the intensity of external light and is used to aid recovery from a shift in environmental time, for instance, caused by rapid travel across multiple time zones. Representative results are shown in Fig. 12.

10.2. Adaptive-phase-isostable reduction and control of entrained oscillations

In order to accommodate large magnitude inputs associated with the nominal 24-hour light-dark cycle, the control strategy from Section 10.1 applied the phase-isostable reduction approach to consider the entrained orbit itself. Wilson (2021) employs an alternative strategy, using the adaptive phase-isostable reduction strategy to first obtain an accurate reduced order model and subsequently treating the light-dark cycle as an input. Once again letting $u_{\text{nom}}(t_s)$ be a nominal light-dark cycle and $u(t)$ be a control input, the objective is to identify a control input $u_{\text{tot}}(t) = u_{\text{nom}}(t_s + \Delta t_s) + u(t)$ that minimizes $C = \int_{t_0}^{t_0+T_f} u^2(t)dt$ subject to the end point constraint $x(t_0+T_f) = x^\gamma_{\text{ent}}(t_0+\Delta t_s)$ and the input constraint $u_{\min} \le u_{\text{nom}}(t_s+\Delta t_s)+u(t) \le u_{\max}$. In other words, following a shift in the external time Δt_s, the goal is to find a minimal energy input to drive the state to the new fully entrained solution after T_f time units.

Using a model of 3000 coupled SCN oscillators described by Eqs. (B.3) and (B.4), this optimal control problem is considered by first transforming to an adaptive phase-isostable model of the form (A.12). Ultimately, this strategy results in a two-dimensional representation of the 3000 oscillator model for which the optimal control problem can be formulated and solved using a calculus of variations approach to minimize the cost functional. Using this approach, recovery times of less than 24 h can be achieved in many cases.

As illustrated in Wilson (2021), the adaptive phase-isostable reduction yields solutions that work well when applied to the full order model. Note that this optimal control problem can also be formulated and solved using a first order accurate phase-isostable reduced model of the form (A.4). However, when using the first order accurate reduction, the resulting optimal inputs do not work when applied to the full order model. Additionally, an interesting result was observed when considering large magnitude time advances and delays; the optimal inputs initially desynchronize the oscillators and then resynchronize them with the target phase of oscillation. These results are highlighted in Fig. 13 for a representative time shift of $\Delta t_s = +10$ hours. Note that the control objective did not explicitly consider relative degree synchronization of the oscillators within the population. Nevertheless, this phenomenon was observed for both large magnitude time advances and time delays. It would be of interest to investigate if this is a common feature when considering aggregate oscillations of a large population of oscillators, i.e, if an initial desynchronizing impulse can allow for more rapid phase shifts in general models.

10.3. Using the Ott–Antenson ansatz to investigate recovery from circadian misalignment

The Ott–Antenson ansatz, as discussed in Section 3.2 can also accommodate sinusoidal external forcing in the reduction from an infinite population of oscillators to a single complex-valued ordinary differential equation. In contrast to phase–amplitude-based methods, which require relatively fast decay of the truncated amplitude coordinates, the Ott–Antenson approach yields an exact reduction for the behavior on an invariant manifold. As such, it can provide highly accurate results even when very high amplitude inputs are used making it particularly well suited to study bifurcations that emerge using the low-order system. This approach is applied in Lu, Klein-Cardeña, Lee, Antonsen, Girvan, and Ott (2016) to investigate the sinusoidally forced Kuramoto model using model parameters chosen to represent the behavior of a large population of oscillators from the suprachiasmatic nucleus. In the context of jet-lag recovery, the authors attribute the commonly observed east–west asymmetry in jet-lag recovery times to the presence of a saddle-type fixed point that has a larger influence on the dynamics when the time shifts are larger.

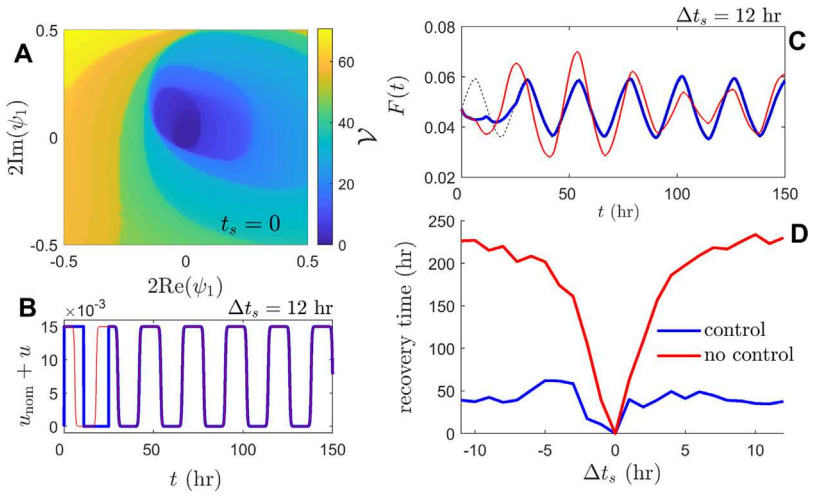

Fig. 12. A population of 3000 coupled SCN oscillators is reduced to a phase-isostable representation of the form (65) with two complex-conjugate isostable coordinates; the goal is to drive these isostable coordinates near zero as quickly as possible following a shift in the external time by Δt_s hours, thereby reducing the time required to reentrain to the new time zone. A Hamilton–Jacobi–Bellman approach is used to solve the minimal-time-to-reach problem by minimizing the cost functional given by Eq. (66). The control input $u(t)$ represents the intensity of an external light source. Panel A shows the value function for initial conditions with $t_s = 0$. Panel B shows the resulting optimal control input in blue for recovery from a $\Delta t_s = 12$ hour time shift. This control formulation results in a bang–bang type controller. The red trace shows the nominal uncontrolled light-dark cycle. Panel C shows the associated recovery of the mean field coupling parameter. The dashed line associated with the fully entrained solution is shown for reference. Note that in the uncontrolled case, the state eventually reaches the fully entrained orbit because of the nominal 24-hour light-dark cycle. Panel D shows the recovery time for various shifts in external time using the optimal control (blue lines) and for the uncontrolled recovery (red lines). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: This figure is adapted from results originally presented in Wilson (2020a).

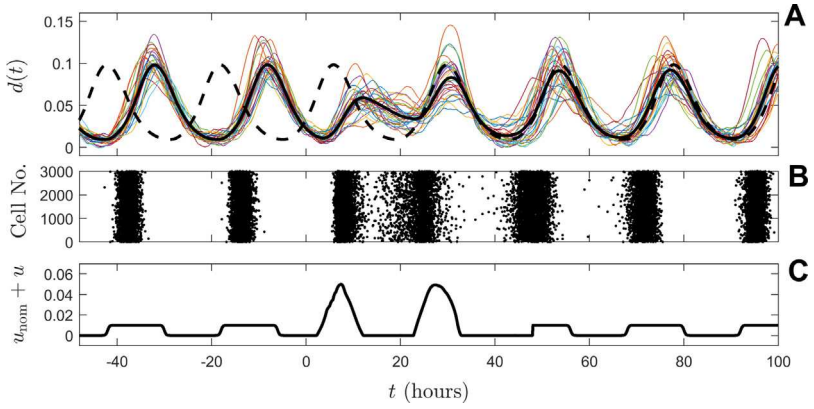

Fig. 13. An optimal control problem is formulated using the adaptive phase-isostable-reduced model of the form (A.12) and solved for a $\Delta t_s = +10$ hour shift in the external time. Colored lines in panel A show individual traces of $d(t)$, the variable that contributes to the mean field coupling. The solid black line shows the average value in the population, with the dashed line corresponding to a solution that is fully entrained to the new time zone. Under the application of the optimal input shown in panel C, the preliminary large magnitude pulse initially desynchronizes the individual oscillators, as can be seen in the raster plot from panel B and from the individual traces in panel A. A second large magnitude pulse resynchronizes the oscillators with the appropriate phase. This desynchronize-then-resynchronize strategy is not mandated by the optimal control strategy, but it occurs for many choices of Δt_s. It would be of further interest to see if this represents a common mechanism for rapid phase shifting in coupled populations of oscillators.
Source: This figure is adapted from results originally presented in Wilson (2021).

10.4. Optimal waveforms for entrainment and reentrainment

Phase-based methods have been widely used to approach problems that consider entrainment to an exogenous input. For instance, Pyragas et al. (2018), Tanaka (2014) and Zlotnik and Li (2012) consider the problem of designing inputs that are optimal relative to various performance metrics in order to entrain an oscillator to a desired frequency. From an implementation perspective, these strategies generally leverage phase difference coupling strategies discussed in Section 3.1 in order design an appropriate coupling function to achieve entrainment. Provided that inter-oscillator coupling is negligible relative to the entraining stimulus, such strategies can readily accommodate larger populations. Similar strategies were considered in Wilson et al. (2015) to design optimal inputs for entrainment of heterogeneous oscillator populations in the presence of noise.

A variety of problems that consider reentrainment, for instance, following a sudden shift in the timing of the applied stimulus, can also be formulated using phase-based methods. Such an approach was taken in Zlotnik, Chen, Kiss, Tanaka, and Li (2013) where the optimization strategy explicitly considers the rate at which entrainment is achieved. Related strategies were applied in Takata, Kato, and Nakao (2021) which proposes the use of an amplitude feedback term in conjunction with the design of optimal inputs for entrainment in order to yield faster decay rates of the transient solutions and hasten reentrainment.

11. Practical considerations for experimental implementation of control algorithms

The majority of the control techniques discussed in this review have only been developed and tested *in silico*. While such simulations are a

necessary first step to demonstrate proof of concept, it is necessary to keep in mind practical issues pertaining to implementation in living systems. Here we discuss critical assumptions required for implementations of each of the control algorithms discussed in previous sections, focusing on potential issues that would need to be addressed in the transition from computer simulations to experimental illustrations.

11.1. General phase reduction techniques

Phase-based reduction techniques provide a foundation on which the many of the control algorithms discussed in this work are based, including those highlighted in Sections 5, 6, 7, 9, and 10. Such methods consider dynamical behavior in reference to a limit cycle which may be difficult to detect in experimental applications, especially in highly noisy systems. Provided a reference limit cycle can be identified, the direct method (Glass & Mackey, 1988; Netoff et al., 2012; Winfree, 2001) has been successfully used in many applications to measure the phase response curves necessary to represent dynamical behavior in terms of a reduced order model of the form (8).

