Biological Cybernetics
https://doi.org/10.1007/s00422-024-00997-w

Biological Cybernetics

ORIGINAL ARTICLE

Phase response curves and the role of coordinates

Simon Wilshin[1] · **Matthew D. Kvalheim**[2] · **Shai Revzen**[3]

Received: 18 April 2024 / Accepted: 22 August 2024

Abstract

The "infinitesimal phase response curve" (PRC) is a common tool used to analyze phase resetting in the natural sciences in general and neuroscience in particular. We make the observation that the PRC with respect to a coordinate v actually depends on the choice of other coordinates. As a consequence, a complete delay embedding reconstruction of the dynamics using v which would allow phase to be computed still does not allow the v PRC to be computed. We give a coordinate-free definition of the PRC making this observation obvious. This leads to an experimental protocol: first collect an appropriate ensemble of measurements by intermittently controlling neuron voltage. Then, for any suitable current carrier dynamic postulated, we show how the ensemble can be used to compute the voltage PRC with that current carrier. The approach extends to many oscillators measured and controlled through a subset of their coordinates.

Keywords Phase response · Dynamical systems · Neural oscillator · Fitzhugh–Nagumo · Delay coordinates

Contents

Communicated by Peter Thomas.

✉ Shai Revzen
shrevzen@umich.edu

Simon Wilshin
swilshin@rvc.ac.uk

Matthew D. Kvalheim
kvalheim@umbc.edu

[1] Royal Veterinary College, London, UK

[2] University of Maryland, Baltimore County, MD, USA

[3] University of Michigan, Ann Arbor, MI, USA

1 Motivation

Taking voltage measurements of neurons and using time delay embedding to reconstruct the neuronal dynamics can lead to erroneous estimates of phase response curves (PRCs), even the PRC of the voltage which was being directly measured. Here we discuss how to correctly recover the PRC from any choice of state measurements, show that reconstruction of the dynamics from, e.g. voltage delays, does not fully specify the PRC, and show how the PRC can be reconstructed once missing information is added.

Many physical systems are oscillators, i.e. they exhibit stable, long-term repeating cycles which persist even when the system is perturbed. Spike train generating neurons fall into this class, and it is common for neuro-scientists to attempt to study these neurons using voltage measurements or by clamping the voltage to a desired value and then releasing the clamp as a way of resetting initial conditions and exam-

Published online: 30 October 2024

ining the recovery. Mathematical models of coupled neurons often use the *infinitesimal phase response curve* which represents the ratio of infinitesimal phase change to infinitesimal external perturbation. For example, the voltage-PRC shows how much neuronal timing (phase) will change for a small voltage perturbation applied as a function of the phase in the cycle at which this perturbation is applied.

These PRCs are used extensively in neuroscience models (Izhikevich 2007), dating back at least to the 1970s and 1980s, which saw significant developments applying the theory of weakly coupled oscillators to biological problems, e.g. Ermentrout (1986), Guckenheimer (1975) and Winfree (1980). A PRC can be used to create a Fokker–Planck equation for the distribution of observed phases of a population after a perturbation; this relationship can sometimes be inverted (Wilson and Moehlis 2015) to obtain the PRC. It is also common to study the neuronal phase response to an incoming spike at various phases of the cycle. This non-infinitesimal phase response cannot, in general, be computed from the (infinitesimal) PRC because additional non-linearities come into play. However, if the "isostable reduction[1]" is known, one can replace the estimation of asymptotic phase response curves with the measurement of "operational phase" response and obtain the similar coupling behaviors (Wilson and Ermentrout 2018).

In this paper we are concerned with the process of estimating a PRC, e.g. the voltage PRC, using a state-space model constructed via delay embeddings, e.g. from the aforementioned voltage-only measurements and clamping experiments. While it is typically possible to reconstruct the dynamics of neurons from voltage and its delays, the PRC obtained from taking the (partial) derivative of phase with respect to voltage does not yield the correct result for the voltage-PRC. This is true even for seemingly trivial changes of co-ordinate, consider for example the first row of Fig. 1, where a simple linear change of co-ordinates for a neural oscillator radically changes the resoluting PRC. This problem extends to the PRC with respect to arbitrary perturbation types. Somewhat counter-intuitively, even an ensemble of voltage measurements that fully characterizes the dynamics of the neuron is not sufficient for computing so much as the voltage-PRC. Despite the fact that the input variable voltage, the output variable phase, and the dynamics governing both in time are fully known, the voltage-PRC is not defined without knowing what states to hold constant while voltage is perturbed. By postulating the dynamics of hidden state variables, e.g. gating or current variables, the problem can be resolved and the appropriate coordinate changes can be correctly computed and produce the voltage-PRC, or any other PRC needed—but the same ensemble of voltage experiments can be compatible with multiple such postulates leading to different results.

Our contributions are: (1) We point out that the straightforward approach of reconstructing the dynamics from voltage measurements and differentiating phase with respect to voltage at various phases will not produce the correct PRC, because it is coordinate dependent. (2) We show how an ensemble of (noisy, experiment-like) voltage measurements can be combined with a model of the ionic currents or gating variables to produce the correct voltage PRC. (3) We discuss what the PRC is in a physical, coordinate-free sense, for oscillators of arbitrary dimension, using tools of differential geometry.

After the general background in Sect. 2, this paper can be read in two ways. A neuroscience practitioner might read Sect. 3 and would almost certainly wish to bypass the mathematical technicalities of Sects. 4 and 4.1 for the practical methods described in Sect. 5 and the neuroscience motivated example of Sect. 6. For the mathematician interested in oscillator theory, the coordinate invariant formulation in Sects. 4 and 8 form the beating heart of the paper, the examples of Sects. 4.1 and 6 serve as numerical illustrations, and Figs. 5 and 6 serve as cautionary tales regarding the numerical challenges that could be encountered.

2 Oscillators in the physical sciences

This section contains a (somewhat formal) mathematical definition of an "oscillator" model suitable for modeling physical systems. Consider a system described by an ordinary differential equation

$$\dot{x} := dx/dt = f(x) \qquad x(t) \in B \subset \mathbb{R}^n,\ f \in \mathcal{C}^1,\ B \text{ open.} \tag{1}$$

Intuitively, we wish to focus on the part of the dynamics of our system—the stability basin B—that converges to a specific periodic solution. For our purposes here we add some conditions on the convergence rates that ensure the existence of phase. We therefore consider an "*oscillator*" such that: (1) there exists a trajectory $\gamma : \mathbb{R} \to \Gamma \subset B$, and (2) an associated "*(minimal) period*" $T > 0$ such that for any $t \in \mathbb{R}$, $\gamma(t + T) = \gamma(t)$, and (3) all initial conditions in B converge to Γ (the image of γ) at an exponential rate,[2] (4) Γ is Lyapunov stable.

While the phase of oscillators is conventionally represented as a real number in either $[0, 1)$ or $[0, 2\pi)$, these

[1] In its general form, this is the same as "data-driven floquet analysis" representation of Revzen (2009) and Revzen and Kvalheim (2015), originally due to Floquet (1883).

[2] Exponential convergence is sufficient for the existence of phase in oscillators, and necessary for the structural stability of the equations. Without structural stability the equations would be challenging to use in physical modeling.

representations suffer from a discontinuity when cycles wrap, breaking the inherent symmetry whereby no particular phase is "special". Since in this paper we are concerned with differentials of phase and such discontinuities do violence to both the underlying geometry and the very differentials we seek, we adopt a less common representation. We represent a scalar phase using a complex valued "phasor", i.e. a point on the unit circle S^1 in the complex plane $\mathbb{C}$. A conventional real valued phase $\theta \in [0, 2\pi]$ is represented as $e^{i\theta} \in \mathbb{C}$.

From its periodicity, it is trivial to see that Γ can be put in correspondence $\varphi : \Gamma \to S^1 \subset \mathbb{C}$ with S^1 the complex unit circle, such that $\frac{d}{dt}\varphi(\gamma(t)) = (2\pi i/T)\varphi(\gamma(t))$ (where $i := \sqrt{-1}$), i.e. the "*(complex valued asymptotic) phase*" φ circles at a constant rate, and is uniquely defined up to one arbitrary choice of "*gauge*"—the value of phase for one point on the cycle, e.g. $\varphi(\gamma(0))$, can be arbitrarily chosen thereby fixing all the other values of phase. Often, instead of taking the phase as a complex valued phasor, its argument in radians is used instead to give a "*(real valued asymptotic) phase*" which has $\frac{d}{dt}\arg(\varphi(\gamma(t))) = 2\pi/T$ wherever $\arg(\varphi(\gamma(t)))$ is differentiable.

Definition 1 The classical work of Guckenheimer (1975) showed that φ can be extended (C^1) smoothly and uniquely to the entirety of the interior of B, defining the "*asymptotic phase map*". By abuse of notation we denote this map by $\varphi : B \to S^1$, and note that it satisfies $\frac{d}{dt}\varphi(x(t)) = (2\pi i/T)\varphi(x(t))$ for all trajectories $\dot{x} = f(x)$ of the system.

3 Coordinate-dependence of the standard PRC definition

If $(x_1, x_2, \ldots, x_n)$ are coordinates for $\mathbb{R}^n$, then the standard definition of the "*(infinitesimal) phase response curve (PRC)*" with respect to, say, x_1 is a map $\rho_1 : \Gamma \to \mathbb{R}$ defined by the partial derivative

$$\rho_1 := \frac{1}{i\varphi}\frac{\partial\varphi}{\partial x_1} \tag{2}$$

Equivalently, $\rho_1(x) = \frac{\partial}{\partial x_1}\arg(\varphi(x))$ wherever $\arg(\varphi(\cdot))$ is differentiable. See, e.g., (2010, p. 176).

Suppose now that the last $(n-1)$ coordinates are modified to produce a coordinate system of the form $(x_1, y_2, \ldots, y_n)$ with the same phase $\tilde{\varphi}(x_1, y_2, \ldots, y_n) = \varphi(x_1, \ldots, x_n)$ with respect to these new coordinates. We can again define a PRC with respect to x_1, now denoted $\tilde{\rho}_1 : \Gamma \to \mathbb{R}$, via

$$\tilde{\rho}_1 := \frac{1}{i\tilde{\varphi}}\frac{\partial\tilde{\varphi}}{\partial x_1} \tag{3}$$

In general, $\rho_1 \neq \tilde{\rho}_1$, even though both are similarly defined by limits of the form:

$$\frac{\partial\varphi}{\partial x_k}(x) := \lim_{h\to 0}\frac{\varphi(x + \mathbf{e}_k h) - \varphi(x)}{h}.$$

In the next section we introduce a coordinate-free definition of the PRC which makes this coordinate dependence obvious. This coordinate dependence can be seen in Fig. 1, which shows two examples illustrating how the PRC changes with the choice of projection in which the dynamics are presented. Not only that, but even the PRC with respect to a variable shared between the projections may change.

As a preliminary explanation, we quote the following chastening from V. I. Arnold (1989, p. 258, foot. 81):

> "It is important to note that the quantity $\frac{\partial u}{\partial x}$ on the x, y-plane depends not only on the function which is taken for x, but also on the function y: in new variables (x, z) the value of $\frac{\partial u}{\partial x}$ will be different. One should write
>
> $$\left.\frac{\partial u}{\partial x}\right|_{y=\text{const.}} \qquad \left.\frac{\partial u}{\partial x}\right|_{z=\text{const.}}."$$

Taking this to heart, we should have written:

$$\rho_1 := \frac{1}{i\varphi}\left.\frac{\partial\varphi}{\partial x_1}\right|_{x_2,\ldots,x_n=\text{const.}} \qquad \tilde{\rho}_1 := \frac{1}{i\tilde{\varphi}}\left.\frac{\partial\tilde{\varphi}}{\partial x_1}\right|_{y_2,\ldots,y_n=\text{const.}}, \tag{4}$$

making the difference obvious. The coordinate-free view of the PRC in Sect. 4 yields another way to make the difference obvious.

