A State Space Modeling Approach to Real-Time Phase Estimation

Authors: Anirudh Wodeyar[1], Mark Schatza[2], Alik S. Widge[2,3], Uri T. Eden[1], Mark A. Kramer[1]

1. Department of Mathematics and Statistics, Boston University
2. Department of Psychiatry, University of Minnesota
3. Department of Neuroscience, University of Minnesota

Abstract

Brain rhythms have been proposed to facilitate brain function, with an especially important role attributed to the phase of low frequency rhythms. Understanding the role of phase in neural function requires interventions that perturb neural activity at a target phase, necessitating estimation of phase in real-time. Current methods for real-time phase estimation rely on bandpass filtering, which assumes narrowband signals and couples the signal and noise in the phase estimate, adding noise to the phase and impairing detections of relationships between phase and behavior. To address this, we propose a state space phase estimator for real-time tracking of phase. By tracking the analytic signal as a latent state, this framework avoids the requirement of bandpass filtering, separately models the signal and the noise, accounts for rhythmic confounds, and provides credible intervals for the phase estimate. We demonstrate in simulations that the state space phase estimator outperforms current state-of-the-art real-time methods in the contexts of common confounds such as broadband rhythms, phase resets and co-occurring rhythms. Finally, we show applications of this approach to in vivo data. The method is available as a ready-to-use plug-in for the OpenEphys acquisition system, making it widely available for use in experiments.

Introduction

Rhythms, consistently periodic voltage fluctuations, are an ubiquitous phenomenon observed in brain electrophysiology across scales and species. Many studies have described a relationship between rhythms and behavior (Buzsáki, 2006; Haegens & Zion Golumbic, 2018; VanRullen, 2016). A prominent feature of rhythms proposed as relevant for neural processing is the phase. Phase has been proposed to coordinate neural spiking locally and potentially, through coherent networks, even globally (Fries, 2015; Maris et al., 2016). For example, the phase of slow rhythms (3-15 Hz) has shown relationships to dynamics in perception (Busch et al., 2009; Gaillard et al., 2020; Gregoriou et al., 2009; Helfrich et al., 2018). Phase-amplitude synchrony is common (Canolty & Knight, 2010; Hyafil et al., 2015), and has been suggested to

change in neurodegenerative disease (Hemptinne et al., 2013). Cross-regional, low frequency phase synchrony is consistently observed as a correlate of top-down executive control (Widge et al., 2019). Most work examining the importance of phase for neural dynamics and behavior has been correlative in nature. To better understand the functional relevance of the phase of a rhythm requires an ability to monitor and perturb phase relationships in real-time.

Several methods exist to estimate phase in real-time (Blackwood et al., 2018; L. L. Chen et al., 2013; Mansouri et al., 2017; Rivero & Ditterich, 2021; Rutishauser et al., 2013; Siegle & Wilson, 2014; Zrenner et al., 2020). At their core, these methods rely on bandpass filtering, a mechanism to forward predict data (using autoregressive models or linear extrapolation of phase) and the Hilbert transform to estimate the phase. Brain rhythms are often non-sinusoidal (Cole & Voytek, 2017) and broadband (Buzsaki, 2004; Roopun et al., 2008), making development of an accurate bandpass filter difficult. While the contemporary modeling approach using filters accurately estimates the phase across a range of contexts, several limitations exist that limit the phase estimate accuracy (Matsuda & Komaki, 2017; Siegle & Wilson, 2014): (i) By depending on bandpass filters, existing real-time phase estimators are susceptible to non-sinusoidal distortions of the waveform, and inappropriate filter choices may miss the center frequency or total bandwidth of the rhythm. (ii) Phase resets, moments when the phase slips as a result of stimulus presentation or spontaneous dynamics, cannot be tracked using filters, which control the maximum possible instantaneous frequency. (iii) By filtering the observed data directly, these approaches do not model signal and noise processes separately. Phase estimates thus represent signal and noise phase together. (iv) Many extant approaches depend on buffered processing (e.g., Fourier or related transforms), meaning that current phase estimates are delayed in ways that may not support real-time intervention. Further, buffer-based processing is susceptible to edge effects that are not of concern in traditional, offline analysis. (v) Contemporary approaches do not define a measure of confidence in the phase estimate. Despite these limitations, real-time phase simulations using current techniques have produced interesting results (Cagnan et al., 2017; Desideri et al., 2019; Hyman et al., 2003; Kundu et al., 2014; Schaworonkow et al., 2018; Siegle & Wilson, 2014; Zrenner et al., 2018), motivating the value of generating better methods to estimate phase causally.