11.2. General phase-isostable-based reduction techniques

Standard phase reduction strategies can be augmented with isostable coordinates that capture amplitude-based effects. This coordinate system was used for the applications discussed in Sections 9.1, 9.2, 9.3, 10.1, and 10.2. Implementation of these methods requires the inference of isostable response curves that comprise Eq. (A.4). The key challenge here is that isostable coordinates encode for the transient behavior. As such, it is necessary to be able to infer this information during the relaxation to the limit cycle. Strategies akin to the direct method have been proposed in Wilson and Ermentrout (2018b) and Wilson and Ermentrout (2019a) which apply a pulse perturbation at a known phase, consider either the oscillation amplitude of an observable or the change in the interspike interval during the relaxation back to the limit cycle, use this information to provide a pointwise estimate of the isostable response curve for that single phase, and repeat for different initial phases. A related approach proposed in Wilson (2020a) combines a delayed embedding strategy with a least-squares fitting approach to infer phase-isostable-based reduced order models for systems with multiple isostable coordinates. Because isostable coordinates represent level sets of the slowest decaying eigenmodes of the Koopman operator (Budišić, Mohr, & Mezić, 2012; Mezić, 2019), it is likely that dynamic mode decomposition algorithms (Kutz, Brunton, Brunton, & Proctor, 2016; Schmid, 2010; Tu, Rowley, Luchtenberg, Brunton, & Kutz, 2014) could be adapted for phase-isostable-based reduced order models in a data-driven setting.

11.3. The role of heterogeneity

In experimental settings, model-based control frameworks must often consider heterogeneity. As such, experimental implementation of a given control strategy generally requires an initial model identification step (for example, to measure phase response curves and isostable response curves). Note that phase response curves taken from individual neurons within the same brain region can exhibit a large amount of heterogeneity (Ota, Omori, Watanabe, Miyakawa, Okada, & Aonishi, 2011; Wang, Musharoff, Canavier, & Gasparini, 2013). When considering control of populations of neurons, it can be necessary to explicitly consider the heterogeneity within the population. Wilson et al. (2015) and Wilson and Moehlis (2014a) account for heterogeneity by considering an envelope within which phase response curves might fall rather than any single phase response curve. Alternatively, larger populations can be analyzed in terms of their collective behaviors (Kawamura, Nakao, Arai, Kori, & Kuramoto, 2008; Ko & Ermentrout, 2009; Levnajić & Pikovsky, 2010), eliminating the need to consider heterogeneity between individual neurons. Of course, heterogeneity between populations would still need to be considered.

11.4. Model independent control strategies

Model independent control strategies, such as those discussed in Section 4 are generally the easiest to implement in experimental settings. Indeed, the adaptive deep brain stimulation strategy from Section 4.1 has already yielded promising results in multiple human trials (Little, et al., 2013; Meidahl, et al., 2017; Priori et al., 2013; Rosa, et al., 2015). The power of model independent control strategies comes from the fact that no model identification is needed. Provided a useful observable (e.g., LFP power, accelerometer data from patient tremor rhythm, etc.) can be identified, the desired control algorithm can be readily implemented. The lack of an underlying model does come with drawbacks, however. Many of these control algorithms come with parameters that must be tuned in an *ad hoc* manner, and it can be difficult to truly understand the mechanisms underlying the changes produced by the controllers.

11.5. Chaotic desynchronization and cluster control

Considering the chaotic desynchronization strategies discussed in Section 5, experimental implementation would require the measurement of a representative set of phase response curves associated with a population of pathologically synchronized oscillators. While this approach does not explicitly account for heterogeneity among neurons, simulations from Wilson and Moehlis (2014c) and Wilson and Moehlis (2014b) illustrated that such heterogeneity generally increases the rate of desynchronization. Accurate characterization of clustering using the methods discussed in Sections 6.1–6.3 and 6.5 is contingent on the ability to identify phase response curves for a population of oscillators. Alternatively, the coordinate reset strategy from Section 6.4 uses multiple electrodes each applying different inputs to achieve clustering. This control strategy does not necessarily require an underlying model, but rather, requires the identification of an input that can effectively modulate the behavior of each subpopulation.

11.6. Control of phase distributions

Large populations of neurons can often be represented by a probability distribution with dynamics that follow a nonlinear Fokker–Planck equation. Rather than focusing on the behavior of individual neurons, it may be simpler to focus on controlling the temporal evolution of the probability distribution itself. In an experimental setting, it would be necessary to find a meaningful mapping between a given probability distribution and the observable measurements. For example, the control algorithm discussed in Section 7 may be difficult to implement because it requires knowledge of the probability distribution at all times and such a measurement would be difficult to infer from data. Phase-triggered approaches (such the one proposed in Wilson & Moehlis, 2016b) that attempt to relate the phase of a probability distribution to a single observable and accordingly modulate the control input may be more feasible. The open loop control strategies such as those considered in Sections 9.2 and 9.3 would obviate the need for real-time measurement of outputs, but these come with their own drawbacks (i.e., the requirement that undesired (resp., desired) behavior to be suppressed (resp., enhanced) is associated with weakly stable or unstable fixed points).

11.7. Phase randomization

The control algorithms discussed in Section 8 suggest driving a population of synchronized oscillators to a phaseless set using a precomputed stimulus where they can subsequently be desynchronized by inherent noise. The specific implementation requires full knowledge of the underlying model equations and would be computationally difficult to implement in a model neuron with dimension larger than 3. These issues alone would likely preclude the computation of energy optimal

phase randomizing inputs in experimental settings. Nonetheless these results provide proof of concept that targeting either a phaseless set, or a location in phase space where the gradient of the isochrons is large could be a useful objective when considering problems where eliminating synchronization is important. In an experimental setting, rather than computing optimal phase randomizing inputs it may be more feasible to find the required inputs using either trial-and-error or machine learning approaches.

11.8. Open loop versus closed loop methods

Open loop control techniques can be attractive in neurological applications when real-time sensing is difficult. While open loop methods will generally be less efficient than their closed loop counterparts, they are usually much easier to implement, especially in situations where it is difficult to obtain an accurate dynamical model that relates the inputs to the outputs. The techniques from Section 9 consider open loop control problems for modifying the stability of either a weakly stable periodic orbit weakly stable fixed point. Both of these strategies use phase-based reduction methods that come with the usual caveats discussed in Sections 11.1 and 11.2.

12. Conclusions and future outlook

The dynamical equations describing large populations of oscillatory, conductance-based neurons are usually very high dimensional, non-negligibly nonlinear, and capable of displaying a wide variety complex behaviors. As such, a diverse array of control and analysis techniques have been developed to solve control problems involving populations of neural oscillators. This review article discusses a collection of control problems and techniques for applications involving oscillatory neural populations. This review is far from comprehensive, however. Other notable approaches to controlling populations of neurons include Ching and Ritt (2013), Ehrens, Sritharan, and Sarma (2015), Feng, Greenwald, Rabitz, Shea-Brown, and Kosut (2007), Feng, Shea-Brown, Greenwald, Kosut, and Rabitz (2007), Kiss, Rusin, Kori, and Hudson (2007), Schiff (2010) and Weerasinghe, Duchet, Bick, and Bogacz (2021); see also Ritt and Ching (2015), Schiff (2012) and Tass (2007).

Due to the high dimensionality of the underlying dynamical equations, model order reduction is often an imperative first step for mathematical analysis and control design. Many of the applications considered in this review article use phase-based reduction strategies to represent the model equations in a more analytically tractable form. In many experimental neural control applications, however, robust oscillations may not be present or may be difficult to detect. Moreover, the evolution equations and/or the required response curves for the algorithms to work may be challenging to obtain experimentally. And, as is always the case in control theory, the robustness of control algorithms to noise and heterogeneity is often not assured.

Nonetheless, computational studies show that the algorithms presented in this review hold great promise for controlling neural oscillator populations with a variety of control objectives. Our hope is that these successes will motivate more research on how to implement them in experimental and clinical studies, opening the door to more effective and more efficient treatments for Parkinson's disease and other neurological disorders.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This material is based on work supported by National Science Foundation, USA Grants No. NSF-0547606, NSF-1000678, NSF-1264535, NSF-1635542, NSF-1602841, and NSF-1933583.

Appendix A. Model order reduction techniques using phase and isostable coordinates

When considering oscillatory dynamics, phase reduction is one strategy that is often used to understand complex emergent patterns in weakly perturbed oscillators. More recently, a collection of phase–amplitude-based coordinate systems have also been proposed (Castejón et al., 2013; Diekman & Bose, 2016; Letson & Rubin, 2020; Shirasaka et al., 2017; Wedgwood et al., 2013; Wilson & Ermentrout, 2018b) that incorporate information about the transient dynamics during the decay toward the limit cycle.

Here, we provide a review of recent phase–amplitude-coordinate-based reduction strategies that are also considered in the examples in this review article. This review emphasizes the isostable coordinate framework (Mauroy, Mezić, & Moehlis, 2013; Wilson & Moehlis, 2016a) which encodes for the level sets of the slowest decaying eigenmodes of the Koopman operator (Budišić et al., 2012; Mezić, 2019). Below, we explicitly highlight which methods are used in the applications discussed in detail in Sections 9–10.

A.1. Isostable coordinates to encode for amplitude based effects

Isostable coordinates are used to characterize amplitude effects in the model order reduction strategies described in Sections 9.1, 9.2, 9.3, 10.1, and 10.2. Standard phase reduction techniques (8) implicitly assume that the input is sufficiently weak and that the unperturbed relaxation to the limit cycle is fast enough so that the state remains close to the underlying limit cycle at all times. This restriction is often prohibitive in applications where large magnitude inputs are required to achieve a desired control objective. In these situations, it becomes necessary to consider a set of amplitude coordinates that represent the dynamics transverse to the nominal limit cycle.

Different amplitude coordinate frameworks have been proposed in recent years (Castejón et al., 2013; Diekman & Bose, 2016; Letson & Rubin, 2020; Shirasaka et al., 2017; Wedgwood et al., 2013; Wilson & Ermentrout, 2018b), each having established practical utility in a variety of applications. Here, we focus on the isostable coordinate framework (Mauroy et al., 2013; Wilson & Ermentrout, 2018a; Wilson & Moehlis, 2016a). Isostable coordinates are closely related to the Koopman operator (Budišić et al., 2012; Mezić, 2019), which can be used to represent a general nonlinear dynamical system as a linear but potentially infinite dimensional operator. It is often challenging to identify a finite dimensional representation of the full Koopman operator and as such, Koopman eigenfunctions are often used to characterize the underlying dynamics. The isostable coordinate framework considers level sets of the slowest decaying Koopman eigenfunctions.