Perhaps the most familiar examples of this issue appear in undergraduate thermodynamics. Consider the work done by a periodically compressing piston containing an ideal gas. For a periodic motion like $V(t) = \sin(t)$, we could get complete equations of motion by only looking at the volume of the piston over time and expressing dV/dt as a function of $V(t)$ and $V(t-\delta)$. Despite knowing these complete equations of motion, we cannot know the work of the piston. The work of the piston is the integral of the increments of work with respect to volume, but the increment of work dW performed at a given volume V cannot be computed from the increment of volume change dV without additional constraints on the thermodynamic state variables. Classical constraints show qualitatively different work outcomes: in isobaric compression, with pressure constant, $dW = -PdV$ and the work increment is only a function of volume increment; in isothermal compression, by substitution of $P = nRT/V$ from the ideal gas law, we get $dW = -nRTdV/V$, and the work increment is inversely proportional to volume. Just as the increment of work dW associated with a change of volume depends on the choice of coordinates held constant,

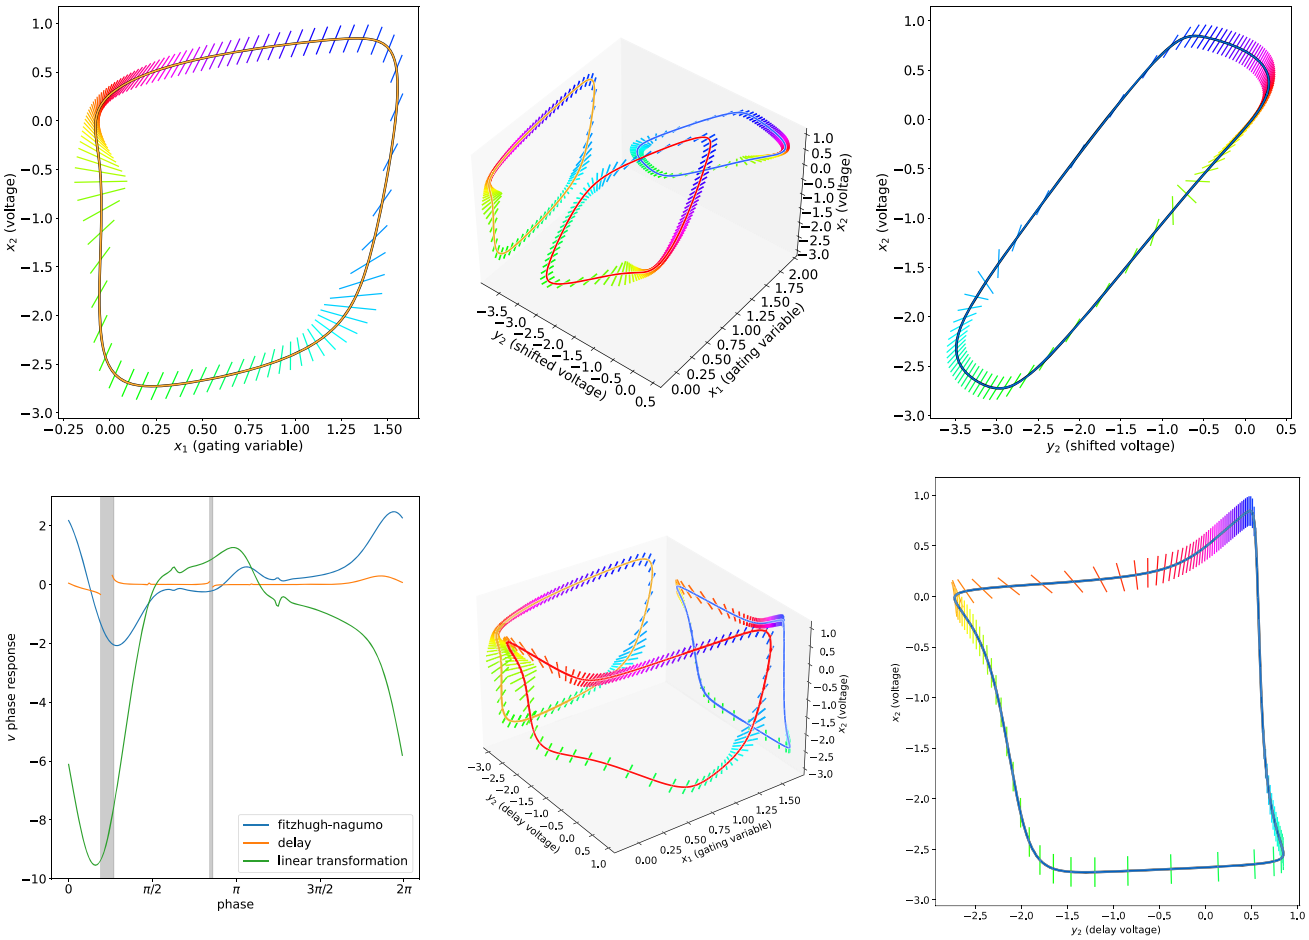

Fig. 1 Two examples of how the PRC depends on the choice of coordinates. Consider the Fitzhugh–Nagumo oscillator (top left, equations of motion given in equation (17)) presented in terms of x_1 (voltage) and x_2 (gating variable), and consider a y_2 replacing x_2. Because y_2 is a function of x_1 and x_2, the oscillator's state-space occupies a two dimensional ribbon in the 3-dimensional (x_1, x_2, y_2) space (center column, rainbow colors). The isochrons lie in this ribbon and intersect the limit cycle at an angle, the slope of which is the PRC. Consequently, the same dynamics can produce very different PRC-s even with respect to the x_1 shared by both (x_1, x_2) and (x_1, y_2) representations. As a simple case, take y_2 linear in x_1 and x_2 (center and right, top), with $y_2 = x_1 - \left(2x_{1,max}/x_{2,max}\right) x_2$, where the "max" subscript indicates the maximum value on the limit cycle). A more complicated smooth y_2 dependence can lead to even more complicated results (center and right, bottom); here we used $y_2(t) := x_1(t - \tau)$ a delayed x_1. The x_1 (voltage) PRC-s obtained in these three coordinate systems are quite different (bottom left). Note that the use of delay coordinates has great advantages from an experimental perspective, because a single time series measurement of voltage provides both x_1 and y_2. Takens' theorem (Takens 1980; Sauer et al. 1991) guarantees that for deterministic dynamical systems, almost every choice of delay coordinates will allow the dynamics to be reconstructed. In practice, finding a set of delays producing coordinates with Jacobians that have a moderate condition number can be challenging and due to this poor condition number portions of the PRC have been omitted (this is indicated by dark vertical bars)

the increment of phase associated with a voltage change, the voltage-PRC, depends on the choice of coordinates held constant.

4 A coordinate-free view of the PRC

In this section we present a geometric notion of the PRC and show how it can be represented in a coordinate-free way over a broad class of state-spaces. Recall from Definition 1 that the complex phase φ takes its value on the unit circle in the complex plane. For any path through the state space $p : [0, L] \to B$, one can hope to formulate the change of phase along the path as an integral of infinitesimal changes of phase along that path, accumulating the total angle through which phase has rotated. Viewed geometrically, the phase φ partitions ("foliates" Guckenheimer and Holmes 1983) the stability basin into smooth surfaces of constant phase called "*isochrons*". The change of phase experienced along a path segment $\dot{p}(t)dt$ is a measure of the number of isochrons crossed in that segment, and in Euclidean spaces it can be estimated as $(-i/\varphi(p(t))\nabla\varphi(p(t))\cdot\dot{p}(t)dt$, which is a vector normal to the isochrons and scaled inversely proportional to the density of the isochrons at that point. Expressing this with

the gradient ∇ requires the underlying space to have a notion of length in the form of a Riemannian metric, since that is needed for defining a gradient, but this Riemannian structure is not required for having a PRC.

The mathematical object that represents the infinitesimal change in phase available at a point $x \in B$ in a metric-independent way is the "*differential 1-form*" (Rudin 1976, Ch.10) which we name the temporal 1-form, $d\tau$, the d indicating that it is a differential 1-form. However it is not the (exterior) derivative of any real-valued function, i.e. it is not obtained by applying the operator d to a real-valued function τ. Like all differential 1-forms, $d\tau$ is a section of a co-tangent bundle, here denoted T^*B.

The tooling of differential forms allows us to define and understand $d\tau$, as we describe below. One important feature of differential forms is that they admit a "*pullback*" operation: if a differential form is defined on the co-domain Y of a smooth map $\psi : X \to Y$, it defines a differential form on the domain X of that map which, informally speaking, "does the same thing". Whatever velocities, such as the rate of change of phase, that the form can measure on paths in the co-domain Y, the pullback form can measure on paths in the domain X mapping to those paths in Y. It does so by mapping the point of evaluation from X to Y using ψ, and the velocity at that point using $\mathsf{D}\psi$, and then applying the original form to the result.

The form $d\tau$ can most easily be seen as pullback of a form measuring the rotation velocity on the unit circle, which is the image of φ. Any change of the coordinates for B does not change the velocity traveled in the image of φ, and thus $d\tau$ is a coordinate invariant mathematical object. The details of this formulation using the notation of differential forms, and the equations governing its transformation between coordinate frames, have been relegated to Sect. 8.

4.1 An extended example

We now turn to an illustrative example of how PRCs change coordinates and in particular with the use of delay coordinates as would be the case in reconstructing neuronal dynamics from voltage measurements and their delays. Perforce, this requires an artificially simple oscillator, whose dynamics are fully integrable, so that we can compute all the relevant expressions in closed form.

The more mathematically inclined reader should view this example in light of Example 1 in Sect. 8 which describes a more general case. Consider a Hopf-like oscillator in polar coordinates on $\mathbb{R}^2\backslash\{\mathbf{0}\}$:

$$\dot{\theta} = w \qquad \dot{r} = 1 - r \tag{5}$$

where dots indicate differentiation with respect to time t. This has the general solution:

$$\theta = wt + \theta_0 \qquad r = 1 + \rho_0 \exp(-t), \tag{6}$$

where $\rho_0 := r(0) - 1$. Note that $\{r = 1\}$ is the image of an exponentially stable periodic orbit for this system; we will also refer to this periodic orbit as a "*limit cycle*".

Let us switch to Cartesian coordinates and imagine that our y coordinate is some variable we can observe (for example, a voltage of a neuron), while the x coordinate is some variable we do not observe (for example, the associated current). We rewrite the solutions as:

$$x = [1 + \rho_0 \exp(-t)] \cos(wt + \theta_0) \tag{7}$$

$$y = [1 + \rho_0 \exp(-t)] \sin(wt + \theta_0)\,. \tag{8}$$

Recalling that $r = \sqrt{x^2 + y^2}$, and noting that the phase of this oscillator can be taken to be θ, we can compute the phase 1-form in this coordinate system as:

$$d\theta(x, y) = \left(-y/r^2, x/r^2\right), \tag{9}$$

where the right side should be viewed as a row vector or co-vector.[3] Using for simplicity the terminology of Sect. 3, the restrictions to the unit circle of the components of this 1-form are the PRCs with respect to x and y in the coordinate system (x, y). We also know the isochrons explicitly, since the asymptotic phase map is given by θ and its level sets are radial lines. This gives rise to the coordinate-independent observation that a velocity v is tangent to some isochron if, and only if, it satisfies $\langle d\theta, v\rangle = 0$. In Cartesian coordinates this is satisfied for $v = (v_0, v_1)$ if $yv_0 = xv_1$, that is if the velocity is radial.

We will now consider delay coordinates (v, u) and how we obtain them from (x, y). The coordinate $u(t)$ is the same coordinate which we observe, $y(t)$. For the coordinate $v(t)$ we take the value of the y coordinate of our system at some lag time $\delta > 0$, i.e., $y(t - \delta)$. This gives

$$u(t) := y(t) = [1 + \rho_0 \exp(-t)] \sin(wt + \theta_0) \tag{10}$$

$$v(t) := y(t - \delta) = [1 + \rho_0 \exp(-t + \delta)] \sin(wt - w\delta + \theta_0) \tag{11}$$

We would like to know how the PRCs with respect to v and u in the coordinates (v, u) are related to those we previously calculated with respect to the (x, y) coordinates. In the present example we have an explicit formula for the coordinate transformation, so we can eliminate time from the above relations and carry out the remaining computations explicitly.

[3] Despite the fact that the closed 1-form defined by the right side of (9) is not exact on $\mathbb{R}^2\backslash\{\mathbf{0}\}$, "$d\theta$" is still common notation for this 1-form. See for example Spivak (1971, p. 93).

The u coordinate change is trivial since, by construction, $u = y$. For v, note the following identities and definitions, all written in terms of x, y, and $r := \sqrt{x^2 + y^2}$: $\exp(-t) = (r - 1)/\rho_0$, $y/r = \sin(wt + \theta_0)$, $x/r = \cos(wt + \theta_0)$, $C := \cos(w\delta)$, and $S := \sin(w\delta)$.