To address the limitations of existing approaches, we implement here a dynamical systems approach to real-time phase estimation. Rather than filtering the data to identify the rhythm of interest, we focus on tracking the analytic signal (a complex valued entity, representing the amplitude and phase of a rhythm (Kramer & Eden, 2016)) as a latent process in the observed time series using a model proposed by (Matsuda & Komaki, 2017). In doing so, we remove the requirement of bandpass filtering, and instead model the signal as a state space harmonic oscillator (or combination of such oscillators) driven by noise. This modeling approach allows: (1) a separate estimate of observation noise, (2) a procedure to model and reduce the confounding activity of other rhythms and (3) a principled method for assessing uncertainty in the real-time phase estimates. After first fitting parameters of the state space harmonic oscillator model acausally, the method then tracks the phase in real-time assuming spectral stationarity. We call this approach the State Space Phase Estimate (SSPE).

To benchmark the SSPE technique against the current state-of-the-art methods of real-time phase estimation, we design several simulations and analyze case studies of *in vivo* data. We limit our comparison to the real-time methods proposed in Zrenner et al. (2020) and Blackwood et al. (2020), methods which provide an estimate of phase on receipt of each new sample of data and have readily available reference code. We additionally apply a standard method for non-real-time phase estimation (an acausal finite impulse response (FIR) filter followed by a Hilbert transform (Lepage et al., 2013)). We note that these (and other existing real-time phase estimation techniques (L. L. Chen et al., 2013; Mansouri et al., 2017; Rivero & Ditterich, 2021; Rutishauser et al., 2013; Siegle & Wilson, 2014)) apply a bandpass filter and forecasting technique, and, by that measure, the acausal FIR approach by utilizing more information suggests the room for improvement available for these estimators. We test the ability of these different methods to track the phase of rhythms with different spectral profiles - including narrowband rhythms, broadband rhythms, multiple rhythms, and rhythms with phase resets. We show that the SSPE method outperforms the existing real-time phase estimation methods in almost all cases. Finally, we illustrate the application of the SSPE method to *in vivo* LFP and EEG recordings, and introduce a plug-in of the SSPE method for OpenEphys (Siegle et al., 2017), making the method widely available for use in real-time experiments. Our work demonstrates that the SSPE method improves the ability to track phase accurately, in real-time, across a diverse set of contexts encountered in data.

Methods

State Space Model Framework

To estimate the phase from an observed signal, we implement a linear state space model framework. A state space model structure consists of two equations: a state equation, and an observation equation. The state equation captures the dynamics of a potentially hidden (or latent) signal we wish to track, while the observation equation maps the state to an observable (here, neural activity). Here, we implement the state equation in (Matsuda & Komaki, 2017) to track the real and imaginary part of the complex valued analytic signal of rhythms present in the neural data. For each rhythm modeled, the state equation dynamics rotate the analytic signal at that rhythm's frequency. The observation equation sums the real part of the analytic signals of every modeled rhythm to predict the observed neural activity. We treat as real-valued the real and imaginary parts of the complex valued analytic signal for ease of analysis and interpretation. This state space model structure allows us to derive the amplitude and phase of each modeled rhythm, and estimate the analytic signal in the temporal domain, thus avoiding complications associated with windowing (i.e., the need to choose large enough window sizes to support sufficient frequency resolution and the associated unavoidable temporal delay in phase estimates) that occur in related frequency domain approaches (Kim et al., 2018).