A constructive definition for slowest decaying isostable coordinate can be obtained by first considering the dynamics near the linearized periodic orbit. To this end, consider a general dynamical system of the form (1) and let $\Delta x = x - x^\gamma(\theta)$ where $x^\gamma(\theta)$ is the intersection of the periodic orbit and the Γ_θ isochron. Linearization with respect to the periodic orbit yields to leading order

$$\frac{d\Delta x}{dt} = DF(\theta)\Delta x, \tag{A.1}$$

where $DF(\theta)$ is the Jacobian evaluated at $x^\gamma(\theta)$ with $\theta(t) = \omega t$. Noting that Eq. (A.1) is periodic and linear time-varying, solutions can be characterized according to Floquet theory as $\Delta x(T) = Y\Delta x(0)$ where Y is the fundamental matrix for initial conditions for which $\theta(x(0)) \approx 0$. Let w_j, v_j, and λ_j correspond to left eigenvectors, right eigenvectors, and associated eigenvalues, respectively, of Y. Additionally, let $\kappa_j = \log(\lambda_j)/T$ denote the Floquet exponent associated with the Floquet multiplier λ_j. Finally, take $\kappa_n = 0$, corresponding to perturbations in the direction of the periodic orbit and sort the other Floquet exponents so that $|\mathrm{Re}(\kappa_k)| \leq |\mathrm{Re}(\kappa_{k+1})|$; assuming that κ_1 is unique, an associated isostable coordinate can be defined according to

$$\psi_1(x) = \lim_{k\to\infty} \left[w_1^T(\eta(t_\Gamma^k, x) - x_0^\Gamma)\exp(-\kappa_1 t_\Gamma^k)\right], \tag{A.2}$$

where $t_{\Gamma_0}^k$ is the time of the kth transversal of the Γ_0 isochron, $\eta(t,x)$ is the flow that maps an initial condition $x(0)$ to $x(t)$, and x_0^Γ is the intersection of the periodic orbit and the Γ_0 level set. While phase coordinates as defined by isochrons encode for the infinite time behavior, the isostable coordinate as defined by (A.2) gives a sense of distance from the periodic orbit. Isostable coordinates associated with more rapidly decaying components of the solution can also be defined, however, there is not always a constructive definition available (Kvalheim & Revzen, 2021). Instead they can be defined implicitly as level sets of their associated Koopman eigenfunctions. Computation of approximations of these higher order eigenfunctions is discussed in Wilson (2020c). Among the advantages of using this coordinate basis is that isostable coordinates decay according to $d\psi_j/dt = \kappa_j\psi_j$ in the entire basin of attraction of the limit cycle; the utility of unperturbed decay becomes clearer when considering phase-isostable-based reduced order equations as discussed in the next section.

A.2. Phase-isostable-based reduced order models

Phase-isostable reduction, as described here, is used in the development of the control strategy discussed in Section 10.1. From the perspective of control theory, the central utility of isostable coordinates is in the augmentation of standard phase-based methodologies in order to obtain a better approximation of the dynamics when the applied input drives the state far from the nominal limit cycle. Indeed, starting from the perturbed system of the form (4), one can transform to isostable coordinates:

$$\begin{aligned}\frac{d\psi_j}{dt} &= \frac{\partial\psi_j}{\partial x}\cdot\frac{dx}{dt}\\ &= \frac{\partial\psi_j}{\partial x}\cdot(F(x)+U(x,t))\\ &= \kappa_j\psi_j + \frac{\partial\psi_j}{\partial x}\cdot U(x,t).\end{aligned} \tag{A.3}$$

Above, the third line is obtained by noting that $d\psi_j/dt = \kappa_j\psi_j$ in the absence of input. Computing $\partial\psi_j/\partial x$ directly on the periodic orbit and combining with the standard phase reduction yields

$$\begin{aligned}\frac{d\theta}{dt} &= \omega + Z(\theta)\cdot U(x^\gamma,t),\\ \frac{d\psi_j}{dt} &= \kappa_j\psi_j + I_j(\theta)\cdot U(x^\gamma,t), \qquad j=1,\dots,\beta,\end{aligned} \tag{A.4}$$

where $I_j(\theta) = \frac{\partial\psi_j}{\partial x}\Big|_{x^\gamma(\theta)}$. Here a subset of the $\beta \le n-1$ most slowly decaying isostable coordinates are considered, and the rest are truncated. Note that because $Z(\theta)$ and $I(\theta)$ are computed as the gradient of the phase and isostable coordinates, respectively, evaluated on the periodic orbit, the phase-isostable formulation from (A.4) is identical to Floquet coordinate formulations that are often used to analyze periodic linear time-varying equations (Jordan & Smith, 2007). Once again, when the input can be written as a rank-1 perturbation so that as $U(x^\gamma,t) = \delta u(t)$ where $\delta\in\mathbb{R}^n$, it is often convenient to rewrite Eq. (A.4) as

$$\begin{aligned}\frac{d\theta}{dt} &= \omega + z(\theta)u(t),\\ \frac{d\psi_j}{dt} &= \kappa_j\psi_j + i_j(\theta)u(t), \qquad j=1,\dots,\beta,\end{aligned} \tag{A.5}$$

where $z(\theta) = Z(\theta)\cdot\delta$ and $i_j(\theta) = I_j(\theta)\cdot\delta$.

Recall that standard phase reduction techniques tend to break down when the state travels too far from the limit cycle. As such, information about the isostable coordinates can be used to limit the distance from the periodic orbit and consequently extend the applicability of the phase reduced equations. Indeed, recent works (Monga & Moehlis, 2019a; Takata et al., 2021) implement various optimal control algorithms on phase-isostable-based reduced order equations in exactly this manner by penalizing the magnitude of the isostable coordinates in the associated cost functional. Eq. (A.4) can be useful in some situations, but is still relatively limited in terms of the amplitude of the inputs that it can accommodate. In cases where even larger magnitude inputs are necessary, additional modifications described in the following sections can be implemented.

A.2.1. Higher order phase and isostable response curves

Computation of phase-isostable reduced order equations to higher orders of accuracy is an essential starting point for the applications discussed in Sections 9.1, 9.2, and 9.3. The standard phase reduction (6) and the augmented phase-isostable reduction (A.4) evaluate gradients of their reduced order coordinates on the periodic orbit. In order to obtain better estimates that retain accuracy as the state is driven far from the periodic orbit, it is convenient to consider a representation in terms of the reduced order isostable coordinates. Following the strategy presented in Wilson (2020c), one can use Eq. (4) as a starting point and change to phase and isostable coordinates:

$$\begin{aligned}\frac{d\theta}{dt} &= \omega + \mathcal{Z}(\theta,\psi_1,\dots,\psi_\beta)\cdot U(x,t),\\ \frac{d\psi_j}{dt} &= \kappa_j\psi_j + \mathcal{I}_j(\theta,\psi_1,\dots,\psi_\beta)\cdot U(x,t), \qquad j=1,\dots,\beta,\end{aligned} \tag{A.6}$$

where $\mathcal{Z} \equiv \frac{\partial\theta}{\partial x}$ and $\mathcal{I}_j \equiv \frac{\partial\psi_j}{\partial x}$ are not evaluated at $x^\gamma(\theta)$, but rather at $x(\theta,\psi_1,\dots,\psi_\beta)$. Once again, the rapidly decaying isostable coordinates $\psi_{\beta+1},\dots,\psi_{n-1}$ are assumed to be zero and truncated, thereby yielding a reduced order set of equations. The constraint that the state x remain close to x^γ is lifted in the formulation (A.6); as such, the state must also be represented in terms of the reduced order coordinates:

$$x \approx x^\gamma(\theta) + \mathcal{G}(\theta,\psi_1,\dots,\psi_\beta), \tag{A.7}$$

where $\mathcal{G}$ gives the deviation between the current state and $x^\gamma(\theta)$. Asymptotically expanding $\mathcal{Z}$, $\mathcal{G}$ and each $\mathcal{I}_j$ in a basis of the isostable coordinates yields

$$\begin{aligned}\mathcal{G}(\theta,\psi_1,\dots,\psi_\beta) &\approx \sum_{k=1}^{\beta}\left[\psi_k g^k(\theta)\right] + \sum_{j=1}^{\beta}\sum_{k=1}^{j}\left[\psi_j\psi_k g^{jk}(\theta)\right] + \dots,\\ \mathcal{Z}(\theta,\psi_1,\dots,\psi_\beta) &\approx Z(\theta) + \sum_{k=1}^{\beta}\left[\psi_k Z^k(\theta)\right] + \sum_{j=1}^{\beta}\sum_{k=1}^{j}\left[\psi_j\psi_k Z^{jk}(\theta)\right] + \dots,\\ \mathcal{I}_n(\theta,\psi_1,\dots,\psi_\beta) &\approx I_n(\theta) + \sum_{k=1}^{\beta}\left[\psi_k I_n^k(\theta)\right] + \sum_{j=1}^{\beta}\sum_{k=1}^{j}\left[\psi_j\psi_k I_n^{jk}(\theta)\right] + \dots,\end{aligned} \tag{A.8}$$

for $n=1,\dots,\beta$. Above, $g^k(\theta)$ is a Floquet eigenfunction of the linearized Eq. (A.1) associated with the Floquet exponent κ_k. Terms of the form g^{jk} provide higher order corrections from the asymptotic expansion of $\mathcal{G}$ in the basis of isostable coordinates. Likewise, $Z(\theta)$ (resp., $I_j(\theta)$) is the phase response curve (resp., isostable response curve), while terms of the form Z^k, Z^{jk}, … (resp., I_j^k, I_j^{jk}, …) provide higher order corrections. By computing the required terms up to a desired order of accuracy in the expansion in the isostable coordinates, a phase-isostable-based reduced order model can be obtained. Indeed, Fig. 14 illustrates the accuracy of isochrons approximated from Eq. (A.6) to the indicated order of accuracy when considering a dynamical model with oscillations that emerge due to a Hopf bifurcation. Details about the computation of the necessary terms of the higher accuracy expansions are discussed in Wilson (2020c).