We thus have:

$$\begin{aligned} v &= \left[1 + \rho_0 e^{\delta} \cdot (r-1)/\rho_0\right] \sin(wt - w\delta + \theta_0) \\ &= \left[1 + (r-1)e^{\delta}\right] (\sin(wt+\theta_0)C - \cos(wt+\theta_0)S) \\ &= \left[(1-e^{\delta}) + re^{\delta}\right] \frac{1}{r}(yC - xS) \\ &= \left[(1-e^{\delta})/r + e^{\delta}\right] (yC - xS) \end{aligned} \tag{12}$$

From this transformation we can compute the Jacobian J of the transformation $T_\delta : (x, y) \mapsto (v, u)$ The components of the Jacobian are (using the notation "$|_x$" as shorthand for "$|_{x=\text{const.}}$" and similarly for "$|_y$"):

$$\begin{aligned} \left.\frac{\partial u}{\partial x}\right|_y &= 0; \\ \left.\frac{\partial v}{\partial x}\right|_y &= \frac{1}{r^3}\left(-Sr^2\left(e^{\delta}r - e^{\delta} + 1\right) + x\left(e^{\delta} - 1\right)(Cy - Sx)\right); \\ \left.\frac{\partial u}{\partial y}\right|_x &= 1; \\ \left.\frac{\partial v}{\partial y}\right|_x &= \frac{1}{r^3}\left(Cr^2\left(e^{\delta}r - e^{\delta} + 1\right) + y\left(e^{\delta} - 1\right)(Cy - Sx)\right). \end{aligned}$$

For certain values of the parameters w and δ, the quantity $\left.\frac{\partial v}{\partial x}\right|_y$ is nowhere vanishing on the image of the limit cycle.[4] For such parameter values, this leads to the invertibility of the Jacobian J at each point on the limit cycle, with expressions for J and J^{-1} as follows:

$$\begin{aligned} J(x, y) &= \begin{pmatrix} \left.\frac{\partial v}{\partial x}\right|_y & \left.\frac{\partial v}{\partial y}\right|_x \\ 0 & 1 \end{pmatrix}, \\ J^{-1}(x, y) &= \begin{pmatrix} \left(\left.\frac{\partial v}{\partial x}\right|_y\right)^{-1} & -\left(\left.\frac{\partial v}{\partial y}\right|_x\right)\left(\left.\frac{\partial v}{\partial x}\right|_y\right)^{-1} \\ 0 & 1 \end{pmatrix}. \end{aligned} \tag{13}$$

[4] Setting $r = 1$ in $\left.\frac{\partial v}{\partial x}\right|_y$ and viewing the resulting expression as a quadratic function of x with parameter y, the discriminant Δ of this quadratic is $\Delta = (e^{\delta} - 1)^2 C^2 y^2 - 4(e^{\delta} - 1)S^2$. Hence $\Delta < 0$ with $\delta > 0$ if and only if $(e^{\delta} - 1)C^2 y^2 < 4S^2$. This implies that, with $\delta > 0$, $\Delta < 0$ for all $(x, y) \in \{r = 1\}$ if and only if $(e^{\delta} - 1)C^2 < 4S^2$. Since a quadratic has a real root if and only if its discriminant is nonnegative, the latter inequality is necessary and sufficient for the quantity $\left.\frac{\partial v}{\partial x}\right|_y$ to be nowhere vanishing on $\{r = 1\}$. If it so happens that $w\delta \notin \frac{\pi}{2} + \pi\mathbb{Z}$, then this necessary and sufficient condition reads

$$e^{\delta} - 1 < 4\tan^2(w\delta).$$

Since $\tan^2(w\delta) = w^2\delta^2 + o(\delta^5)$ as $\delta \to 0$ while $e^{\delta} - 1 = \delta + o(\delta)$, we see that, for this condition to hold, w must be very large if δ is very small.

If additionally $\delta \notin \frac{2\pi}{w}\mathbb{Z}$ so that the restriction of T_δ to $\{r = 1\}$ is invertible, it can be shown that the nonlinear map T_δ is smoothly invertible on some open neighborhood of $\{r = 1\}$ (Guillemin and Pollack 2010, p. 19, Ex. 10). Thus, for such parameter values, we may use the transformation law for 1-forms to transform $d\theta(x, y)$ on a neighborhood of $\{r = 1\}$ according to

$$d\theta_\delta(v, u) := d\theta(T_\delta^{-1}(v, u))J^{-1}(T_\delta^{-1}(v, u)) \tag{14}$$

Since $u = y$, one might assume mistakenly that the change in phase with respect to an infinitesimal perturbation in y in the (x, y) coordinates is the same as the change in phase with respect to an infinitesimal perturbation in u in the (u, v) coordinates. However, the former is x/r^2, whereas the latter is $x/r^2 + y/r^2 \left.\frac{\partial v}{\partial y}\right|_x \left(\left.\frac{\partial v}{\partial x}\right|_y\right)^{-1}$.

Taking one specific example, choose $\delta := \ln 2$ and $w = \pi/(4 \ln 2)$, giving $C = S = 1/\sqrt{2}$, and examine J on the limit cycle $\{r = 1\}$, setting $x := \cos(\phi)$ and $y := \sin(\phi)$. After some algebra, we obtain:

$$\left.\frac{\partial v}{\partial x}\right|_y = -\frac{1}{2}\cos\left(2\phi + \frac{\pi}{4}\right) - \frac{3\sqrt{2}}{4}, \tag{15}$$

$$\left.\frac{\partial v}{\partial y}\right|_x = -\frac{1}{2}\sin\left(2\phi + \frac{\pi}{4}\right) + \frac{3\sqrt{2}}{4}. \tag{16}$$

The resulting difference in PRCs between y and u is shown in Fig. 2.

5 PRC recovery with delay coordinates

The difficulty of the process in the example of the previous section, supported by the detailed analysis in Sect. 8, at equation (24), suggests that the prospect of computing a desired PRC from measurements of a scalar-valued observable such as voltage is grim, since equation (24) involves the (usually unknown) derivatives of the flow ($\mathrm{D}\Phi^{-\delta_k}$ in Sect. 8).

However, in this section we show that certain PRCs *can* be recovered by measurements of a scalar-valued observable under certain additional assumptions. This allows us to provide a potential solution to the problem of recovering a PRC in an experimental system, even if the system provides limited options for observing and controlling its state.

5.1 Assumptions and goal

We continue to make the same assumptions as in Sects. 2 and 4, but need additional requirements that allow us to consider external control of some aspects of the state. We begin by assuming that $B \subseteq M \subseteq X \times Y$, i.e. where M points have a canonical representation as pairs from X and Y, where

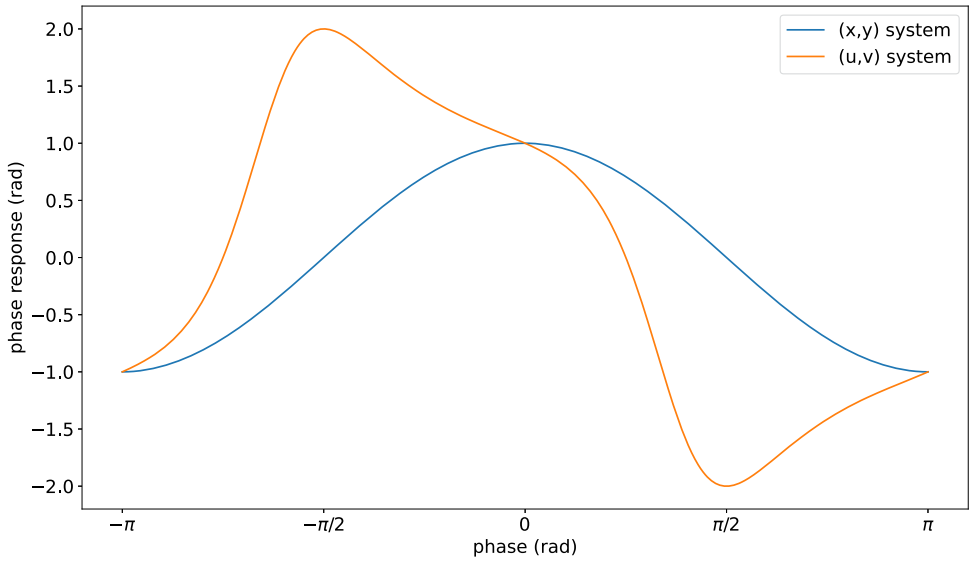

Fig. 2 Comparison of the phase response with respect to an infinitesimal perturbation of the u-coordinate in (v, u), and the y-coordinate in (x, y). Note that we have chosen $u = y$, and we are perturbing the same oscillator, with parameters $\delta := \ln 2$ and $w = \pi/(4 \ln 2)$

we will both observe and control the Y part. We write the vector field f on M as $f = (f_x, f_y)$ so that an integral curve $(x(t), y(t))$ of f satisfies the ODE $\dot{x} = f_x(x, y)$; $\dot{y} = f_y(x, y)$. As is traditional in control theory, we think of x as being a "hidden state" whereas y is "directly observable", i.e., available for us to measure at any instant.

Because our goal is to examine cases where x and y may fall outside the range of pairings allowed by this un-forced ODE, we consider a space of admissible controls $\mathcal{Y} \subset \mathbb{R}_+^Y$ a space of functions from $t \in \mathbb{R}$, $t \geq 0$ to Y such that for each $y_i \in \mathcal{Y}$, there is a bounded interval $[t_i^0, t_i^1]$ such that outside this interval y_i is the Y component of the system $\dot{x} = f_x(x, y_i)$; $\dot{y}_i = f_y(x, y_i)$, and within the interval $\dot{x} = f_x(x, y_i)$ but the second (Y) equation need not hold. We require M to be an open set containing all pairs of (x, y_i) arising from $y_i \in \mathcal{Y}$ (and therefore also B).

We take $\gamma : \mathbb{R} \to B$ be a specific T-periodic solution of the ODE with image $\gamma(\mathbb{R}) = \Gamma$, and assume γ is (not merely orbitally) exponentially stable. From T-periodicity, $\gamma(T) = \gamma(0)$ and $\gamma|_{[0,T)}$ is injective. We use the notation $\gamma_x(t) \in X, \gamma_y(t) \in Y$ for the components of γ so that $\gamma(t) = (\gamma_x(t), \gamma_y(t))$.

Considering the dynamics of $x(t)$, we refer to the control system $\dot{x} = f_x(x(t), u)$ as the $\mathbf{X}(u)$ system. It is a control system with control input u when u is unspecified. If u is a function of time $u(t)$, $\mathbf{X}(u)$ becomes a non-autonomous differential equation. One such non-autonomous system on X is the system $\mathbf{X}(\gamma_y)$ defined by taking $u(t) := \gamma_y(t)$ and leading to the equation $\dot{x} = f_x(x(t), \gamma_y(t))$.

Goal *Find a method to recover the ∂_y-PRC from y measurements, where $y \in Y = \mathbb{R}$.*

We believe the same approach could be extended to Y of higher dimension by defining "vector-valued PRCs".

Assumption 1 $f_x(\cdot, \cdot)$ is known. We do not assume knowledge of $f_y(\cdot, \cdot)$.

Here we assume that we can postulate the dynamics of the part of the state which we can neither directly observe nor directly manipulate. We assume no knowledge of the part of the state we can measure and manipulate (see below).

Assumption 2 There exist known $0 < \delta_1 < \delta_2 < \cdots < \delta_m$ such that the delay map $R : M \to \mathbb{R}^m$ given by $R : (x(t), y(t)) \mapsto (y(t), y(t-\delta_1), \ldots, y(t-\delta_m)) \in \mathbb{R}^m$ is a continuously differentiable map with a differentiable inverse on its image for all trajectories $(x(t), y(t))$ evolving under $\dot{x} = f_x(x, y)$; $\dot{y} = f_y(x, y)$.

This assumes that we have delay coordinates of y which could allow us to reconstruct the dynamics of the observable y in a data-driven manner. The differentiability of the map and its inverse ensure that the dynamics can be written in the transformed coordinates. Takens' theorem (Takens 1980; Sauer et al. 1991) tells us that such a collection $\{\delta_i\}_{i=0}^m$ of delays can usually be easily found. In fact, given some technical conditions are satisfied, almost all collections of delays will do.

Assumption 3 For any C^1 curve $u : [t_0, t_1] \to Y$ satisfying $u(t_0) = y(t_0)$, we can enforce $y(t) = u(t)$ for all $t \in [t_0, t_1]$.

Here we assume we have the ability to "clamp" the Y portion of the state to any trajectory we want.

Assumption 4 $\gamma_x(t)$ is the globally (within M) asymptotically stable attractor of the non-autonomous system $\mathbf{X}(\gamma_y)$.

This ensures that by feeding in the y dynamics of the limit cycle we can reset the system in its entirety to the limit cycle.

Assumption 5 There exists a neighborhood $U \supset \Gamma$ and a neighborhood $V \ni \gamma(0)$ such that, for every $(x_f, y_f) \in U$ and $(x_0, y_0) \in V$, there is a known C^1 control $u : [0, N] \to Y$ satisfying $u(0) = y_0$, $u(N) = y_f$, and steering $x(0) = x_0$ to x_f for the $\mathbf{X}(u)$ dynamics.

This implies that from some specially chosen limit cycle state $\gamma(0)$, we have the ability to create an open set of states around the limit cycle, i.e. that we can produce a sufficiently rich set of perturbations to reproduce all possible behaviors that can occur near the limit cycle.

Remark 1 Assumption 4 may raise some concerns as it is not immediately obvious that a system like $\mathbf{X}(\gamma_y)$ would be stable or have $\gamma_x(t)$ as its attractor. In the appendix we show, as a motivating example, that the required property holds for the Hopf oscillator.

Remark 2 In general, we do not know the delay map R. However, any trajectory of $\mathbf{X}(u)$ that we compute with a segment of time $[s, t]$, $t - s \geq \delta_m$, that evolves autonomously (i.e., such that during $[s, t]$ we choose $u(t) = y(t)$ with $\dot{y}(t) = f_y(x, y)$, thereby not externally forcing the system) gives us a delay coordinate point by observing $y|_{[s,t]}$. For any hypothesised compatible $x(t)$, this would allow us to have sampled the value of R at the point $(x(t), y(t))$ and found it to be $(y(t), y(t - \delta_1), \ldots, y(t - \delta_m))$, giving a sampling of a R map compatible with this hypothesis.