The state space model is linear in the state evolution (dynamics of state transitions) and observation equations. Further, the covariance for the state (which determines the bandwidth for

the rhythms) and noise for the observation is Gaussian (white noise). These assumptions (linearity and Gaussianity) imply that we can use the Kalman filter to correct the state estimate based on the true observation (i.e., to filter the state) (Kalman, 1960). This type of modeling approach has been successfully applied in many contexts to track latent structure from observed data while accounting for state and observation noise effects (Brockwell et al., 2004; Eden et al., 2018; Galka et al., 2004; Yousefi et al., 2019).

To model N rhythms, we define ***X*** as the state variable, a *2N x T* matrix over total time *T*. An individual time point with state X_t (*2N x 1* vector) is rotated according to each rhythm's (x_j^t, a *2 x 1* complex vector for rhythm *j*) central frequency ($\omega_j = 2\pi f_j \Delta t \in (0, \pi)$, where f_j is frequency in Hz and Δt is sampling rate) and damped by a constant (a_j). For each rhythm x_j^t we define θ_j^t as a scalar quantity tracking the phase, and **Θ** as a *N x T* matrix of phase values for all rhythms over all time. For the N rhythms, ***O*** is the rotation matrix, a *2N x 2N* and block-diagonal, defined by the central frequency of each rhythm. ***A*** is the scaling matrix , a diagonal *2N x 2N* matrix, with scaling factor $a_j \in (0,1)$to ensure stability, the same for each rhythm, and different across rhythms. The driving noise for the state U is Gaussian, a *2N x T* matrix with covariance structure **Q,** a *2N x 2N* diagonal matrix with the same variance for the real and imaginary parts of each rhythm. The state is collapsed into the observation *Y*(a *T x 1* vector) through M (a *1 x 2N vector*) which summates the real parts of the analytic signals of all rhythms. For any time point, the observation y_t is a scalar. The noise for the observation v_t is Gaussian with covariance σ_R^2. The model equations are:

$$X_t = [(x_t^1)^T, ..., (x_t^N)^T]$$
$$x_t^j = a_j O(\omega_j) x_{t-1}^j + u_t^j, u_t^j \sim N(0, Q_j) \tag{1}$$

$$\theta_j^t = arg(x_j^t) \tag{2}$$

$$O(\omega_j) = \begin{pmatrix} \cos(\omega_j) & -\sin(\omega_j) \\ \sin(\omega_j) & \cos(\omega_j) \end{pmatrix} \quad Q_j = \begin{pmatrix} \sigma_j^2 & 0 \\ 0 & \sigma_j^2 \end{pmatrix} \tag{3}$$

$$y_t = MX_t + v_t, v_t \sim N(0, \sigma_R^2) \tag{4}$$

$$M = [1, 0, 1, 0, ...]. \tag{5}$$

We note that, while we term ${x_j}^t$ as tracking the analytic signal, this nomenclature differs from the signal processing literature. In signal processing, the analytic signal is defined as a signal that has zero power at negative frequencies. However, given the stochastic frequency modulation permitted in the state vector ${x_j}^t$, we cannot guarantee that ${x_j}^t$ has zero power at negative frequencies. Nevertheless, since for sufficiently complex signals there are (potentially) several ways one could represent the dynamics of the amplitude and phase (Boashash, 1991;

Rosenblum et al., 1997), we use the term *analytic signal* here for ease of understanding our intent with this model and comparing it to classical Hilbert-based methods.

Real-time Phase Estimation

We perform real-time estimation of phase as follows. First, we use an existing interval of data to acausally fit the parameters of the state space model a_j, ω_j, Q_j, R using an expectation-maximization (EM) algorithm as proposed by (Soulat et al., 2019); for details see the Supplementary Information in (Soulat et al., 2019). In simulations, unless otherwise specified, we use the first 2 s of data to fit the model parameters. With these model parameters estimated and fixed, we then apply a Kalman filter to predict and update the state estimates, and estimate the phase and amplitude for each oscillator, each representing a different rhythm, for every sample (Figure 1).