A.3. Adaptive phase-isostable reduction

The adaptive phase-isostable reduction is used as a starting point to pose and solve the optimal control problem discussed in Section 10.2. Asymptotic expansion in a basis of isostable coordinates provides one method of obtaining more accurate characterizations of the reduced order dynamics. However, diminishing returns are typically observed at higher orders of accuracy, and these reduction strategies still break down when the underlying state travels too far from the underlying limit cycle. Recent work (Wilson, 2021, 2022) considers a continuous family of limit cycles that result when using different parameter sets. The strategy summarized below actively chooses an appropriate nominal parameter set that keeps the state close to the underlying limit cycle so that truncation errors can be mitigated, resulting in highly accurate reduced order

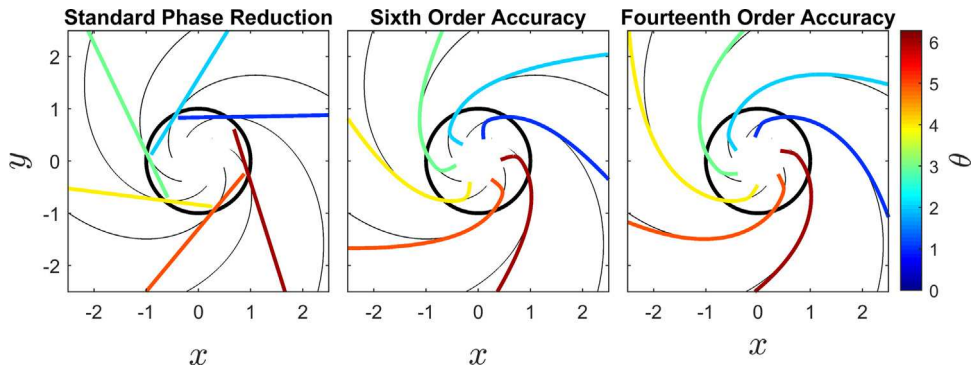

Fig. 14. For a 2-dimensional model with oscillations that emerge due to a Hopf bifurcation, the necessary terms of the phase-isostable formulation (A.6) are computed to the indicated order of accuracy. Isochrons are subsequently inferred as level sets of the phase. Colored lines represent individual isochrons and black lines show the true values of the isochrons computed directly according to the definition (2). The approximation becomes better as higher orders of accuracy are used. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: This figure is adapted from results originally presented in Wilson (2020c).

models. Similar strategies were considered in Kurebayashi, Shirasaka, and Nakao (2013) and Park and Ermentrout (2016), which limited the analysis to dynamical systems with slowly varying parameters.

Implementation of the adaptive phase-isostable reduction requires the explicit consideration of a parameter set $p \in \mathbb{R}^m$. As such, we will rewrite the perturbed system from Eq. (4) as

$$\frac{dx}{dt} = F(x, p_0) + U(x, t). \tag{A.9}$$

Here, Eqs. (4) and (A.9) are identical, with the only difference being in the explicit consideration of a nominal parameter set p_0. Suppose that in some allowable range of parameters $p \in P$, the p-limit cycle $x^\gamma(\theta, p)$ emerges in the absence of input. Here, we invoke the notion of a generalized phase $\theta(x, p)$ and a set of generalized isostable coordinates $\psi_1(x, p), \ldots, \psi_\beta(x, p)$, each associated with both the state and the p-limit cycle. To proceed, consider a rewritten version of Eq. (A.9)

$$\frac{dx}{dt} = F(x, p) + U_e(t, p, x), \tag{A.10}$$

with the extended input

$$U_e(t, p, x) = U(x, t) + F(x, p_0) - F(x, p). \tag{A.11}$$

One can verify that Eqs. (A.9) and (A.10) are indeed identical. Nonetheless, the term $F(x, p)$ gives the dynamics of the underlying system using the parameter set p, and U_e captures both the externally applied input and the mismatch in the unperturbed dynamics between the two parameter sets. Supposing that $x^\gamma(\theta, p)$ is continuously differentiable with respect to both θ and p and that $\theta(x, p)$ and each $\psi_j(x, p)$ are appropriately defined so that they are continuously differentiable with respect to x and p, one can consider p to be a free variable yielding the transformation

$$\begin{aligned}
\frac{d}{dt}\theta(x, p) &= \frac{\partial \theta}{\partial x} \cdot \frac{dx}{dt} + \frac{\partial \theta}{\partial p} \cdot \frac{dp}{dt}, \\
\frac{d}{dt}\psi_j(x, p) &= \frac{\partial \psi_j}{\partial x} \cdot \frac{dx}{dt} + \frac{\partial \psi_j}{\partial p} \cdot \frac{dp}{dt}, \qquad j = 1, \ldots, \beta, \\
\frac{dp}{dt} &= G_p(p, \theta, \psi_1, \ldots, \psi_\beta),
\end{aligned} \tag{A.12}$$

where G_p is an update function for p that will be discussed momentarily. Notice the similarities between Eq. (A.12) and both (5) and (A.3); Eqs. (5) and (A.3) assume a constant parameter set and as such, the terms $\frac{\partial \theta}{\partial x} \cdot \frac{dx}{dt}$ and $\frac{\partial \psi_j}{\partial x} \cdot \frac{dx}{dt}$ are identical to the terms from Eqs. (5) and (A.3). As shown in Wilson (2022), by considering the remaining terms and continuing to simplify (A.12), to leading order one can write

$$\begin{aligned}
\frac{d\theta}{dt} &= \omega(p) + Z(\theta, p) \cdot U_e(t, p, x) + D(\theta, p) \cdot \frac{dp}{dt}, \\
\frac{d\psi_j}{dt} &= \kappa_j(p)\psi_j + I_j(\theta, p) U_e(t, p, x) + E_j(\theta, p) \cdot \frac{dp}{dt}, \qquad j = 1, \ldots, \beta,
\end{aligned}$$

$$\frac{dp}{dt} = G_p(p, \theta, \psi_1, \ldots, \psi_\beta). \tag{A.13}$$

Here, $\omega(p)$ and $\kappa_j(p)$ are natural frequencies and Floquet exponents associated with the p-limit cycle, and $Z(\theta, p)$ and each $I_j(\theta, p)$ are given by the gradients of $\theta(x, p)$ and $\psi(x, p)$ evaluated on the p-limit cycle. Additionally, the ith component of $D(\theta, p)$ is identical to $-\frac{\partial x^\gamma}{\partial p_i} \cdot Z(\theta, p)$ where $\partial x^\gamma / \partial p_i|_{\theta_0, p} \equiv \lim_{a \to 0}(x^\gamma(\theta_0, p + e_i a) - x^\gamma(\theta_0, p))/a$, with e_i being the ith component of the standard unit basis. Likewise, the ith component of $E_j(\theta, p)$ is given by $-\frac{\partial x^\gamma}{\partial p_i} \cdot I_j(\theta, p)$. Furthermore, to leading order of accuracy in the non-truncated isostable coordinates, the state is well-approximated given by

$$x \approx x^\gamma(\theta, p) + \sum_{j=1}^{\beta} \psi_j g^j(\theta, p), \tag{A.14}$$

where $g^j(\theta, p)$ is a Floquet eigenfunction associated with the p-limit cycle.

Provided a function G_p can be obtained that keeps each ψ_j sufficiently small, the truncation errors associated with neglecting the higher order terms from the phase and isostable dynamics can be kept negligibly small yielding a reduced order set of equations that is substantially more accurate than standard phase reduction strategies. This general idea is highlighted in the schematic shown in Fig. 15. General heuristics of the design of G_p are discussed in Wilson (2022). For instance, when only one non-truncated, real-valued isostable coordinate, ψ_1, is required and $p \in \mathbb{R}$, one can often use

$$G_p(p, \theta, \psi_1) = -\alpha \psi_1 E_1(\theta, p), \tag{A.15}$$

where α is a positive constant so that the isostable coordinate dynamics from Eq. (A.13) become

$$\frac{d\psi_1}{dt} = (\kappa - \alpha E_1^2(\theta, p))\psi_1 + I_1(\theta, p) U_e(t, p, x). \tag{A.16}$$

This choice of G_p serves to drive the isostable coordinate to smaller magnitude values with an effect that is proportional to the isostable coordinate itself. Alternatively, in the special case that $E_1(\theta, p)$ is bounded away from zero, one can also choose

$$G_p(p, \theta, \psi_1) = -\frac{I_1(\theta, p) U_e(t, p, x)}{E_1(\theta, p)}. \tag{A.17}$$

Substituting the above relation into Eq. (A.13) the isostable dynamics become $d\psi_1/dt = \kappa_1(p)\psi_1$ and can ultimately be truncated, further reducing the dimension of the reduction.

Appendix B. Models

In this appendix, we give details of the mathematical models considered in the main text.

Thalamic neuron model

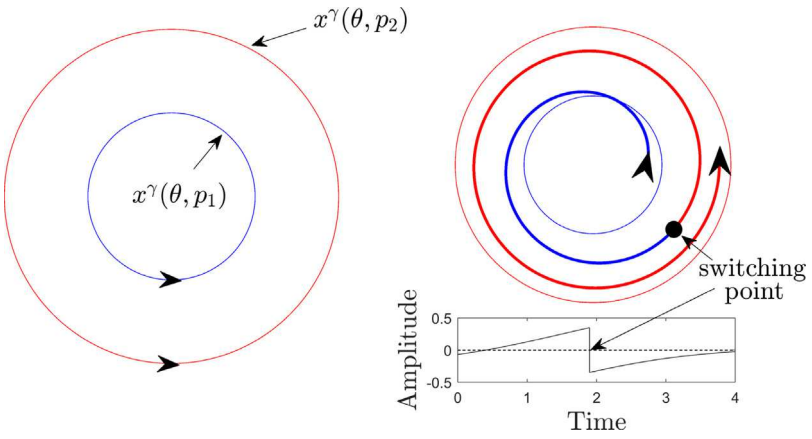

Fig. 15. An schematic depicting the adaptive reduction scheme. The left side shows two adjacent periodic orbits that emerge for two different parameter sets p_1 and p_2. The right side shows an example trajectory. This trajectory starts closer to the blue orbit but gradually travels towards the red orbit. By converting phase and isostable coordinates from the blue orbit to those of the red orbit at an appropriate time, the amplitude coordinate can be kept small, subsequently limiting errors induced by truncating higher order terms of the phase-isostable reduction. The adaptive phase-isostable reduction implements this general idea in a continuous manner, i.e., by continuously updating the adaptive parameters in order to keep the isostable coordinates small.