5.2 PRC recovery procedure

Under these assumptions we propose to recover the PRC from delay coordinate measurements and Y clamping experiments as follows.

5.2.1 Preparation

Compute from y-observations over a long enough time an estimate of the period T and γ_y. Knowing the period enables us to match the noisy observations of γ_y that are an integer number of periods apart and thus correspond to observing the same underlying value multiple times. The repeated observation allows for better estimation, e.g. via averaging. Having γ_y also provides us with the image $R(\Gamma)$ of the limit-cycle point-set Γ under the coordinate change R into delay coordinates. We can also obtain the phase on the limit cycle itself, $\varphi|_\Gamma$ through the use of a phase estimator on the γ_y data.

5.2.2 Sampling

Now consider a fixed u and a time N.

(1) Let the $\mathbf{X}(u)$ (physical, un-disturbed) system run autonomously, producing $z(t) := (x(t), y(t))$ comprising our not known as of yet $x(t)$, and our directly observable $y(t)$. After δ_m units of time, the trajectory is long enough to allow the delay map image $R(z(t))$ to be computed, and by Assumption 2 this image is smoothly embedded. Continue autonomously until $|R(z) - R(\gamma(0))|$ is sufficiently small to allow us to treat it as $z = \gamma(0)$. The system has "reset"; we therefore define the time satisfying the reset condition as a new $t = 0$. (For our Fitzhugh–Nagumo example the reset condition occurs when an arbitrary point on the limit cycle is passed; we selected this point by integrating five cycles from the position $v = 0, u = 1$.)

(2) Create a perturbed initial condition by forcing $y(t) := u(t)$ for $t \in [0, N]$, leading to $z(N) \in U \backslash \Gamma$. This should be possible given Assumptions 3 and 5. We will later select these $z(N)$ to comprise a collection of points that enable us to estimate partial derivatives of $\varphi(\cdot)$.

(3) Let the system run autonomously for at least δ_m more time units. $R(z(t))$ is known again for any $t > N + \delta_m$. Thus we can determine when $z(t)$ returns to Γ to within sufficiently good accuracy by checking $R(z)$ for approximate membership in $R(\Gamma)$ through finding the minimum distance from a fourier series model of $R(\gamma(\varphi))$. Assume the time at which this happens is t_F, i.e., we approximately have $z(t_F) \in \Gamma$. We can approximately obtain $\varphi_F := \varphi(z(t_F))$ because in Sect. 5.2.1 we obtained the phase on Γ from $R(z(t_F))$. From this and the period we get the phases for the entire time segment as $\varphi(z(t)) = \exp(-i2\pi(t_F - t)/T)\varphi_F$.

Consider now from these trajectories the subset of points $R(z(t))$ sufficiently close to $R(\Gamma)$ for which we also know $\varphi(z(t))$. This set of points will allow us to estimate the ∂_y-PRC with respect to delay coordinates (cf. Arnold's chastening in Sect. 3).

At this time our data collection is complete. Note that, importantly, we have not used any knowledge about a specific x or f_x other than presuming Assumptions 2, 3, and 5. Despite x being neither known nor selected, our experimental data is fixed from here on.

5.2.3 Reconstructing *R*

Now we choose a specific f_x that would meet our assumptions. Using Assumptions 1 and 4 and knowledge of γ_y from Sect. 5.2.1, we numerically integrate the $\mathbf{X}(\gamma_y)$ system for a long enough time to determine $\gamma_x(0)$ as accurately as desired.

Consider again the process of Sect. 5.2.2. For each trajectory computed therein, starting with the initial condition $z(0) = \gamma(0)$ which is now fully known, we integrate the control system $\mathbf{X}(y)$ to give the trajectory $z(t)$. For the times $t > N + \delta_m$, the now known $x(t)$ and $R(z(t))$ provide us with samples of the map $R(\cdot, \cdot)$. Using this collection of samples, we interpolate the map R and its Jacobian $\mathsf{D}R$. By recon-

structing R and its Jacobian we have acquired the means to convert any geometric object between delay coordinates and the original coordinates.

This allows us to use (24) to recover the desired PRC in the (x, y) coordinates. The PRC is the application of the temporal 1-form in delay coordinates $d\tau_{\text{delay}}$ to whatever vector field moves in the direction of the perturbation with respect to which we want the PRC; here $[\partial_y]_{\text{delay}}$ for the y-PRC[5]. This PRC in delay coordinates is $\langle d\tau_{\text{delay}}, [\partial_y]_{\text{delay}}\rangle$. The correct transformation to (x, y) coordinates maintaining the inner product is $\langle (\mathsf{D}R)^{\mathsf{T}} d\tau_{\text{delay}} \circ R, (\mathsf{D}R)^{-1}[\partial_y]_{\text{delay}} \circ R\rangle$, which equals $\langle d\tau_{x,y}, [\partial_y]_{x,y}\rangle$.

6 An example from neuroscience

Whereas the mathematical observations in the former sections and Sect. 8 is perhaps noteworthy, it is hardly surprising to a geometrically-minded mathematician. However, these observations could have significant practical applications. We therefore switch our focus to numerical approaches that could be applied to (possibly noisy) experimental data. For this reason, in this section we will analyze data produced from numerically simulating a model system as if we were experimentalists without knowledge of the underlying equations of motion. We will conduct this analysis based on plausible assumptions informed by numerical tests, highlighting the places where analytical or computer-assisted proofs would perhaps be nice to have, but are beyond the scope of typical experimental work.

For our example we consider the Fitzhugh–Nagumo (FHN) system, often used as a model of a simple neuronal oscillator. A standard form of the equations for this system is:

$$\dot{v} = c\left(v + w - \frac{v^3}{3} + z\right) \qquad \dot{w} = \frac{a - v - bw}{c} \tag{17}$$

Here voltage v, gating variable w, and time have all been suitably normalised. Differentiation with respect to time is indicated by a dot. We adopted the values 0.7, 0.8, 3 and -0.4 for a, b, c and z respectively.

The time series data we obtained by numerically integrating (17) at these parameter values strongly suggests that there is an exponentially stable limit cycle with a large basin of attraction; we henceforth assume this is the case without proof. We modeled the limit cycle with two univariate splines (one for each coordinate; `scipy.interpolate.UnivariateSpline`) based on 2000 points with the number of knots set to interpolate through all points.

Following the method of Sect. 5, we attempted to use delay coordinates (v, v_d) with a single delay of length 0.1 units to build a full picture of the dynamics using only observations of v. The rationale behind using such a voltage-centered model is that when observing a neuron we are typically able to record its voltage but not the current flowing through it or the status of its ion channels (here represented by the gating variable). The delay coordinates (v, v_d) are implicitly defined as functions of (v, w):

$$v = v \qquad v_d = g(v, w)\,. \tag{18}$$

Using trajectory data, we generated an approximation of the graph of the map $(v, w) \mapsto g(v, w) = v_d$ shown in Fig. 3. This approximation strongly suggests that $\frac{\partial g}{\partial w}|_{v=\text{const.}} > 0$ in a neighborhood of the limit cycle, which we henceforth assume without proof. This assumption implies that the map $(v, w) \mapsto (v, v_d)$ is a diffeomorphism from a (possibly smaller) neighborhood of the limit cycle onto its image, and so the delay coordinates (v, v_d) yield a smooth and invertible change of coordinates on this neighborhood.

As in Sects. 3 and 4, we let $\varphi\colon B \to S^1 \subset \mathbb{C}$ be the asymptotic phase map defined on the basin of attraction B of the limit cycle. The restriction of $\frac{1}{i\varphi}\frac{\partial \varphi}{\partial v}|_{w=\text{const.}}$ to the limit cycle is the PRC curve associated with rapidly perturbing the system's voltage, so quickly that the (unobserved) gating variable remains effectively fixed, i.e. changes at an asymptotic order slower, where the perturbation is $\mathcal{O}(t)$ and the gating changes $\mathcal{O}(t^2)$. If we were to mistakenly use $\frac{1}{i\varphi}\frac{\partial \varphi}{\partial v}|_{v_d=\text{const.}}$ instead, our result will be in error by a term obtainable from the Jacobian of the coordinate transformation (18), as explained in Example 1 and demonstrated in the example of Sect. 4.1. The notion of computing the partial derivative of v while holding v_d constant may seem confusing because the deterministic nature of the dynamics does not allow a delayed value to be held constant while the value at any other time changes. The partial derivative in question is better viewed as the slope in the v direction of the manifold representing the $(v, v_d, \dot{v}, \dot{v}_d)$ graph—it contains part of the information needed to reconstruct the dynamics, but is not a direction of change that corresponds to anything the system could do, because v and v_d are intertwined by those dynamics and cannot be independently modified.

We now describe a simulated experiment using this FHN model. The experiment consisted of a noise free simulation of a FHN model, which we used to establish the "true" PRCs, followed by a stochastic simulation of the same model, producing simulated experimental data, which we then use to demonstrate that similar results can be obtained from realistically noisy data.

[5] The notation ∂_z, a vector, is not the partial derivative $\frac{\partial}{\partial z}$; it is the derivation along z, i.e. given any path $p(t)$, $d/dt\, z(p(t)) = \langle \partial_z, \dot{p}(t)\rangle$. Thinking of z as scalar valued function, and ignoring issues of the presence or lack thereof of an underlying metric, ∂_z is the gradient of z, and the previous equation is an expression of the chain rule.

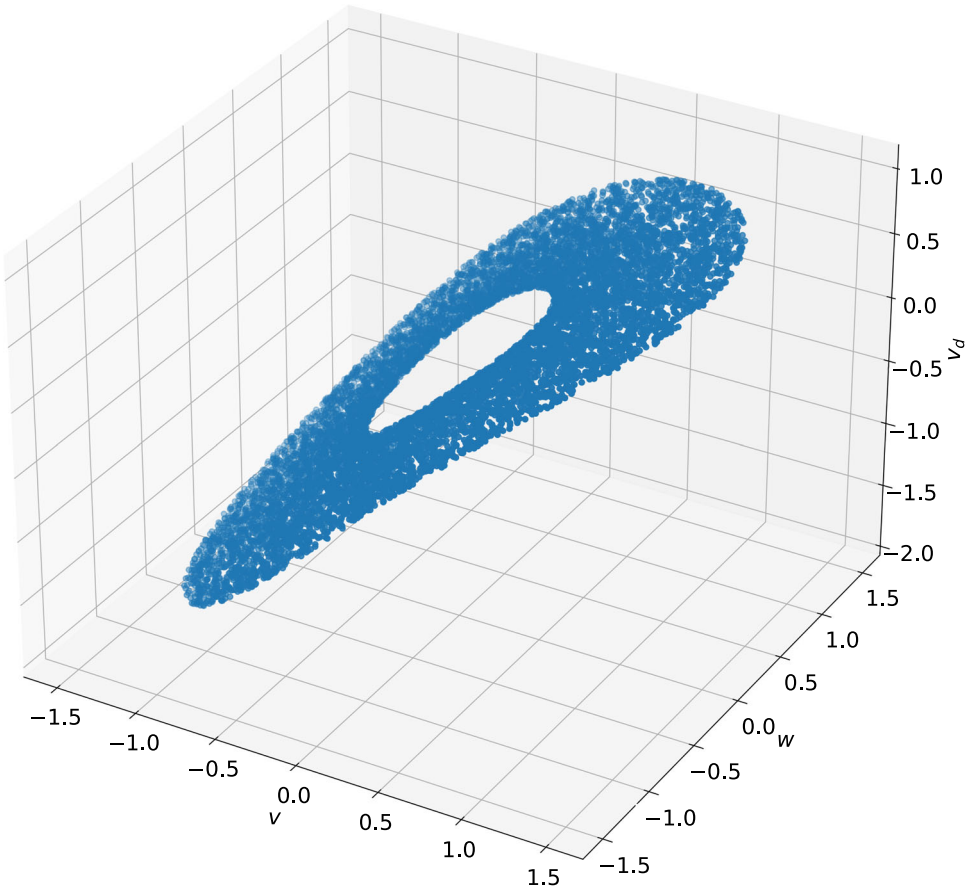

Fig. 3 Shown here are sample points illustrating the coordinate transformation $(v, w) \mapsto (v, v_d)$ for the Fitzhugh–Nagumo system (17) implicitly defined by delaying the v coordinate. Here v is the normalised voltage, w is the gating variable, and v_d is the delayed normalised voltage

6.1 Noise-free "ground truth"

Because we conducted a simulated experiment we obtained the phase response curve via a perturbation approach which could be available to an experimentalist, rather than e.g. using the adjoint method (Malkin 1949, 1959; Hoppensteadt and Izhikevich 1997) which would require knowledge of the dynamical equations of the system.