We now define the stages of Kalman filtering. To predict the future state and state estimation error at time *t* prior to filtering (x_t^{t-1} and P_t^{t-1}) given the filtered state and estimation error (x_{t-1}^{t-1} and P_{t-1}^{t-1}) from the previous sample (*t-1*), we compute:

$$x_t^{t-1} = aO(\omega)x_{t-1}^{t-1} \tag{6}$$

$$P_t^{t-1} = aO(\omega_j)P_{t-1}^{t-1}(aO(\omega_j))^T + Q \tag{7}$$

We then estimate the Kalman gain for the current sample (*t*):

$$K_t = P_t^{t-1}M^T(MP_t^{t-1}M^T + R)^{-1} \tag{8}$$

Finally, we filter the current sample using the Kalman gain to compute filtered state (x_t^t) and state estimation error (P_t^t) at time *t*:

$$x_t^t = x_t^{t-1} + K_t(y_t - Mx_t^{t-1}) \tag{9}$$

$$P_t^t = P_t^{t-1} - K_tMP_t^{t-1} \tag{10}$$

The SSPE method also tracks the state estimation error (P_t^t), which allows estimation of confidence in the mean phase estimates. To do so, we sample 10,000 samples from the posterior distribution of each state *j*, and from this estimate the width of the 95% confidence bounds for the phase. Note that, alternatively, one could attempt to define a closed form for the phase distribution to estimate confidence bounds, as we estimate the posterior distribution for the analytic signal under the Kalman filter. However, this strategy requires an expectation of asymptotic convergence and the assumption of independence of amplitude and phase (Withers & Nadarajah, 2013).

AR-based Forecasting to Estimate Phase

We compare our approach to two existing algorithms for real-time phase estimation. Both methods estimate phase by first asynchronously estimating the parameters of an autoregressive (AR) model, then forecasting data from the current sample, before finally utilizing a filter and Hilbert transform to estimate phase (see Blackwood et al., 2018 and Zrenner et al., 2020). We implement the version of the Zrenner et al. (2020) algorithm accessible at https://github.com/bnplab/phastimate and term this method as *Zrenner*. Please refer to Zrenner et al. (2020) for algorithm details. Compared to the default settings, we set the algorithm to track a 6 Hz rhythm by increasing the window size for the filter to 750 ms, the filter order to 192, and the frequency band to 4 to 8 Hz. We utilize the implementation of the Blackwood et al. (2018) algorithm as provided by the authors. In Blackwood et al (2018), the authors develop a bandpass filter - AR-based forecasting and then the Hilbert transform - to compute the phase. Here, we instead apply a Hilbert transformer that estimates the phase by directly computing the Hilbert transform while filtering the data as implemented in (Blackwood, 2019). We do so because this implementation is more computationally efficient, functionally equivalent to the original method and is the code made widely available. We verified in sample data that application of the original and modified Blackwood methods produced consistent phase estimates. We refer to this method as *Blackwood*.