The thalamic neuron model is given as

$$\begin{aligned}
\dot{v} &= \frac{-I_L - I_{Na} - I_K - I_T + I_b}{C_m} + u(t),\\
\dot{h} &= \frac{h_\infty - h}{\tau_h},\\
\dot{r} &= \frac{r_\infty - r}{\tau_r},
\end{aligned}$$

where

$$\begin{aligned}
h_\infty &= 1/(1+\exp((v+41)/4)),\\
r_\infty &= 1/(1+\exp((v+84)/4)),\\
\alpha_h &= 0.128\exp(-(v+46)/18),\\
\beta_h &= 4/(1+\exp(-(v+23)/5)),\\
\tau_h &= 1/(\alpha_h+\beta_h),\\
\tau_r &= (28+\exp(-(v+25)/10.5)),\\
m_\infty &= 1/(1+\exp(-(v+37)/7)),\\
p_\infty &= 1/(1+\exp(-(v+60)/6.2)),\\
I_L &= g_L(v-e_L),\\
I_{Na} &= g_{Na}({m_\infty}^3)h(v-e_{Na}),\\
I_K &= g_K((0.75(1-h))^4)(v-e_K),\\
I_T &= g_T(p_\infty^2)r(v-e_T),
\end{aligned}$$

$$C_m = 1, \quad g_L = 0.05, \quad e_L = -70, \quad g_{Na} = 3, \quad e_{Na} = 50,$$

$$g_K = 5, \quad e_K = -90, \quad g_T = 5, \quad e_T = 0, \quad I_b = 5.$$

In Section 5.1, we consider:

$$\begin{aligned}
\dot{V}_i &= (-I_L(V_i) - I_{Na}(V_i,h_i) - I_K(V_i,h_i) - I_T(V_i,r_i)\\
&\quad + I_{SM} + \frac{1}{N}\sum_{i=1}^{N}\alpha_{ij}(V_j - V_i) + u(t) + \eta_i(t))/C,\\
\dot{h}_i &= (h_\infty(V_i) - h_i)/\tau_h(V_i),\\
\dot{r}_i &= (r_\infty(V_i) - r_i)/\tau_r(V_i), \quad i = 1,\dots,N.
\end{aligned} \tag{B.1}$$

Here we have augmented the voltage equation by additively including electrotonic coupling (Johnston & Wu, 1995), DBS input, and Gaussian white noise. Other coupling paradigms (e.g., synaptic coupling) could also be implemented instead of electrotonic coupling. Here, N is the total number of neurons, V_i, h_i, and r_i are membrane voltage and gating variables for neuron i, α_{ij} characterizes the coupling strength between electrotonically coupled neurons i and j, with $\alpha_{ij} = \alpha_{ji}$ and $\alpha_{ii} = 0$ for all i, $\eta_i(t) = \sqrt{2D}\mathcal{N}(0,1)$ is the i.i.d. noise associated with each neuron, assumed to be zero-mean Gaussian white noise with variance $2D$, and $u(t) = I(t)/C$ represents a common control input. In this equation I_{SM} represents the baseline current which we take to be $5\mu\text{A/cm}^2$.

Hodgkin–Huxley neuron model

A model for a coupled population of Hodgkin–Huxley model neurons is given below. Note, that each neuron from the population uses the two-dimensional reduction proposed in Rinzel (1985)

$$\begin{aligned}
\dot{V}_i &= f_V(V_i,n_i) + \frac{1}{N}\sum_{i=1}^{n}\alpha_{ij}(V_j - V_i) + u(t) + \eta_i(t),\\
\dot{n}_i &= f_n(V_i,n_i),
\end{aligned} \tag{B.2}$$

for $i = 1,\dots,N$. Here, V_i and n_i are membrane voltage and gating variables for neuron i, α_{ij} characterizes the coupling strength between electrotonically coupled neurons i and j, $\eta_i(t) \in \sqrt{2D}\mathcal{N}(0,1)$ is zero-mean Gaussian white noise associated with each neuron with variance $2D$, $u(t) = I(t)/C$ represents a common control input in $\mu A/\mu F$ where $C = 1\mu F/\text{cm}^2$ is the membrane capacitance and $I(t)$ is a DBS input current. Auxiliary equations are given by:

$$\begin{aligned}
f_V &= (I_b - \bar{g}_{Na}[m_\infty(V)]^3(0.8-n)(V-V_{Na})\\
&\quad - \bar{g}_K n^4(V-V_k) - \bar{g}_L(V-V_L))/C,\\
f_n &= a_n(v)(1-n) - b_n(V)n.
\end{aligned}$$

Other functions and parameters for the reduced model are:

$$\begin{aligned}
m(V) &= \frac{a_m(V)}{a_m(V)+b_m(V)},\\
a_m(V) &= 0.1(V+40)/(1-\exp(-(V+40)/10)),\\
b_m(V) &= 4\exp(-(V+65)/18),\\
a_n(V) &= 0.01(V+55)/(1-\exp(-(V+55)/10)),\\
b_n(V) &= 0.125\exp(-(V+65)/80),\\
V_{Na} &= 50\text{ mV}, V_K = -77\text{ mV}, V_L = -54.4\text{ mV},\\
\bar{g}_{Na} &= 120\text{ mS/cm}^2, \bar{g}_K = 36\text{ ms/cm}^2,\\
\bar{g}_L &= 0.3\text{ mS/cm}^2, I_b = 10\ \mu\text{A/cm}^2.
\end{aligned}$$

Here, $\bar{g}_{Na}$, $\bar{g}_K$, and $\bar{g}_L$ represent the conductances of the sodium, potassium and leakage channels, respectively, and V_{Na}, V_K, and V_L are their respective reversal potentials. I_b, is a baseline current chosen to ensure that the neuron is in an oscillatory (periodically spiking) regime.

For this set of parameters, the natural period of oscillation for a single neuron is 11.81 ms in the absence of noise and control input.

Coupled Oscillations of SCN Neurons

A model describing gene expression of N coupled oscillators from the suprachiasmatic nucleus is given below. This model is adapted from Gonze, Bernard, Waltermann, Kramer, and Herzel (2005):

$$\begin{aligned}
\dot{a}_i &= h_1 \frac{K_1^n}{K_1^n + c_i^n} - h_2 \frac{a_i}{K_2 + a_i} + h_c \frac{KF(t)}{K_c + KF(t)} \\
&\quad + S_i \left[L(t_s) + u(t) \right] + \sqrt{2D}\eta_i, \\
\dot{b}_i &= h_3 a_i - h_4 \frac{b_i}{K_4 + b_i}, \\
\dot{c}_i &= h_5 b_i - h_6 \frac{c_i}{K_6 + c_i}, \\
\dot{d}_i &= h_7 a_i - h_8 \frac{d_i}{K_8 + d_i}, \quad i = 1, \ldots, N, \\
\dot{t}_s &= 1, \\
F(t) &= (1/N) \sum_{j=1}^{N} d_i(t).
\end{aligned} \tag{B.3}$$

Above, variables a_i, b_i and c_i represent the concentration of the mRNA clock gene, the associated protein and the nuclear form of the protein, respectively, for the ith neuron. d_i represents the concentration of a neurotransmitter that sets the coupling. Diffusion is assumed to occur on fast enough time scales so that the average value $F(t)$ can be used to determine the coupling. $L(t_s) \in \mathbb{S}^1$ is the environmental time that sets the nominal 24-hour light-dark cycle governed by

$$L_{\text{nom}}(t_s) = L_0 \left[\frac{1}{1 + \exp(-\nu(t_s - 6))} - \frac{1}{1 + \exp(-\nu(t_s - 18))} \right], \tag{B.4}$$

where L_0 is the nominal light intensity during daylight hours and ν sets the transition rate of the light intensity between day and night. A control input $u(t)$ is applied identically to each oscillator and can be used to implement a light exposure or avoidance strategy to promote reentrainment after a sudden shift to the environmental time. $\sqrt{2D}\eta_i$ is an independent and identically distributed zero-mean white noise process with intensity D. Nominal parameters are chosen so that the unperturbed population settles to a near 24-hour periodic oscillation in steady state. S_i represents the sensitivity to light of neuron i and is used to incorporate heterogeneity into the model. Additional heterogeneity is incorporated by drawing h_1, h_2, h_3, h_4, h_5, and h_6 from normal distributions with a mean equal to the nominal parameter value. This model is used in the applications considered in Section 10. For a full description of the relevant model parameters, the reader is referred to either (Wilson, 2020a) or (Wilson, 2021).

References

Abouzeid, A., & Ermentrout, B. (2009). Type-II phase resetting curve is optimal for stochastic synchrony. *Physical Review E*, *80*, Article 011911.

Adamchic, I., Hauptmann, C., Barnikol, U. B., Pawelczyk, N., Popovych, O., Barnikol, T. T., et al. (2014). Coordinated reset neuromodulation for Parkinson's disease: proof-of-concept study. *Movement Disorders*, *29*(13), 1679–1684.

Anderson, R. W., Farokhniaee, A., Gunalan, K., Howell, B., & McIntyre, C. C. (2018). Action potential initiation, propagation, and cortical invasion in the hyperdirect pathway during subthalamic deep brain stimulation. *Brain Stimulation*, *11*(5), 1140–1150.

Ashwin, P., & Swift, J. W. (1992). The dynamics of n weakly coupled identical oscillators. *Journal of Nonlinear Science*, *2*, 69–108.

Åström, K. J., & Bernhardsson, B. (2003). Systems with Lebesgue sampling. In A. Rantzer, & C. I. Byrnes (Eds.), *Directions in mathematical systems theory and optimization. Vol. XIII*. Springer.

Athans, M., & Falb, P. L. (1966). *Optimal control: an introduction to the theory and its applications*. McGraw-Hill New York.

Azodi-Avval, R., & Gharabaghi, A. (2015). Phase-dependent modulation as a novel approach for therapeutic brain stimulation. *Frontiers in Computational Neuroscience*, *9*, 26.

Bagheri, N., Stelling, J., & III, F. J. Doyle (2008). Circadian phase resetting via single and multiple control targets. *PLoS Computational Biology*, *4*(7), Article e1000104.

Bardi, M., & Capuzzo-Dolcetta, I. (1997). *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Birkhauser.

Benabid, A. L., Pollak, P., Louveau, A., Henry, S., & Rougemont, J. De (1987). Combined (thalamotomy and stimulation) stereotactic surgery of the VIM thalamic nucleus for bilateral Parkinson disease. *Stereotactic and Functional Neurosurgery*, *50*(1–6), 344–346.

Bick, C., Goodfellow, M., Laing, C. R., & Martens, E. A. (2020). Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: a review. *The Journal of Mathematical Neuroscience*, *10*, 1–43.

Braiman, Y., & Goldhirsch, I. (1991). Taming chaotic dynamics with weak periodic perturbations. *Physical Review Letters*, *66*(20), 2545.

Brown, E., Holmes, P., & Moehlis, J. (2003). Globally coupled oscillator networks. In E. Kaplan, J. Marsden, & K. R. Sreenivasan (Eds.), *Nonlinear science: a celebratory volume in honor of larry sirovich* (pp. 183–215). Springer.

Brown, E., Moehlis, J., & Holmes, P. (2004). On the phase reduction and response dynamics of neural oscillator populations. *Neural Computation*, *16*, 673–715.