6.1.1 Obtaining the period *T*

For forward integration we used `scipy.integrate.odeint`, with the Jacobian provided, and tolerances for the error control parameters `atol` and `rtol` set to a low tolerance of 10^{-10}. We recorded a sufficiently large collection of steady-state oscillations, so large as to enable an initial estimate of the period of the limit cycle. We then estimated the period by forward integration from a point on the limit cycle. This initial point was obtained by integrating forward in time approximately one thousand periods.

We then used `scipy.optimize.fmin` to find the integration time that minimises the function returning the distance between the start and end point of an integration using with 10^{-12} tolerance for both the function value and the location of the optima. This provided a highly accurate sampling of our periodic solution, γ of Sect. 2.

This periodic solution also enabled us to model the limit cycle of the system in delay coordinates, as per Sect. 5.2.1, to our desired accuracy.

6.1.2 Voltage perturbations

We then went on to evaluate the PRCs by sampling. We allowed the system's state to evolve until we were confident it was close to the limit cycle, at which point we applied a small change to voltage (10^{-2} in the normalised units of the Fitzhugh–Nagumo system), instantly driving the voltage slightly above or below the voltage when resting on the limit cycle, while at the same time leaving the gating variable unchanged.[6] We allowed the system to settle for a fixed amount of time equal to five periods (as estimated above),

[6] We note that many characterisations of neural oscillators in terms of phase do not apply a small voltage in order to estimate the PRC, but rather apply a spike (large voltage, short duration) to obtain a PRC which might be inconsistent with a first order approximation. Such an experiment is not considered here, although we note that what is required to calculate the induced phase change across the stability basin is some method of integrating the accumulated phase change resulting from travelling from the limit cycle to the perturbed state. Because the

and then recorded the resulting phase shift as described in the next section. Because we waited an integer number of periods, the unperturbed base-point on the limit cycle would have returned to itself, simplifying the calculation of the phase response.

6.1.3 Obtaining phases

We took our estimate of the limit cycle and assigned to points on this curve a phase which advances uniformly in time.

We evaluated phase on other trajectories by: (1) integrating forward in time until they were "sufficently" close to the limit cycle; (2) obtaining an estimate of the phase of the final point on the trajectory; and (3) computing the phases of all previous points from their times and the period T.

To estimate the phase of a final trajectory point we found the nearest two points in our high resolution model of the limit cycle described above, and orthogonally projected the final trajectory point onto the line segment connecting them. We took the linear interpolation of the phases of the limit cycle points at the projected point as the phase of the final trajectory point.

Given the phase of the final point in a trajectory, we used the period and the sample times of the other trajectory points to assign self-consistent phases to all the points in each trajectory.

A pseudo-code summary of this approach to phase estimation would be:

(1) Compute the limit cycle and uniformly distribute phases on it
(2) Obtain a point of the limit cycle
(3) Perturb the point in voltage by a small amount
(4) Simulate the system through until sufficiently close to the limit cycle
(5) Assign phases to the points obtained from this simulation by uniformly distributing phase changes through time and ending on the same phase as the point of closest approach on the limit cycle

6.1.4 Ground truth PRC-s

All our perturbed initial points were obtained by an instantaneous change in voltage from a related point on the limit cycle. We therefore took the phase of the perturbed initial point and subtracted from it the phase of the related limit cycle point, giving a phase change. We divided this phase change by the change in voltage that created the perturbed initial point to obtain the PRC. Because the perturbation only changed v, but not w, this procedure gave us an estimate of the voltage-PRC with the gating variable held constant.

The "ground truth" PRC-s are shown in the Fig. 4 "small perturbation" line.

6.2 Simulated experimental data

Our simulated experimental data included measurement noise (we added a sample from a Gaussian distribution to our observed voltages, with standard deviation 0.1 units), but otherwise evaluated phase in a similar process. This produced a distribution of "measured" phases, as shown in Fig. 4 "phase change", which shows up as a coloured region.

Since our phase estimates for these noisy trajectory points relied on finding the nearest limit cycle point to the (noisy) final point of each trajectory, we also computed a windowed median of these estimates (see Fig. 4 "median phase change" line. We used the median rather than mean since our system produced results with variability that seemed to be heavy tailed; rank statistics are more stable under such conditions than moments tend to be).

For actual experimental data we should presume the presence of many types of noise: numerical errors, experimental system noise, and experimental measurement noise. Under such conditions, the finite difference approximation of the derivatives we wished to compute would have poor statistical properties. A more noise-resilient method is to take small segments of our data at points around the limit cycle in delay coordinates, confine ourselves to slices of the data in the delay coordinate v_d, and perform a linear regression of the final phase against change in voltage v. However, this method will calculate the PRC with a fixed delay coordinate (Fig. 4 "constant delay" plot) which, for the reasons we explained in Example 1 and demonstrated in Sect. 4.1, will not be the voltage PRC with fixed gating variable which we desire.

Recall that here we attempted to analyze numerically generated data as if this data came from experimental measurements, for which we have no underlying equations of motion (17). Unlike in the example of Sect. 4.1, there is no way to construct the delay coordinate transformation in closed form.[7]

6.2.1 Estimating PRC using nearby points

We constructed a Taylor approximation of the delay coordinate transformation from points in small regions. Each region

Footnote 6 Continued
ensuing state change is typically "large", i.e., of the order of the size of the limit cycle itself, infinitesimal approximations are unlikely to be accurate.

[7] Even if we were assuming knowledge of the equations of motion, it seems likely to us that (17) is sufficiently complicated that constructing the delay coordinate transformation in closed form is intractable. We were able to construct the delay coordinate transformation in closed form for the toy example of Sect. 4.1 since its form was chosen for this very purpose.

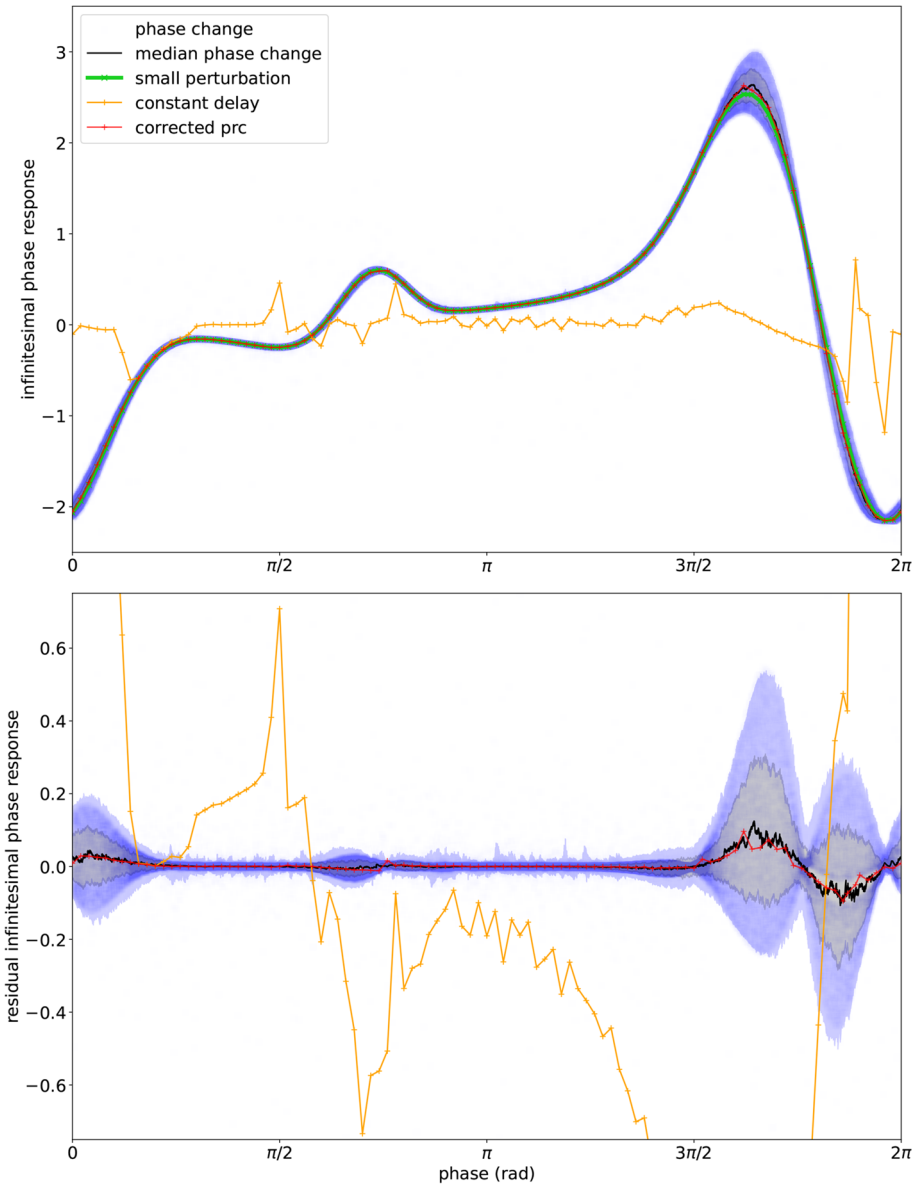

Fig. 4 Plot of the normalised voltage PRC of the Fitzhugh–Nagumo system against the phase of the oscillator. We show a ground truth PRC (green) which we calculated numerically from a small perturbation, our sampled observations, a rolling median of the observations (black), $\frac{1}{i\varphi}\partial_v\phi|_{v_d=\text{const.}}$ (orange), and $\frac{1}{i\varphi}\partial_v\phi|_{v_d=\text{const.}}$ corrected with the Jacobian of the delay map to give $\frac{1}{i\varphi}\partial_v\phi|_{w=\text{const.}}$ (red). We computed the sampled observations by taking the difference between the phase a fixed time after the perturbation and the phase predicted by simulating the unperturbed system forward in time the same amount, dividing by the magnitude of the voltage perturbation, leading to variable results (blue dots; percentiles 2 and 98 light blue lines; 25 and 75 grey lines). Because some of the plots are hard to distinguish, we also plotted the same, but with the forward simulation (green, top subplot) subtracted (bottom subplot; all other line types unchanged)

was centered at a limit cycle base-point and contained the 0.5% observations in our training set nearest (by Euclidean distance) to that base-point (Fig. 5). Initially, we evaluated the Jacobian of the delay coordinate transformation by total least squares regression of the region points against their transformed images. However, this produced several numerical anomalies we wished to explore further.

To make it easier for us to interpret and sanity-check the coordinate transformation, we further sub-divided the regions into high aspect-ratio horizontal and vertical strips (Fig. 5, left sub-plot, blue and red). For a region sufficiently small to have a nearly constant Jacobian, a scatter plot of the major axis coordinate of a strip against a transformed output coordinate value is expected to be a line segment whose slope is the respective element of the Jacobian (Fig. 6 bottom 4×4 grid). Inspection of this kind of scatter plot for behaviour distinct from this linear structure allowed us to resolve the numerical issues we encountered.

A pseudocode summary of this method would be:

(1) Assign phases to observations as in the previous subsections
(2) For every point on the limit cycle:
    (a) Take a thin slice of voltages at a point close to the limit cycle and regress the assigned phases against the delayed voltage to obtain the phase response for changes in the delayed voltage at fixed voltage
    (b) Repeat the above step for thin slices in delayed voltage regressing phase against voltage

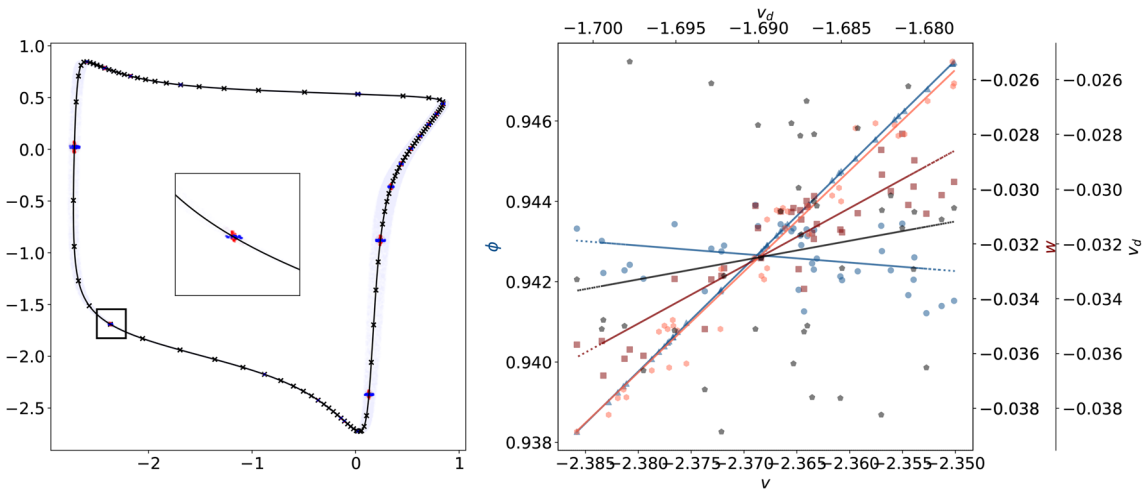

Fig. 5 Process for calculating the PRC in delay coordinates by correcting using the Jacobian of the transformation from conventional to delay coordinates. We plotted the delay coordinate data for 100,000 points (left subplot; sampled points in pale blue; x-axis is v, y-axis is v_d), indicated the limit cycle (black line) and the points at which we calculated the PRC (black crosses). At every fifth such point we indicated vertical and horizontal subsets of points we used to compute the Jacobian with (gem blue and red). We plotted an illustrative zoomed in view of one such location (boxed on limit cycle; center inset). To demonstrate the quality of the various linear fits this process required, we plotted the individual linear input–output relationships (right subplot). This includes the linear fits of the phase and the Jacobian of the transformation between (v, v_d) and (v, w) coordinates at the location of the inset. The plot shows both functions of v (with respect to the bottom x-axis; dotted lines), and functions of v_d (top x-axis; solid lines). We show phase change (blue; circular and triangular markers); w change (orange and red; square and hexagonal markers); and v_d (black; pentagonal markers)

(3) Compute the Taylor expansion of the current and voltage in terms of the gating variable and delay voltage and compute the Jacobian at every point on the limit cycle
(4) Perform the co-ordinate transform to compute the phase response with respect to voltage at a fixed gating variable

6.2.2 Auto-scaled local linear regression

At most points along the limit cycle the following procedure worked well. We selected a region size, and within that region selected 20% of the points in the region with the least change in the variable being held constant. Using linear regression, we calculated the PRC for both v and v_d, with v_d and v held constant, respectively. In addition, we calculated the Jacobian of the transformation from (v, w) to (v, v_d) by linear regression. We then reduced the size of the region of admitted points until the R^2 statistic of the linear fit increased above an arbitrarily chosen threshold of 0.75. This automatic scaling ensured that we were sampling a region within which the higher order Taylor polynomial terms do not dominate the results.