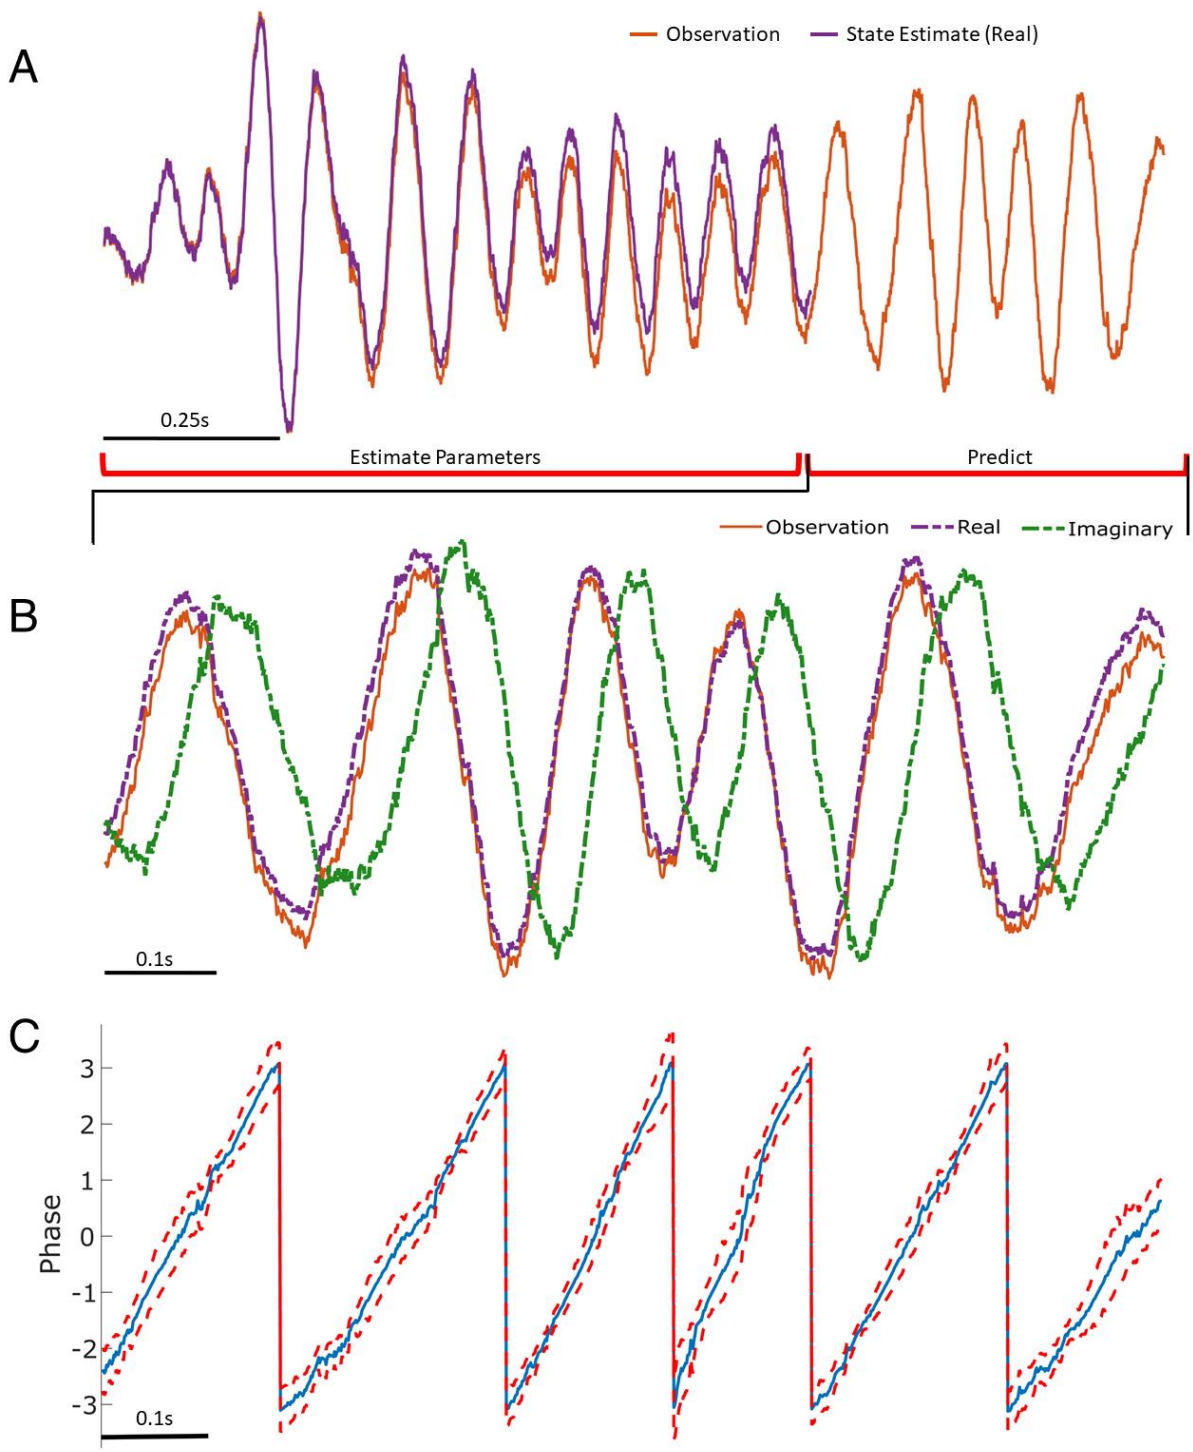

***Figure 1**: **Illustration of State Space Phase Estimate (SSPE) procedure.** Given an observation (**A**, red), we estimate model parameters in an initial interval of data. At subsequent times, we use the observed data to (**B**) predict and update the state estimate (purple), and (**C**) the phase and credible intervals, causally estimated. Note that prediction can be done for any time after parameter estimation; here we show a representative time interval.*