Budišić, M., Mohr, R., & Mezić, I. (2012). Applied Koopmanism. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *22*(4), Article 047510.

Cagnan, H., Pedrosa, D., Little, S., Pogosyan, A., Cheeran, B., Aziz, T., et al. (2017). Stimulating at the right time: phase-specific deep brain stimulation. *Brain*, *140*(1), 132–145.

Castejón, O., Guillamon, A., & Huguet, G. (2013). Phase-amplitude response functions for transient-state stimuli. *The Journal of Mathematical Neuroscience*, *3*, 1–26.

Chacon, R., & Bejarano, J. D. (1993). Routes to suppressing chaos by weak periodic perturbations. *Physical Review Letters*, *71*(19), 3103.

Chen, C. C., Litvak, V., Gilbertson, T., Kuhn, A., Lu, C. S., Lee, S. T., et al. (2007). Excessive synchronization of basal ganglia neurons at 20 Hz slows movement in Parkinson's disease. *Experimental Neurology*, *205*(1), 214–221.

Chesson, A. L., Littner, M., Davila, D., Anderson, W. M., Grigg-Damberger, M., Hartse, K., et al. (1999). Practice parameters for the use of light therapy in the treatment of sleep disorders. *Sleep*, *22*(5), 641–660.

Ching, S., & Ritt, J. T. (2013). Control strategies for underactuated neural ensembles driven by optogenetic stimulation. *Frontiers in Neural Circuits*, *7*, Article 54.

Coddington, E. A., & Levinson, N. (1955). *Theory of ordinary differential equations*. New York: McGraw-Hill.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, *2*, 303–314.

Daido, H. (1996). Onset of cooperative entrainment in limit-cycle oscillators with uniform all-to-all interactions: bifurcation of the order function. *Physica D*, *91*, 24–66.

Danzl, P., Hespanha, J., & Moehlis, J. (2009). Event-based minimum-time control of oscillatory neuron models: phase randomization, maximal spike rate increase, and desynchronization. *Biological Cybernetics*, *101*, 387–399.

Dean, D. A., Forger, D. B., & Klerman, E. B. (2009). Taking the lag out of jet lag through model-based schedule design. *PLoS Computational Biology*, *5*(6), Article e1000418.

Demmel, J. W. (1997). *Applied numerical linear algebra*. Philadelphia: SIAM.

Diekman, C. O., & Bose, A. (2016). Entrainment maps: A new tool for understanding properties of circadian oscillator models. *Journal of Biological Rhythms*, *31*(6), 598–616.

Duchet, B., Weerasinghe, G., Bick, C., & Bogacz, R. (2021). Optimizing deep brain stimulation based on isostable amplitude in essential tremor patient models. *Journal of Neural Engineering*, *18*(4), Article 046023.

Duchet, B., Weerasinghe, G., Cagnan, H., Brown, P., Bick, C., & Bogacz, R. (2020). Phase-dependence of response curves to deep brain stimulation and their relationship: from essential tremor patient data to a Wilson–Cowan model. *The Journal of Mathematical Neuroscience*, *10*(1), 1–39.

Dumont, G., & Gutkin, B. (2019). Macroscopic phase resetting-curves determine oscillatory coherence and signal transfer in inter-coupled neural circuits. *PLoS Computational Biology*, *15*(5), Article e1007019.

Ehrens, D., Sritharan, D., & Sarma, S. V. (2015). Closed-loop control of a fragile network: application to seizure-like dynamics of an epilepsy model. *Frontiers in Neuroscience*, *9*, Article 58.

Ermentrout, B. (1996). Type I membranes, phase resetting curves, and synchrony. *Neural Computation*, *8*, 979–1001.

Ermentrout, G. B., & Kopell, N. (1991). Multiple pulse interactions and averaging in coupled neural oscillators. *Journal of Mathematical Biology*, *29*, 195–217.

Ermentrout, G. B., & Terman, D. H. (2010). *Mathematical foundations of neuroscience*. Berlin: Springer.

Faramarzi, S., & Netoff, T. I. (2021). Closed-loop neuromodulation for clustering neuronal populations. *Journal of Neurophysiology*, *125*, 248–255.

Feng, X.-J., Greenwald, B., Rabitz, H., Shea-Brown, E., & Kosut, R. (2007). Toward closed-loop optimization of deep brain stimulation for parkinson's disease: concepts and lessons from a computational model. *Journal of Neural Engineering*, *4*, L14–L21.

Feng, X.-J., Shea-Brown, E., Greenwald, B., Kosut, R., & Rabitz, H. (2007). Optimal deep brain stimulation of the subthalamic nucleus - a computational study. *Journal of Computational Neuroscience*, *23*, 265–282.

Forger, D. B., & Paydarfar, D. (2004). Starting, stopping, and resetting biological oscillators: in search of optimal perturbations. *Journal of Theoretical Biology*, *230*, 521–532.

Gardiner, C. W. (2004). *Handbook of stochastic methods* (3rd ed.). Berlin: Springer.

Gengel, E., Teichmann, E., Rosenblum, M., & Pikovsky, A. (2020). High-order phase reduction for coupled oscillators. *Journal of Physics: Complexity*, *2*(1), Article 015005.

Glass, L., & Mackey, M. C. (1988). *From clocks to chaos: the rhythms of life*. Princeton: Princeton University Press.

Golombek, D. A., & Rosenstein, R. E. (2010). Physiology of circadian entrainment. *Physiological Reviews*, *90*(3), 1063–1102.

Gonze, D., Bernard, S., Waltermann, C., Kramer, A., & Herzel, H. (2005). Spontaneous synchronization of coupled circadian oscillators. *Biophysical Journal*, *89*(1), 120–129.

Grado, L. L., Johnson, M. D., & Netoff, T. I. (2018). Bayeslan adaptive dual control of deep brain stimulation in a computational model of Parkinson's disease. *PLoS Computational Biology*, *14*(12), Article e1006606.

Greenberg, B. D., Malone, D. A., Friehs, G. M., Rezai, A. R., Kubu, C. S., Malloy, P. F., et al. (2006). Three-year outcomes in deep brain stimulation for highly resistant obsessive–compulsive disorder. *Neuropsychopharmacology*, *31*(11), 2384–2393.

Guckenheimer, J. (1975). Isochrons and phaseless sets. *Journal of Mathematical Biology*, *1*, 259–273.

Guckenheimer, J., & Holmes, P. J. (1983). *Nonlinear oscillations, dynamical systems and bifurcations of vector fields*. New York: Springer-Verlag.

Hahn, P. J., & McIntyre, C. C. (2010). Modeling shifts in the rate and pattern of subthalamopallidal network activity during deep brain stimulation. *Journal of Computational Neuroscience*, *28*(3), 425–441.

Hammond, C., Bergman, H., & Brown, P. (2007). Pathological synchronization in Parkinson's disease: networks, models, and treatments. *Trends in Neurosciences*, *30*, 357–364.

Hansel, D., Mato, G., & Meunier, C. (1995). Synchrony in excitatory neural networks. *Neural Computation*, *7*, 307–337.

Holt, A. B., Kormann, E., Gulberti, A., Pötter-Nerger, M., McNamara, C. G., Cagnan, H., et al. (2019). Phase-dependent suppression of beta oscillations in Parkinson's disease patients. *Journal of Neuroscience*, *39*(6), 1119–1134.

Holt, A. B., Wilson, D., Shinn, M., Moehlis, J., & Netoff, T. I. (2016). Phasic burst stimulation: a closed-loop approach to tuning deep brain stimulation parameters for Parkinson's disease. *PLOS Computational Biology*, *13*, Article e1005011.

Hoppensteadt, F. C., & Izhikevich, E. M. (1997). *Weakly connected neural networks*. New York: Springer-Verlag.

Izhikevich, E. M. (2007). *Dynamical systems in neuroscience: the geometry of excitability and bursting*. London: MIT Press.

Johnston, D., & Wu, S. M.-S. (1995). *Foundations of cellular neurophysiology*. Cambridge, MA: MIT Press.

Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, *13*(4), 455–492.

Jordan, D., & Smith, P. (2007). *Nonlinear ordinary differential equations: an introduction for scientists and engineers. Vol. 10*. Oxford: Oxford University Press.

Josic, K., Shea-Brown, E. T., & Moehlis, J. (2006). Isochron. *Scholarpedia*, *1*(8), 1361.

Kawamura, Y., Nakao, H., Arai, K., Kori, H., & Kuramoto, Y. (2008). Collective phase sensitivity. *Physical Review Letters*, *101*(2), Article 024101.

Khalil, H. K. (2002). *Nonlinear systems*. Upper Saddle River, NJ: Prentice Hall.

Kirk, D. E. (1970). *Optimal control theory: an introduction*. New York: Dover Publications Inc.

Kirk, D. (1998). *Optimal control theory*. New York: Dover Publications.

Kiss, I. Z., Rusin, C. G., Kori, H., & Hudson, J. L. (2007). Engineering complex dynamical structures: sequential patterns and desynchronization. *Science*, *316*, 1886–1889.

Ko, T. W., & Ermentrout, G. B. (2009). Phase-response curves of coupled oscillators. *Physical Review E*, *79*(1), Article 016211.

Koller, W., Pahwa, R., Busenbark, K., Hubble, J., Wilkinson, S., Lang, A., et al. (1997). High-frequency unilateral thalamic stimulation in the treatment of essential and parkinsonian tremor. *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, *42*(3), 292–299.

Kuelbs, D., Dunefsky, J., Monga, B., & Moehlis, J. (2020). Analysis of neural clusters due to deep brain stimulation. *Biological Cybernetics*, *114*, 589–607.

Kuhn, A. A., Kempf, F., Brucke, C., Gaynor Doyle, L., Martinez-Torres, I., Pogosyan, A., et al. (2008). High-frequency stimulation of the subthalamic nucleus suppresses oscillatory beta activity in patients with Parkinson's disease in parallel with improvement in motor performance. *The Journal of Neuroscience*, *28*(24), 6165–6173.

Kuncel, A. M., & Grill, W. M. (2004). Selection of stimulus parameters for deep brain stimulation. *Clinical Neurophysiology*, *115*(11), 2431–2441.

Kuramoto, Y. (1984). *Chemical oscillations, waves, and turbulence*. Berlin: Springer.

Kuramoto, Y. (1997). Phase-and center-manifold reductions for large populations of coupled oscillators with application to non-locally coupled systems. *International Journal of Bifurcation and Chaos*, *7*(04), 789–805.