However, in some regions the coordinate transformation for the delay we chose had singularities ("folds") coming close to the limit cycle. As a consequence, the image points of some of the strips were clearly quadratic (see Fig. 8 or bi-modal (Fig. 6 top 4×4 grid). In these cases too the auto-scaling helped prevent the numerical problems we initially encountered.

Armed with a high quality estimate of the Jacobian of the coordinate change $(v, w) \mapsto (v, v_d)$, we calculated the desired PRC $\frac{1}{i\varphi}\frac{\partial\varphi}{\partial v}|_{w=\text{const.}}$ using the (v, v_d) coordinates and the estimated Jacobian as in Example 1 and in Sect. 4.1 (see Fig. 4 "corrected PRC"). There is excellent agreement between the results of this calculation, using simulated experimental data with delay coordinates, and that obtained numerically by the other calculation methods. Furthermore, the computation we performed allowed us to obtain a voltage PRC with the gating variable constant that very closely follows that which we could have obtained by computing the median of an enormous number of experiments (see Fig. 4 bottom sub-figure for detail of differences between PRC calculation results).

It is *crucial* to note that this entire simulated experimental protocol did not require any measurements of the gating variable.

7 Conclusion

We have pointed out both the danger and opportunity inherent in attempting to construct oscillator models including their phase response curves from measuring only a subset of the state variables. Such an approach could allow us, for example, to fully characterize an FHN-like neuron by only measuring and perturbing "voltage". However, the resulting PRC will then be correct only with respect to the delay embedding we

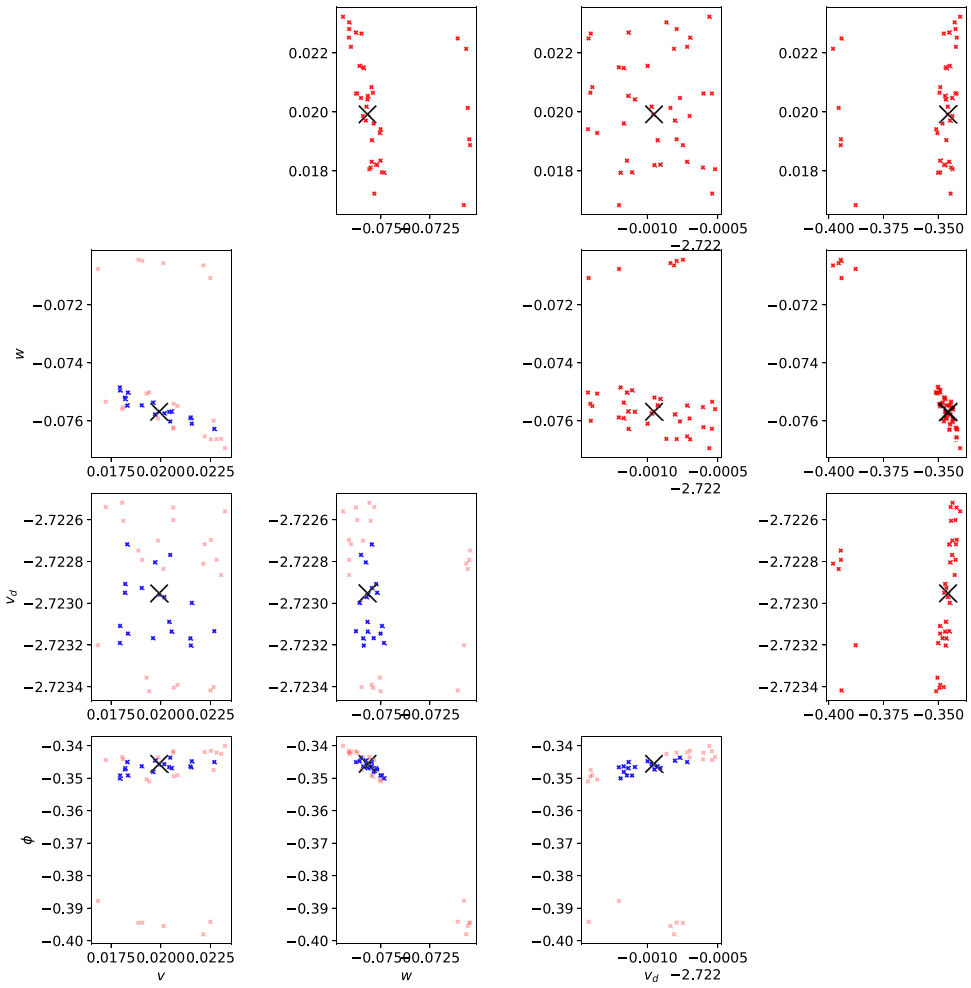

Fig. 6 Scatter plot matrices of the dimensionless voltage, gating variable, delayed voltage; and phase (red, blue) around a point (black) on the limit cycle. The transformation to delay coordinates has a singularity in it which causes bimodality (red, top-right half of matrix, low alpha in the bottom-left); this is removed by excluding those observations which are separated from the point about which we are calculating our derivatives (blue, bottom left)

use to describe the dynamics. Since in practice, some other non-observed quantities (such as "current" or "gating variable") are held constant while coupling to the neuron, we then showed how a model of those quantities can be used to convert the PRC to the correct coordinates. We demonstrated this computation in a form that is robust to noise and potentially suitable for the analysis of experimental data. This process can be fraught with numerical difficulties if the coordinate change required is nearly singular close to the limit cycle. We have shown techniques which helped us ascertain when and where such problems appeared in our data.

Our approach provides a rapid method for obtaining a PRC in cases where the unobserved dynamics (the gating variable in our neuroscience example) are well known. However, from a neuroscience perspective there is an important lesson to be learned here: for any given set of voltage measurements, there can be a variety of models for current or gating variable which could be assumed with equal validity, and the resulting PRC would accordingly change. Our approach gives the "correct" PRC for an experimental system if the model of the state variable that remained constant under perturbation is correct. If the model is not correct, there is no guarantee that the PRC will be accurate. This dependency might be a means for eliminating among multiple such models, provided the voltage measurements yield an alternative method for measuring the PRC against which the predictions from different current or gating variable models can be tested. As new current or gating variable generation dynamics are hypothesized, it may be possible to use the differences in PRC that would be produced by these new models to discern which "current" or "gating variable" is consistent with a given neuron's prerecorded voltage measurements—all without requiring any new data collection.

The approach we present here is by no means limited to neuroscience applications. Future applications of this work include producing reliable and robust models of coupled oscillator systems in any domain in which oscillator models are used.

In this paper we considered only classical asymptotic phase for deterministic oscillators. However, there are also multiple notions of stochastic phase defined for stochastic oscillatory systems (Schwabedal and Pikovsky 2013;

Thomas and Lindner 2014; Cao et al. 2020; Engel and Kuehn 2021) (see Pikovsky 2015; Thomas and Lindner 2015 for a spirited discussion), and it would be interesting to investigate these notions in the context of the present work.

8 The temporal 1-form view of phase response curves

In this final, primarily technical, section we present a detailed account of how the "*temporal 1-form*" can be obtained mathematically, and how this coordinate invariant object relates to all PRC-s. The original idea for the temporal 1-form extends several years back; our first public presentation of this concept was in Wilshin and Revzen (2014), and additional details can be found in Wilshin et al. (2014). The insight that the "gradient" of phase is free of the gauge indeterminacy and jump discontinuity that appears in the commonly used real-valued phase goes back at least to Takajo and Takahashi (1988), where this insight was used for improved recovery of the phase of Fourier transformed images. The treatment there did not consider phase in the dynamical systems sense at all. In Guillamon and Huguet (2009) the importance of this gradient is also reiterated, and attributed also to Izhikevich (Izhikevich (2007), Ch. 10). In a sense the "phase response surface" defined and discussed there is a conjugate object to the temporal 1-form, but the computational approaches used there assume the ability to accurately integrate the dynamics—an approach not very amenable to the analysis of experimental data. More recent work on modeling oscillators using machine learning methods has shown the ability to recover PRC-s from machine learning models trained to simulate the dynamics of the system, e.g. using recurrent neural networks (Cestnik and Abel 2019). It is clearly the case that a faithful simulation of the dynamics, either by integration as in Guillamon and Huguet (2009) or through machine learning as in Cestnik and Abel (2019), can allow one to reconstruct the PRC-s. Knowing the true dynamics precisely, or learning a faithful model of the dynamics from the data are very demanding requirements compared to the requirements for estimating the temporal 1-form as described in Wilshin et al. (2014), where we additionally provided a convergence proof and estimate of convergence rates as a function of the properties of the noise. Furthermore, the coordinate independent nature of 1-forms captures and expresses the implications of coordinate choice which are the topic of the current manuscript.

Consider the differential 1-form $d\theta \in \mathsf{T}^*S^1$ measuring change of angle on the unit circle. Viewing the complex plane as $\mathbb{R}^2$ and the unit circle S^1 geometrically, at any point $\phi \in S^1$, this form is dual to the (unit) vector $i\phi$, pointing tangent to the circle and counter-clockwise. If $\phi : \mathbb{R} \to S^1$ is a function of time, the angle change it undergoes in a short interval dt is $(-i/\phi)\langle\dot{\phi}, d\theta\rangle dt$, where this inner product $\langle\cdot,\cdot\rangle$ can be viewed as the standard $\mathbb{R}^2$ inner product, or more accurately as the contraction of the co-vector $d\theta \in \mathsf{T}^*_{\phi(t)}S^1$ with a vector $\dot{\phi} \in \mathsf{T}_{\phi(t)}S^1$. This contraction perspective is coordinate independent and does not require a Riemannian metric.

The pullback of $d\theta$ through φ (φ as defined in Sect. 2) denoted by $\varphi^* d\theta$ is a differential 1-form on B the domain of φ which satisfies the property that for any (smooth) path $p : \mathbb{R} \to B$, $\langle \frac{d}{dt} p(t), \varphi^* d\theta(p(t))\rangle = \langle \frac{d}{dt}\varphi(p(t)), d\theta(p(t))\rangle$. We define the (real) differential 1-form

$$d\tau(x) := (\varphi^* d\theta)(x). \tag{19}$$

Consider now a trajectory $x(t)$ of the dynamical system $\dot{x} = f(x)$. We already know that the complex valued asymptotic phase φ satisfies $\frac{d}{dt}\varphi(x(t)) = (i2\pi/T)\varphi(x(t))$. It follows that

$$\begin{aligned}\langle f(x), d\tau(x)\rangle = \langle \dot{x}, d\tau(x)\rangle &= \left\langle \frac{d}{dt}\varphi(x), d\theta \right\rangle \\ &= (2\pi/T)\langle i\varphi(x), d\theta\rangle = 2\pi/T,\end{aligned} \tag{20}$$

i.e. the 1-form $d\tau$ measures the rate of change of phase along paths and as such identically gives the constant $2\pi/T$ on the vector field $f : M \to \mathsf{T}M$ that generates the dynamics. Because this 1-form, which we call the "temporal 1-form", measures the rate of change of phase along any direction in $\mathsf{T}M$, it contains the information about all possible PRC-s in a single coordinate-independent object.

We can therefore define the conventional PRC as follows.