Non-Causal Phase Estimation

Existing studies compare causal, real-time phase estimation approaches to an acausal estimate of phase computed by applying an (acausal) FIR filter and the Hilbert transform (Blackwood et al., 2018; L. L. Chen et al., 2013; Zrenner et al., 2020). To maintain consistency with the existing literature, we do so here and apply a least-squares linear-phase FIR filter of order 750, with bandpass 4 to 8 Hz. After forward and backward filtering of the data (using MATLAB's *filtfilt* function), we apply the Hilbert transformer, a filtering based approximation of the Hilbert transform, to compute the analytic signal (using MATLAB's *hilbert* function). We estimate phase using the four-quadrant arctan function implemented by MATLAB's *angle* function and estimate the amplitude envelope using MATLAB's *abs* function. We use the *angle* function throughout the analysis to estimate the phase from a complex valued analytic signal. This acausal estimate of phase has been proposed to serve as a lower-bound in phase estimation error for methods that apply a filter based approach to real-time phase estimation (Zrenner et al., 2020).

Defining Error as Circular Standard Deviation

Identifying a common error metric across alternative algorithms to causally estimate phase is crucial for accurate comparison. We follow Zrenner et al. (2020) and use the circular standard deviation of the difference between the estimated phase and true phase:

$$CircSD = \sqrt{-2\log(|\exp\left(\frac{1}{n}\sum_{k=1}^{n}(i\theta - i\hat{\theta})\right)|)} \tag{11}$$

where n is number of samples, θ is true phase, $\hat{\theta}$is the phase estimate, and | | indicates the absolute value. We transform the final result from radians to degrees. The circular standard deviation (a transformed version of the phase locking value) captures the variance in the circular error for the phase estimate and ignores any potential bias (non-zero mean shift).

Simulated data

Narrow Band to Broad Band Rhythms

Our initial simulation tests the capability of different real-time phase estimators to track different frequency band profiles. We shift the bandwidth of the target rhythm from narrow (sharply peaked in the spectral domain at 6 Hz) to wide (broadly peaked in the spectral domain; peak bandwidth 4 - 8 Hz). We simulate all rhythms for 10 s at a sampling rate of 1000 Hz, with a central frequency of 6 Hz. For each simulated rhythm, we apply each phase estimation method, as described in Results.

The first type of rhythm we simulate is a pure sinusoid with added white noise:

$$Y = 10\cos(2\pi * 6 * T) + \epsilon_1$$
$$\epsilon_1 \sim N(0,1) \tag{12}$$

where *T* is time, and *N(0,1)* indicates a standard normal distribution. The second type of rhythm we simulate is a pure sinusoid with added pink noise, i.e., the power spectral density (PSD) of the noise decreases with the reciprocal of frequency (f):

$$Y = 10\cos(2\pi * 6 * T) + \epsilon_1$$
$$PSD(\epsilon_1(f)) \sim \frac{1}{f^{1.5}} \tag{13}$$

The third type of rhythm we simulate is filtered pink noise with added pink noise. To create a broadband spectral peak, we apply the same FIR filter utilized in the acausal phase estimation analysis (see *Methods: Non-Causal Phase Estimation*) to pink noise ($PSD \sim 1/f^{1.5}$), normalizing by the standard deviation and amplifying the resulting signal 10-fold. The resulting signal consists of a broad spectral peak centered at 6 Hz. To this broadband signal we again add independent pink noise (ϵ_1) as in Equation (13).

For the final rhythm we simulate a signal y_t from the SSPE model (see *Methods: State Space Model Framework*) which is non-sinusoidal and broadband. We fix one latent state (j=1) with frequency 6 Hz ($\omega_j = 6$), and set *a* (scaling factor) = *0.99*, variance *Q* = *10,* and observation noise variance (*R*) to 1. For this rhythm, phase was estimated from the state directly as described in *Methods: State Space Model Framework*. We show example instances of all four simulated rhythm types in Figure 2A.