Kurebayashi, W., Shirasaka, S., & Nakao, H. (2013). Phase reduction method for strongly perturbed limit cycle oscillators. *Physical Review Letters*, *111*(21), Article 214101.

Kuritz, K., Zeng, S., & Allgöwer, F. (2019). Ensemble controllability of cellular oscillators. *IEEE Control Systems Letters*, *3*, 206–301.

Kutz, J. N., Brunton, S. L., Brunton, B. W., & Proctor, J. L. (2016). *Dynamic mode decomposition: data-driven modeling of complex systems*. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Kvalheim, M. D., & Revzen, S. (2021). Existence and uniqueness of global Koopman eigenfunctions for stable fixed points and periodic orbits. *Physica D: Nonlinear Phenomena*, *425*, Article 132959.

Laxton, A. W., Tang-Wai, D. F., McAndrews, M. P., Zumsteg, D., Wennberg, R., Keren, R., et al. (2010). A phase I trial of deep brain stimulation of memory circuits in alzheimer's disease. *Annals of Neurology*, *68*(4), 521–534.

Letson, B., & Rubin, J. E. (2020). LOR for analysis of periodic dynamics: A one-stop shop approach. *SIAM Journal on Applied Dynamical Systems*, *19*(1), 58–84.

Levnajić, Z., & Pikovsky, A. (2010). Phase resetting of collective rhythm in ensembles of oscillators. *Physical Review E*, *82*(5), Article 056202.

Levy, R., Hutchison, W. D., Lozano, A. M., & Dostrovsky, J. O. (2000). High-frequency synchronization of neuronal activity in the subthalamic nucleus of parkinsonian patients with limb tremor. *The Journal of Neuroscience*, *20*(20), 7766–7775.

Li, M. C. H., & Cook, M. J. (2018). Deep brain stimulation for drug-resistant epilepsy. *Epilepsia*, *59*(2), 273–290.

Li, J.-S., Dasanayake, I., & Ruths, J. (2013). Control and synchronization of neuron ensembles. *IEEE Transactions on Automatic Control*, *58*, 1919–1930.

Lima, R., & Pettini, M. (1990). Suppression of chaos by resonant parametric perturbations. *Physical Review A*, *41*(2), 726.

Little, S., Pogosyan, A., Neal, S., Zavala, B., Zrinzo, L., Hariz, M., et al. (2013). Adaptive deep brain stimulation in advanced Parkinson disease. *Annals of Neurology*, *74*(3), 449–457.

Lozano, A. M., & Lipsman, N. (2013). Probing and regulating dysfunctional circuits using deep brain stimulation. *Neuron*, *77*(3), 406–424.

Lu, Z., Klein-Cardeña, K., Lee, S., Antonsen, T. M., Girvan, M., & Ott, E. (2016). Resynchronization of circadian oscillators and the east-west asymmetry of jet-lag. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *26*(9), Article 094811.

Lücken, L., Yanchuk, S., Popovych, O. V., & Tass, Peter A (2013). Desynchronization boost by non-uniform coordinated reset stimulation in ensembles of pulse-coupled neurons. *Frontiers in Computational Neuroscience*, *7*, 63.

Lysyansky, B., Popovych, O. V., & Tass, P. A. (2011). Desynchronizing anti-resonance effect of m: n ON-OFF coordinated reset stimulation. *Journal of Neural Engineering*, *8*(3), Article 036019.

Lysyansky, B., Popovych, O. V., & Tass, P. A. (2013). Optimal number of stimulation contacts for coordinated reset neuromodulation. *Frontiers in Neuroengineering*, *6*, 5.

Malkin, I. G. (1949). *The methods of Lyapunov and poincare in the theory of nonlinear oscillations*. Moscow-Leningrad: Gostekhizdat.

Martens, E. A., Barreto, E., Strogatz, S. H., Ott, E., So, P., & Antonsen, T. M. (2009). Exact results for the Kuramoto model with a bimodal frequency distribution. *Physical Review E*, *79*(2), Article 026204.

Matchen, T., & Moehlis, J. (2018). Phase model-based neuron stabilization into arbitrary clusters. *Journal of Computational Neuroscience*, *44*, 363–378.

Matchen, T. D., & Moehlis, J. (2021). Leveraging deep learning to control neural oscillators. *Biological Cybernetics*, *115*, 219–235.

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

Mayberg, H. S., Lozano, A. M., Voon, V., McNeely, H. E., Seminowicz, D., Hamani, C., et al. (2005). Deep brain stimulation for treatment-resistant depression. *Neuron*, *45*(5), 651–660.

McIntyre, C. C., Mori, S., Sherman, D. L., Thakor, N. V., & Vitek, J. L. (2004). Electric field and stimulating influence generated by deep brain stimulation of the subthalamic nucleus. *Clinical Neurophysiology*, *115*(3), 589–595.

Meidahl, A. C., Tinkhauser, G., Herz, D. M., Cagnan, H., Debarros, J., & Brown, P. (2017). Adaptive deep brain stimulation for movement disorders: the long road to clinical therapy. *Movement Disorders*, *32*(6), 810–819.

Merrill, D., Bikson, M., & Jefferys, J. (2005). Electrical stimulation of excitable tissue: design of efficacious and safe protocols. *Journal of Neuroscience Methods*, *141*, 171–198.

Mezić, I. (2019). Spectrum of the Koopman operator, spectral expansions in functional spaces, and state-space geometry. *Journal of Nonlinear Science*, 1–55.

Mitchell, I. M. (2007). *A toolbox of level set methods*: *Technical report UBC CS TR-2007-11*.

Mitchell, I. M. (2008). The flexible, extensible and efficient toolbox of level set methods. *Journal of Scientific Computing*, *35*(2), 300–329.

Mitchell, B. A., & Petzold, L. R. (2018). Control of neural systems at multiple scales using model-free, deep reinforcement learning. *Scientific Reports*, *8*, 1–12.

Moehlis, J., Shea-Brown, E., & Rabitz, H. (2006). Optimal inputs for phase models of spiking neurons. *ASME Journal of Computational and Nonlinear Dynamics*, *1*, 358–367.

Monga, B., Froyland, G., & Moehlis, J. (2018). Synchronizing and desynchronizing neural populations through phase distribution control. In *Proceedings of the 2018 American control conference* (pp. 2808–2813). Milwaukee.

Monga, B., & Moehlis, J. (2019a). Optimal phase control of biological oscillators using augmented phase reduction. *Biological Cybernetics*, *113*(1–2), 161–178.

Monga, B., & Moehlis, J. (2019b). Phase distribution control of a population of oscillators. *Physica D*, *398*, 115–129.

Monga, B., & Moehlis, J. (2020). Supervised learning algorithms for controlling underactuated dynamcial systems. *Physica D*, *412*, Article 132621.

Monga, B., Wilson, D., Matchen, T., & Moehlis, J. (2019). Phase reduction and phase-based optimal control for biological systems: A tutorial. *Biological Cybernetics*, *113*(1), 11–46. http://dx.doi.org/10.1007/s00422-018-0780-z.

Montbrió, E., Pazó, D., & Roxin, A. (2015). Macroscopic description for networks of spiking neurons. *Physical Review X*, *5*(2), Article 021028.

Nabi, A., Mirzadeh, M., Gibou, F., & Moehlis, J. (2013). Minimum energy desynchronizing control for coupled neurons. *The Journal of Computational Neurology*, *34*, 259–271.

Nagaraj, V., Lamperski, A., & Netoff, T. I. (2017). Seizure control in a computational model using a reinforcement learning stimulation paradigm. *International Journal of Neural Systems*, *27*, Article 1750012.

Netoff, T., Schwemmer, M. A., & Lewis, T. J. (2012). Experimentally estimating phase response curves of neurons: theoretical and practical issues. In *Phase response curves in neuroscience* (pp. 95–129). Springer.

Ota, K., Omori, T., Watanabe, S., Miyakawa, H., Okada, M., & Aonishi, T. (2011). Measurement of infinitesimal phase response curves from noisy real neurons. *Physical Review E*, *84*(4), Article 041902.

Ott, E., & Antonsen, T. M. (2008). Low dimensional behavior of large systems of globally coupled oscillators. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *18*(3), Article 037113.

Park, Y., & Ermentrout, B. (2016). Weakly coupled oscillators in a slowly changing world. *The Journal of Comparative Neurology*, *40*, 269–281.

Park, Y., & Wilson, D. D. (2021). High-order accuracy computation of coupling functions for strongly coupled oscillators. *SIAM Journal on Applied Dynamical Systems*, *20*(3), 1464–1484.

Pontryagin, L. S., Trirogoff, K. N., & Neustadt, L. W. (1962). *The mathematical theory of optimal processes*. Wiley New York.

Popovych, O. V., Hauptmann, C., & Tass, P. A. (2006). Control of neuronal synchrony by nonlinear delayed feedback. *Biological Cybernetics*, *95*(1).

Priori, A., Foffani, G., Rossi, L., & Marceglia, S. (2013). Adaptive deep brain stimulation (aDBS) controlled by local field potential oscillations. *Experimental Neurology*, *245*, 77–86.

Pyragas, K. (1992). Continuous control of chaos by self-controlling feedback. *Physics Letters. A*, *170*(6), 421–428.

Pyragas, K. (2006). Delayed feedback control of chaos. *Philosophical Transactions of the Royal Society of London A (Mathematical and Physical Sciences)*, *364*(1846), 2309–2334.

Pyragas, K., Fedaravičius, A. P., Pyragienė, T., & Tass, P. A. (2018). Optimal waveform for entrainment of a spiking neuron with minimum stimulating charge. *Physical Review E*, *98*(4), Article 042216.

Ramesh, M., & Narayanan, S. (1999). Chaos control by nonfeedback methods in the presence of noise. *Chaos, Solitons & Fractals*, *10*(9), 1473–1489.

Reppert, S. M., & Weaver, D. R. (2002). Coordination of circadian timing in mammals. *Nature*, *418*(6901), 935.

Rinzel, J. (1985). Excitation dynamics: insights from simplified membrane models. In *Federation proceedings. Vol. 44* (15), (pp. 2944–2946).

Ritt, J. T., & Ching, S. (2015). Neurocontrol: Methods, models and technologies for manipulating dynamics in the brain. In *Proceedings of the 2015 American control conference* (pp. 3765–3780). Chicago, IL.

Rosa, M., Arlotti, M., Ardolino, G., Cogiamanian, F., Marceglia, S., Fonzo, A. Di, et al. (2015). Adaptive deep brain stimulation in a freely moving parkinsonian patient. *Movement Disorders*, *30*(7), 1003–1005.