Definition 2 Let $Z : \Gamma \to \mathsf{T}B|_\Gamma$ assign to each $x \in \Gamma$ a vector $Z(x) \in \mathsf{T}_x B$. Then the *infinitesimal Z phase response curve (PRC)* $\rho_Z : \Gamma \to \mathbb{R}$ is defined by

$$\rho_Z(x) := \langle d\tau, Z\rangle(x). \tag{21}$$

Remark 3 The connection of Definition 2 with the PRC definition of Sect. 3 is as follows: take $B = \mathbb{R}^n$ and let $Z = \partial_{x_1}$ be the first coordinate vector field defined by a choice of coordinates $(x_1, x_2, \ldots, x_n)$ for $\mathbb{R}^n$.

If the last $(n-1)$ coordinates are changed to produce a coordinate system $(x_1, y_2, \ldots, y_n)$, then with respect to these new coordinates the first coordinate vector field $\tilde{Z}$ is generally not equal to Z; it is the pushforward of Z along the coordinate change $\Psi : (x_1, x_2, \ldots, x_n) \mapsto (x_1, y_2, \ldots, y_n)$. Explicitly,

$$\tilde{Z} = \Psi_* Z := \mathsf{D}\Psi \circ Z \circ \Psi^{-1}.$$

In light of Definition 2 and Remark 3, this explains the observation made at the end of Sect. 3.

Remark 4 More generally, let $B \subseteq M$, $\Psi\colon M \to \tilde{M}$ be a diffeomorphism. The vector field f pushes forward to a vector field $\tilde{f} = \Psi_* f$ on $\tilde{M}$ having a stable hyperbolic T-periodic orbit with image $\tilde{\Gamma} = \Psi(\Gamma)$, basin of attraction $\tilde{B} = \Psi(B)$, and asymptotic phase map $\tilde{\varphi}\colon \tilde{B} \to S^1$ given by $\tilde{\varphi} = \varphi \circ \Psi^{-1}$. Given $Z\colon \Gamma \to \mathsf{T}B|_\Gamma$, we define $\tilde{Z}\colon \tilde{\Gamma} \to \mathsf{T}\tilde{B}|_{\tilde{\Gamma}}$ by $\tilde{Z} := \mathsf{D}\Psi \circ Z \circ \Psi^{-1}|_{\tilde{\Gamma}}$. The definitions of pullbacks and pushforwards immediately imply that, on $\tilde{\Gamma}$:

$$\begin{aligned}\rho_Z \circ \Psi^{-1} &= \langle \varphi^* d\theta, Z\rangle \circ \Psi^{-1} = \langle (\Psi^{-1})^* \varphi^* d\theta, \Psi_* Z\rangle \\ &= \langle (\varphi \circ \Psi^{-1})^* d\theta, \Psi_* Z\rangle \\ &= \langle \tilde{\varphi}^* d\theta, \tilde{Z}\rangle =: \tilde{\rho}_{\tilde{Z}}. \end{aligned} \tag{22}$$

Hence one can obtain the PRC ρ_Z from data on $\tilde{M}$ using $\rho_Z(x) = \tilde{\rho}_{\tilde{Z}}(\Psi(x))$.

Example 1 Here we will illustrate this phase extraction process, and the limitation previously identified, using a simple worked example. For the reader interested in comparing the general case of examples discussed here with the detailed example worked out in Sect. 4.1, the following table provides some key correspondences.

Section 4.1	Here
$\mathbb{R}^2\backslash\{\mathbf{0}\}$	M
$(w, 1-r)$	$f(x)$
$\{r = 1\}$	Γ
$\mathbb{R}^2\backslash\{\mathbf{0}\}$	B
$d\theta$	$\varphi^* d\theta$
$(0, 1)$	Z
x/r^2	$\rho_Z(x)$
$(\frac{\partial u}{\partial y}, \frac{\partial v}{\partial y}) \neq (0, 1)$	$\tilde{Z}$

Let Φ be the local flow generated by f; since we are only interested in dynamics on B, we may assume without loss of generality that $\Phi\colon \mathbb{R} \times M \to M$ is a flow. Given $0 = \delta_0 < \delta_1 < \delta_2 < \cdots < \delta_m$ and a smooth function $h\colon M \to \mathbb{R}$, define the delay embedding coordinates taken from delays of observable h, $R\colon M \to \mathbb{R}^{m+1}$ via

$$R(x) := (h(x), h \circ \Phi^{-\delta_1}(x), h \circ \Phi^{-\delta_2}(x), \ldots, h \circ \Phi^{-\delta_m}(x)). \tag{23}$$

Let us assume that R is a smooth embedding so that the codomain-restricted map $\Psi := R\colon M \to R(M) =: \tilde{M}$ is a diffeomorphism.

Because R is locally just a choice of coordinates like any other, we can express PRCs associated with the vector field f from observations on $\tilde{M}$ just as easily as on M using Remark 4: given any vector field $Z\colon \Gamma \to \mathsf{T}B|_\Gamma$ over Γ, the Z-PRC $\rho_Z\colon \Gamma \to \mathbb{R}$ is given by

$$\begin{aligned}\rho_Z(x) &= \langle \tilde{\varphi}^* d\theta, \Psi_* Z\rangle(\Psi(x)) \\ &= \left\langle \tilde{\varphi}^* d\theta, \sum_{k=0}^{m} (\mathsf{D}h \circ \mathsf{D}\Phi^{-\delta_k} Z(x)) \partial_{x_k} \right\rangle, \end{aligned} \tag{24}$$

where ∂_{x_k} is the derivation (i.e. partial derivative vector field) with respect to the k-th coordinate, x_k, in whatever coordinates $(x_0, x_1, \ldots, x_m)$ we choose for representing $\tilde{M} \subseteq \mathbb{R}^{m+1}$. The first equality follows here from Eqs. 21 and 22. The right hand equality follows from Eq. 23.

Appendix A. The Hopf oscillator satisfies Assumption 4

Consider the Hopf oscillator whose attractor is a trajectory around the unit circle with velocity 1.

Here $\gamma_x(t) := \sin(t)$, leading to

$$\dot{y} = (1 - \sin^2(t) - y^2)y - \sin(t)$$

as the non-autonomous y dynamics, which we expect to converge to $y = \cos(t)$.

Changing to an error variable $z := y - \cos(t)$, and shifting t by $\pi/2$ for algebraic convenience gives

$$\dot{z} = -z(z^2 + 3z\sin(t) + 2\sin^2(t)).$$

without loss of generality we can assume $z > 0$, since an initial negative z corresponds to a time shift of π in t, and we seek to prove a result for all t.

We need to show $z \to 0$ for all initial z. Define $L := \ln(z)$. We shall show that over a period of 2π, L decreases by at least some constant, thereby showing $z \to 0$ exponentially. We do this by taking a separate bound for each half-cycle,

$$\dot{L} = \frac{\dot{z}}{z} = -(z^2 + 3z\sin(t) + 2\sin^2(t)) \tag{25}$$

$$\begin{aligned}\int_0^{2\pi} \dot{L} dt &= \int_0^{2\pi} \left(\frac{1}{4}\sin^2(t) - (z + \frac{3}{2}\sin(t))^2\right) dt \\ &= \int_0^{\pi} \left(\frac{1}{4}\sin^2(t) - (z + \frac{3}{2}\sin(t))^2\right) dt \\ &\quad + \int_\pi^{2\pi} \left(\frac{1}{4}\sin^2(t) - (z + \frac{3}{2}\sin(t))^2\right) dt \\ &\leq \left[\left(\frac{1}{4} - \frac{9}{4}\right) + \frac{1}{4}\right] \int_0^{\pi} \sin^2(t) = -\frac{7}{8}\pi. \end{aligned} \tag{26}$$

Appendix B. Estimating the period of the Fitzhugh-Nagumo system

Below is a listing of the python code used to calculate the period of the Fitzhugh-Nagumo oscillator. This oscillator is slow-fast and quite stiff, so some care is needed when integrating it; the default settings of the scipy odeint command do not produce reliable results.

It is useful to supply the Jacobian and to have a large number of intermediate steps. In addition, for the precision we required for the period in this work, a higher than default tolerance was needed.

```
from numpy import stack,linspace
from scipy.integrate import odeint
from scipy.optimize import fmin

# Derivatives and partial derivatives of the Fitzhugh-Nagumo system.
def vdot(v,w,z,c): return(c*(w+v-((v*v*v)/3.)+z))
def wdot(v,w,a,b,c): return(-(v-a+b*w)/c)
def vdotdv(v,w,z,c): return(c*(1.-v*v))
def vdotdw(v,w,z,c): return(c)
def wdotdv(v,w,a,b,c): return(-1./c)
def wdotdw(v,w,a,b,c): return(-b/c)

def xdotfhn(p,t=0,z=-0.4,a=0.7,b=0.8,c=3.,tau=1.):
  '''Vectorised form of the Fitzhugh-Nagumo system'''
  x,y = p[...,0],p[...,1]
  return(tau*stack([vdot(x,y,z,c), wdot(x,y,a,b,c)],-1))

def Jxfhn(p,t=0,z=-0.4,a=0.7,b=0.8,c=3.,tau=1.):
  '''Jacobian of the Fitzhugh-Nagumo system (vectorised)'''
  x,y = p[...,0],p[...,1]
  return(tau*stack([
    stack([vdotdv(x,y,z,c),vdotdw(x,y,z,c)],-1),
    stack([wdotdv(x,y,a,b,c),wdotdw(x,y,a,b,c)],-1) ],-1))

def loopf(tc):
  '''
  Integrate the Fitzhugh-Nagumo system for a period of tc and return the
  difference between the initial and final point.
  Minimising this will give an integral multiple of the period.
  '''
  y1 = odeint(xdotfhn,y0,[0,tc],Dfun=Jxfhn,col_deriv=True,
    rtol=tol, atol=tol, mxstep=mxstep)[-1]
  return(sum((y0-y1)**2))

tc0 = 3*3.743      # Initial guess at the period
tol = 1.e-10       # Tolerance
trelax = 1000      # A long period of time, sufficient for the system to relax
Nstep = 1000       # Need extra steps for long initial integration
mxstep=5000        # Maximum number of intermediate steps for integrator
yinit = [1.0,0.0] # A starting location in stability basin

# Get a point on the limit cycle by integrating for a long period of time
y0 = odeint( xdotfhn, yinit, linspace(0,trelax,Nstep),
  Dfun=Jxfhn, col_deriv=True, rtol=tol, atol=tol, mxstep=mxstep )[-1]

# Minimise the squared difference between the starting position on the limit
# cycle and the final position on the limit cycle by adjusting the integration
# period, should have a minimum of zero to within tolerace after convergence.
tc1 = fmin(loopf,tc0,ftol=1e-12,xtol=1e-12)[0]
print("Period estimated to be:",tc1)
```

Appendix C. Special cases estimating the infinitesimal PRC

There are multiple special cases which must be considered when estimating the partial derivatives used to calculate the infinitesimal phase response curve close to the limit cycle for this system.

We consider two models. First, we select a model where the dependent variable (the phase for the infinitesimal phase response curve, the un-transformed coordinates for the Jacobian) is a polynomial function of the independent variable (the transformed coordinate). We consider the linear and quadratic cases, and select from these models via an AIC. An example where the linear model was superior is included in Figure 7.

For some cases this model is deficient. There is a clear non-linear relationship between the independent variable and dependent variable, such that a quadratic in the dependent variable is a more appropriate model. Such cases can be readily identified by inspection, the non-linearity is obvious and there is frequently bi-modality in the distribution of the dependent variable. In such cases a quadratic in the dependent variable is instead used. This is illustrated in Figure 8.

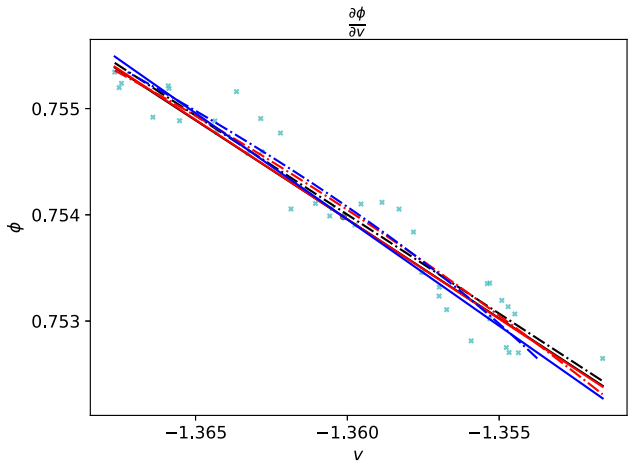

Fig. 7 An example where a simple linear model was used to estimate the rate of change of the phase against change in dimensionless voltage at a fixed point on the limit cycle (black point in center of plot). We plotted the phases of the points at fixed delay voltages v_d across a range of voltages v (cyan crosses). We also plotted various model fits (dot-dash lines): linear (black), quadratic in v (red), and quadratic in phase ϕ (blue). These imply potentially unequal derivatives which we indicate as lines with the associated slope (same color, solid thin lines). In this instance all fits give comparable estimates for the slope

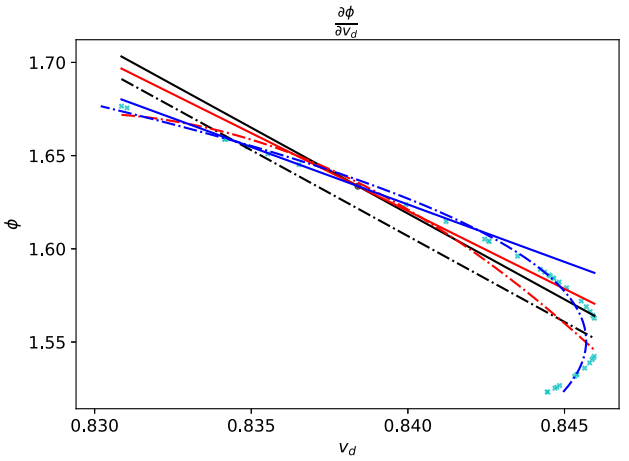

Fig. 8 An example where there is quadratic dependence of phase against delayed dimensionless voltage. The elements of this plot are as in Figure 7. The model with this quadratic dependence is clearly superior

In some cases the transformation from the conventional coordinates of the Fitzhugh-Nagumo system to the delay coordinates contains a singularity which is very close to the limit cycle. This was handled in two ways.