Two Simultaneous Rhythms With Nearby Frequencies

In many physiologic use cases, we may wish to track one rhythm while ignoring another at a nearby frequency. To simulate this case, we consider two sinusoids. We define the target rhythm to track as a 6 Hz sinusoid with consistent amplitude. We define the confounding rhythm as a sinusoid with varying frequency and amplitude. We define the observation (signal plus noise) as:

$$Y = 25\cos(2\pi * 6 * T) + A * 25\cos(2\pi * F * T + \frac{\pi}{4}) + \epsilon_1 \tag{14}$$

$$\epsilon_1 \sim N(0, 0.5)$$

where $A \in [0.2,2]$ and $F \in [1,11]$. The confounding rhythm frequency assumes a range of integer values from 1 to 11 Hz. The confounding rhythm amplitude ranges from 0.2 to 2 times the amplitude of the target rhythm, in increments of 0.2.

Phase Reset

Ongoing slow rhythms in the brain occasionally undergo a sudden shift in phase. This can occur both in response to a stimulus or as a result of ongoing spontaneous neural dynamics (Voloh & Womelsdorf, 2016). To simulate sudden changes in a rhythm's phase, we implement a pure sinusoid (6 Hz) with additive pink noise and impose phase slips at four distinct time points, separated by at least 1 s. We fix the extent of phase slip, relative to the previous phase, at $\pi/2$. However, the phase at which the slip occurs differs across the four time points (see example in Figure 4, Panel A). We simulate the signal in five intervals, as follows:

$$\begin{gathered}
Y(0 < T < 3.5) = 10\cos(2\pi * 6 * T) + \epsilon_1 \\
Y(3.5 < T < 4.75) = 10\cos(2\pi * 6 * (T - 3.5) + \frac{\pi}{2}) + \epsilon_1 \\
Y(4.75 < T < 6.5) = 10\cos(2\pi * 6 * (T - 4.75)) + \epsilon_1 \\
Y(6.5 < T < 8.75) = 10\cos(2\pi * 6 * (T - 6.5) + \frac{\pi}{2}) + \epsilon_1 \\
Y(8.75 < T < 10) = 10\cos(2\pi * 6 * (T - 8.75)) + \epsilon_1 \\
PSD(\epsilon_1(f)) \sim \frac{1}{f^{1.5}}
\end{gathered} \tag{15}$$

In vivo Data

We consider two datasets to demonstrate how the SSPE method performs *in vivo*. The first dataset consists of local field potential (LFP) activity collected from the infralimbic cortex in rats. LFP were acquired continuously at 30 kHz (OpenEphys, Cambridge, MA, USA). An adaptor connected the recording head stage (RHD 2132, Intan Technologies LLC, Los Angeles, CA, USA) to two Mill-max male-male connectors (8 channels each, Part number: ED90267-ND, Digi-Key Electronics, Thief River Falls, MN, USA). For additional details on data collection, please refer to (Lo et al., 2020). We analyze here 250 s of LFP data from one channel. Before real-time phase estimation, we downsample the data to 1000 Hz (with a low pass filter at 100 Hz to prevent aliasing) for computational speed. To model these data, we fit three oscillators with the SSPE method, with the goal of tracking a target theta rhythm (defined as 4 - 11 Hz). Motivated by an initial spectral analysis of these data, we choose one oscillator to capture low

frequency delta band activity ($\omega_j = 1$Hz, based on spectral peak), one to capture the theta band rhythm of interest ($\omega_j = 7$Hz), and a final oscillator to track high frequency activity ($\omega_j = 40$Hz). We estimate these model parameters (so center frequencies may change) using the first 10 s of LFP data, and then fix these model parameters to estimate future phase values (up to 250 s). In doing so, we assume that the same oscillator models in the SSPE method remain appropriate for the entire duration of the LFP recording.