Rosenblum, M. G., & Pikovsky, A. S. (2004a). Controlling synchronization in an ensemble of globally coupled oscillators. *Physical Review Letters*, *92*, Article 114102.

Rosenblum, M., & Pikovsky, A. (2004b). Delayed feedback control of collective synchrony: An approach to suppression of pathological brain rhythms. *Physical Review E*, *70*(4), Article 041904.

Rosenblum, M., & Pikovsky, A. (2019). Numerical phase reduction beyond the first order approximation. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *29*(1), Article 011105.

Rosin, B., Slovik, M., Mitelman, R., Rivlin-Etzion, M., Haber, S. N., Israel, Z., et al. (2011). Closed-loop deep brain stimulation is superior in ameliorating parkinsonism. *Neuron*, *72*(2), 370–384.

Rubin, J. E., & Terman, D. (2004). High frequency stimulation of the subthalamic nucleus eliminates pathological thalamic rhythmicity in a computational model. *Journal of Computational Neuroscience*, *16*(3), 211–235.

Sanders, J. A., Verhulst, F., & Murdock, J. (2007). *Averaging methods in nonlinear dynamical systems* (2nd ed.). New York: Springer-Verlag.

Schiff, S. J. (2010). Towards model-based control of Parkinson's disease. *Philosophical Transactions of the Royal Society, Series A*, *368*, 2269–2308.

Schiff, S. J. (2012). *Neural control engineering*. Cambridge, Massachusetts: MIT Press.

Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, *656*, 5–28.

Schmidt, H., Avitabile, D., Montbrió, E., & Roxin, A. (2018). Network mechanisms underlying the role of oscillations in cognitive tasks. *PLoS Computational Biology*, *14*(9), Article e1006430.

Schöll, E., Hiller, G., Hövel, P., & Dahlem, M. A. (2009). Time-delayed feedback in neurosystems. *Philosophical Transactions of the Royal Society of London A (Mathematical and Physical Sciences)*, *367*(1891), 1079–1096.

Schrock, L. E., Mink, J. W., Woods, D. W., Porta, M., Servello, D., Visser-Vandewalle, V., et al. (2015). Tourette syndrome deep brain stimulation: a review and updated recommendations. *Movement Disorders*, *30*(4), 448–471.

Serkh, K., & Forger, D. B. (2014). Optimal schedules of light exposure for rapidly correcting circadian misalignment. *PLoS Computational Biology*, *10*(4), Article e1003523.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & Freitas, N. De (2015). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, *104*(1), 148–175.

Shirasaka, S., Kurebayashi, W., & Nakao, H. (2017). Phase-amplitude reduction of transient dynamics far from attractors for limit-cycling systems. *Chaos*, *27*, Article 023119.

Stieve, B., Richner, T., Krook-Magnuson, C., Netoff, T., & Krook-Magnuson, E. (2021). *Optimization of closed-loop electrical stimulation enables robust cerebellar-directed seizure control*. Cold Spring Harbor Laboratory, BioRxiv.

Strogatz, S. H. (2000). From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D*, *143*, 1–20.

Takata, S., Kato, Y., & Nakao, H. (2021). Fast optimal entrainment of limit-cycle oscillators by strong periodic inputs via phase-amplitude reduction and Floquet theory. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *31*(9), Article 093124.

Tanaka, H. A. (2014). Optimal entrainment with smooth, pulse, and square signals in weakly forced nonlinear oscillators. *Physica D: Nonlinear Phenomena*, *288*, 1–22.

Tass, P. A. (2003). A model of desynchronizing deep brain stimulation with a demand-controlled coordinated reset of neural subpopulations. *Biological Cybernetics*, *89*(2), 81–88.

Tass, P. (2007). *Phase resetting in medicine and biology: stochastic modelling and data analysis*. Berlin: Springer-Verlag.

Tu, J. H., Rowley, C. W., Luchtenberg, D. M., Brunton, S. L., & Kutz, J. N. (2014). On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, *1*(2), 314–421.

University of Michigan (2018). Entrain yourself. http://entrain.math.lsa.umich.edu. (Accessed 19 September 2018).

Volkmann, J., Herzog, J., Kopper, F., & Deuschl, G. (2002). Introduction to the programming of deep brain stimulators. *Movement Disorders*, *17*(3), S181–187.

Volkmann, J., Moro, E., & Pahwa, R. (2006). Basic algorithms for the programming of deep brain stimulation in Parkinson's disease. *Movement Disorders: Official Journal of the Movement Disorder Society*, *21*(S14), S284–S289.

Vreeswijk, C. Van, Abbott, L. F., & Ermentrout, G. B. (1994). When inhibition not excitation synchronizes neural firing. *Journal of Computational Neuroscience*, *1*(4), 313–321.

Wang, S., Musharoff, M. M., Canavier, C. C., & Gasparini, S. (2013). Hippocampal CA1 pyramidal neurons exhibit type 1 phase-response curves and type 1 excitability. *Journal of Neurophysiology*, *109*(11), 2757–2766.

Waterhouse, J., Reilly, T., Atkinson, G., & Edwards, B. (2007). Jet lag: trends and coping strategies. *The Lancet*, *369*(9567), 1117–1129.

Wedgwood, K. C. A., Lin, K. K., Thul, R., & Coombes, S. (2013). Phase-amplitude descriptions of neural oscillator models. *The Journal of Mathematical Neuroscience*, *3*(1), 1–22.

Weerasinghe, G., Duchet, B., Bick, C., & Bogacz, R. (2021). Optimal closed-loop deep brain stimulation using multiple independently controlled contacts. *PLoS Computational Biology*, *17*, e1009281.

Wiggins, S. (1994). *Normally hyperbolic invariant manifolds in dynamical systems*. New York: Springer.

Wilson, D. D. (2019). An optimal framework for nonfeedback stability control of chaos. *SIAM Journal on Applied Dynamical Systems*, *18*(4), 1982–1999.

Wilson, D. (2020a). A data-driven phase and isostable reduced modeling framework for oscillatory dynamical systems. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *30*(1), Article 013121.

Wilson, D. (2020b). Optimal open-loop desynchronization of neural oscillator populations. *Journal of Mathematical Biology*, *81*(1), 25–64.

Wilson, D. (2020c). Phase-amplitude reduction far beyond the weakly perturbed paradigm. *Physical Review E*, *101*(2), Article 022220.

Wilson, D. (2020d). Stabilization of weakly unstable fixed points as a common dynamical mechanism of high-frequency electrical stimulation. *Scientific Reports*, *10*(1), 1–21.

Wilson, D. (2021). Optimal control of oscillation timing and entrainment using large magnitude inputs: an adaptive phase-amplitude-coordinate-based approach. *SIAM Journal on Applied Dynamical Systems*, *20*(4), 1814–1843.

Wilson, D. (2022). An adaptive phase-amplitude reduction framework without $O(\epsilon)$ constraints on inputs. *SIAM Journal on Applied Dynamical Systems*, *21*(1), 204–230.

Wilson, C. J., Beverlin, B., & Netoff, T. (2011). Chaotic desynchronization as the therapeutic mechanism of deep brain stimulation. *Frontiers in Systems Neuroscience*, *5*, 50.

Wilson, D., & Ermentrout, B. (2018a). Greater accuracy and broadened applicability of phase reduction using isostable coordinates. *Journal of Mathematical Biology*, *76*(1–2), 37–66.

Wilson, D., & Ermentrout, B. (2018b). An operational definition of phase characterizes the transient response of perturbed limit cycle oscillators. *SIAM Journal on Applied Dynamical Systems*, *17*(4), 2516–2543.

Wilson, D., & Ermentrout, B. (2019a). Augmented phase reduction of (not so) weakly perturbed coupled oscillators. *SIAM Review*, *61*(2), 277–315.
Wilson, D., & Ermentrout, B. (2019b). Phase models beyond weak coupling. *Physical Review Letters*, *123*(16), Article 164101.
Wilson, D., Faramarzi, S., Moehlis, J., Tinsley, M. R., & Showalter, K. (2018). Synchronization of heterogeneous oscillator populations in response to weak and strong coupling. *Chaos. An Interdisciplinary Journal of Nonlinear Science*, *28*(12), Article 123114.
Wilson, D., Holt, A. B., Netoff, T. I., & Moehlis, J. (2015). Optimal entrainment of heterogeneous noisy neurons. *Frontiers in Neuroscience*, *9*, 192.
Wilson, D., & Moehlis, J. (2014a). An energy-optimal approach for entrainment of uncertain circadian oscillators. *Biophysical Journal*, *107*, 1744–1755.
Wilson, D., & Moehlis, J. (2014b). Locally optimal extracellular stimulation for chaotic desynchronization of neural populations. *Journal of Computational Neuroscience*, *37*, 243–257.
Wilson, D., & Moehlis, J. (2014c). Optimal chaotic desynchronization for neural populations. *SIAM Journal on Applied Dynamical Systems*, *13*, 276–305.
Wilson, D., & Moehlis, J. (2015). Clustered desynchronization from high-frequency deep brain stimulation. *PLoS Computational Biology*, *11*(12), Article e1004673.
Wilson, D., & Moehlis, J. (2016a). Isostable reduction of periodic orbits. *Physical Review E*, *94*, Article 052213.
Wilson, D., & Moehlis, J. (2016b). Isostable reduction with applications to time-dependent partial differential equations. *Physical Review E*, *94*, 012211.
Winfree, A. T. (1967). Biological rhythms and the behavior of populations of coupled oscillators. *Journal of Theoretical Biology*, *16*, 14–42.
Winfree, A. (1974). Patterns of phase compromise in biological cycles. *Journal of Mathematical Biology*, *1*, 73–95.
Winfree, A. (2001). *The geometry of biological time* (2nd ed.). New York: Springer.
Wright, K. P., Hughes, R. J., Kronauer, R. E., Dijk, D. J., & Czeisler, C. A. (2001). Intrinsic near-24-h pacemaker period determines limits of circadian entrainment to a weak synchronizer in humans. *Proceedings of the National Academy of Sciences*, *98*(24), 14027–14032.
Yu, Y.-C., Narayanan, V., Ching, S., & Li, J.-S. (2020). Learning to control neurons using aggregated measurements. In *Proceedings of the 2020 American control conference* (pp. 4028–4033). Denver, CO.
Zlotnik, A., Chen, Y., Kiss, I. Z., Tanaka, H.-A., & Li, J.-S. (2013). Optimal waveform for fast entrainment of weakly forced nonlinear oscillators. *Physical Review Letters*, *111*, Article 024102.
Zlotnik, A., & Li, J. S. (2012). Optimal entrainment of neural oscillator ensembles. *Journal of Neural Engineering*, *9*(4), Article 046015.