In most cases the approach of this singularity was close, but not so close that a reasonable estimate of the gradient could not be obtained by simply excluding manually those points which came excessively close to the limit cycle. An example of such a case is illustrated in Figure 9.

For one test point the singularity came extremely close to the limit cycle and no reasonable estimate of the phase response was possible. Instead two additional points were selected near to this example and the phase response estimated at these locations instead. This is illustrated in Figure 10.

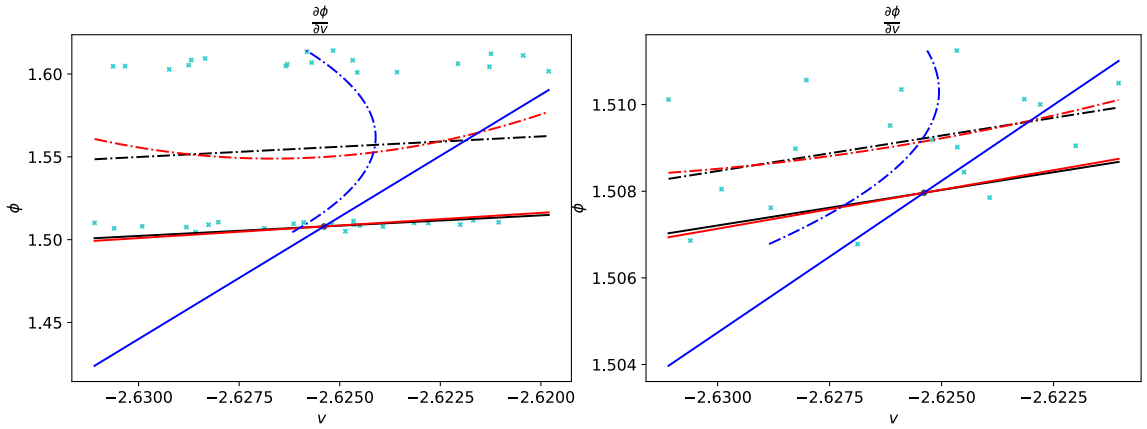

Fig. 9 An example where a singularity in the transformation to delay coordinates is present and close to the limit cycle. The elements of this plot are as in Figure 7. In this case, those points which are on the wrong side of the singularity can simply be excluded (right panel) and a better estimate of the partial derivative is obtained.

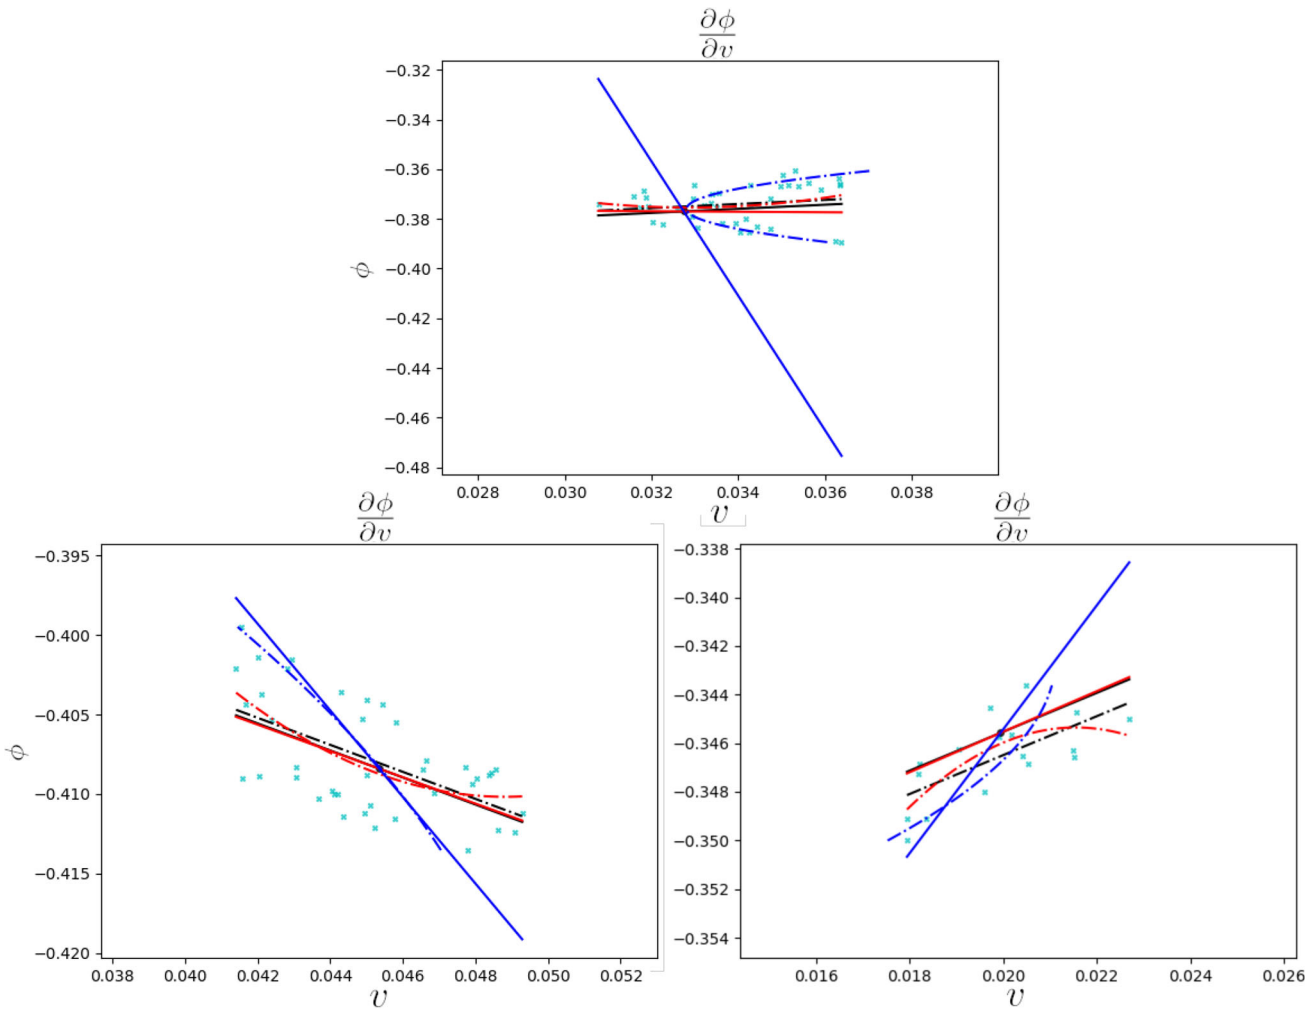

Fig. 10 Here, the singularity in the transformation to delay coordinates is too close to the limit cycle, and no reasonable estimate of the partial derivative can be obtained, illustrated in the top panel. The elements of this plot are as in Figure 7. Instead, two points (bottom left and right panels) close by on the limit cycle are considered and estimates for the partial derivatives are obtained there

Acknowledgements Kvalheim was supported by ARO MURI W911NF-18-1-0327, the SLICE Multidisciplinary University Research Initiatives Program; and by ONR N00014-16-1-2817, a Vannevar Bush Faculty Fellowship sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering. Revzen was supported by ARO W911NF-14-1-0573, ARO MURI W911NF-17-1-0306, NSF CMMI 1825918, and the D. Dan and Betty Kahn Michigan-Israel Partnership for Research and Education. We thank John Guckenheimer for the discussion leading to the temporal 1-form approach to phase quantification.

Author contributions All authors contributed to the writing of the manuscript. The formal proofs were undertaken by M.K. under supervision and with assistance of S.R. The numerical work was conducted by S.W under supervision and with the assistance of S.R.

Data availibility No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

References

Arnold VI (1989) Mathematical methods of classical mechanics, 2nd edn. Springer, Berlin

Cao A, Lindner B, Thomas PJ (2020) A partial differential equation for the mean-return-time phase of planar stochastic oscillators. SIAM J Appl Math 80(1):422–447

Cestnik R, Abel M (2019) Inferring the dynamics of oscillatory systems using recurrent neural networks. Chaos Interdiscip J Nonlinear Sci 29(6):063128

Engel M, Kuehn C (2021) A random dynamical systems perspective on isochronicity for stochastic oscillations. Commun Math Phys 386(3):1603–1641

Ermentrout B (1986) Losing amplitude and saving phase. In: Nonlinear oscillations in biology and chemistry: proceedings of a meeting held at the University of Utah, May 9–11, 1985. Springer, pp 98–114

Ermentrout GB, Terman DH (2010) Mathematical foundations of neuroscience, vol 35. Interdisciplinary applied mathematics. Springer, New York

Floquet G (1883) Sur les équations différentielles linéaires à coefficients périodiques. Ann Sci lÉcole Norm Supér Sér 2:12

Guckenheimer J (1975) Isochrons and phaseless sets. J Math Biol 1:259–273

Guckenheimer J, Holmes P (1983) Nonlinear oscillations, dynamical systems, and bifurcations of vector fields. Springer, Berlin

Guillamon A, Huguet G (2009) A computational and geometric approach to phase resetting curves and surfaces. SIAM J Appl Dyn Syst 8(3):1005–1042

Guillemin V, Pollack A (2010) Differential topology. AMS Chelsea Publishing, Providence (**Reprint of the 1974 original**)

Hoppensteadt FC, Izhikevich EM (1997) Weakly connected neural networks, vol 126. Springer, Berlin

Izhikevich EM (2007) Dynamical systems in neuroscience. MIT Press, Cambridge

Malkin IG (1949) Methods of poincaré and Liapunov in theory of nonlinear oscillations. Gostexizdat, Moscow

Malkin IG (1959) Some problems in the theory of nonlinear oscillations, vol 1. US Atomic Energy Commission, Technical Information Service

Pikovsky A (2015) Comment on "asymptotic phase for stochastic oscillators". Phys Rev Lett 115(6):069401

Revzen S (2009) Neuromechanical control architectures in arthropod locomotion. Ph.D. thesis, University of California, Berkeley. Department of Integrative Biology

Revzen S, Kvalheim M (2015) Data driven models of legged locomotion. In: Micro-and nanotechnology sensors, systems, and applications VII, vol 9467. SPIE, pp 315–322

Rudin W (1976) Principles of mathematical analysis, 3rd edn. McGraw Hill, New York

Sauer T, Yorke JA, Casdagli M (1991) Embedology. J Stat Phys 65(3–4):579–616

Schwabedal JTC, Pikovsky A (2013) Phase description of stochastic oscillations. Phys Rev Lett 110(20):204102

Spivak M (1971) Calculus on manifolds. Perseus Books Publishing, New York

Takajo H, Takahashi T (1988) Least-squares phase estimation from the phase difference. JOSA A 5(3):416–425

Takens F (1980) Detecting strange attractors in turbulence. Dynamical systems and turbulence, Warwick, vol 1981. Springer, Berlin, pp 366–381

Thomas PJ, Lindner B (2014) Asymptotic phase for stochastic oscillators. Phys Rev Lett 113(25):254101

Thomas PJ, Lindner B (2015) Thomas and Lindner reply. Phys Rev Lett 115(6):069402

Wilshin SD, Revzen S (2014) Phase driven models of unperturbed locomotion. Integr Comp Biol 54:e226

Wilshin S, Kvalheim MD, Scott C, Revzen S (2014) Estimating phase from observed trajectories using the temporal 1-form (in prep)

Wilson D, Ermentrout B (2018) An operational definition of phase characterizes the transient response of perturbed limit cycle oscillators. SIAM J Appl Dyn Syst 17(4):2516–2543

Wilson D, Moehlis J (2015) Determining individual phase response curves from aggregate population data. Phys Rev E 92(2):022902

Winfree AT (1980) The geometry of biological time. Springer, New York