The second dataset was derived from the publicly available data in Zrenner et al. (2020). The data consist of spatially filtered and downsampled scalp electroencephalogram (EEG) recordings from electrode C3 (over central areas), which records a strong somatosensory mu rhythm (8-13 Hz). For details on data collection and preprocessing, see Zrenner et al. (2020). We analyze 250 s of data from a single participant for demonstration purposes. Before real-time phase estimation, we high-pass filter the data at 0.5 Hz to remove slow drifts (see *Discussion*), and low-pass filter the data at 100 Hz. We again track three oscillators in the SSPE method, with the goal of estimating phase from the mu rhythm: one in the delta band ($\omega_j = 2$Hz), one centered at the mu rhythm ($\omega_j = 10$Hz), and one in the beta band ($\omega_j = 22$Hz). We note that these frequency values initialize the model, and are then tracked in time. In analyzing the mu rhythm, we ensured that phase was estimated from an oscillator whose frequency was in the mu rhythm range, i.e., 8-12 Hz.

For the *in vivo* data, the true phase is unknown. To estimate phase error, we define the true phase as follows: first, we select a 10 s segment centered on the 1 s interval of the phase estimate. We then fit the State Space Model (Equations (1)-(4)) acausally to this 10 s segment, and estimate the phase in the 1 s interval at the segment's center. We note that here, the phase estimate is acausal because we use forward and backward smoothing of the state estimate when estimating the analytic signal. We also note that, in some segments, the frequency band of interest may not be identified through EM due to low signal strength. When this occurred, we omit error in the phase estimates, as these intervals lack the frequency of interest.

Real-time Implementation in TORTE

To make SSPE more broadly available, we implemented SSPE into the real-time TORTE (Toolkit for Oscillatory Real-time Tracking and Estimation) system (https://github.com/tne-lab/phase-calculator, 10.5281/zenodo.2633295). TORTE is built upon the Open Ephys platform (https://open-ephys.org/), and can be used with a broad range of electrophysiologic acquisition devices. Phase outputs and triggers from the system can be used to drive presentation of psychophysical stimuli or delivery of optical/electrical/magnetic stimulation. We directly integrated the MATLAB based SSPE algorithm described above into C++ for improved efficiency. TORTE uses a buffer based approach to receive data, based on the buffered approach of the underlying Open Ephys system. New buffers are repeatedly received by TORTE containing the most recent neural data. These buffers are subject to system latency causing a delay between activity in the brain being received by the toolkit. Both the buffer size and latency individually vary from microseconds to a few milliseconds depending on the acquisition system used. The phase of each sample within the buffer is estimated using

SSPE. A detector continuously monitors the estimated phase and can trigger output logic at any desired phase. We note that the observation noise parameter (*Q*) is defaulted to 50, but should be updated by the user to the square root of the amplitude. Two modifications were needed to integrate SSPE into TORTE. For computational efficiency, data were downsampled before application of the SSPE algorithm, and then the resulting phase estimates were upsampled back to the native sampling rate. The Open Ephys system can provide buffer sizes that are not exactly divisible into downsampled buffers. When this occurs, the phase estimates at the end of the upsampled buffer cannot be generated from the downsampled SSPE result. To account for this, an additional predicted state sample (with credible intervals) is computed by forward filtering using equations 6 and 7. This extra value is inserted at the end of the downsampled buffer. This provides an extra data point in the SSPE result that can be used to generate the remainder of samples for the upsampled phase output buffer. The second modification results from the real-time data acquisition by TORTE. In this case, low pass filtering must be causal, not the acausal (i.e., forward-backward filtering) used in the MATLAB implementation. To account for the consistent phase shift introduced in a real-time system, TORTE includes an online learning algorithm that adjusts target phase thresholding to improve event timings in closed loop applications. We test the real-time implementation on an AMD FX-8350 CPU with 16GB RAM.

Code availability

The code to perform this analysis is available for reuse and further development at (in MATLAB) https://github.com/wodeyara/stateSpacePhasePredictor and at (for OpenEphys) https://github.com/tne-lab/phase-calculator.

Results

We first examine in simulations the viability of different causal phase estimation algorithms, and how these existing methods compare to the proposed method - the state space phase estimator (SSPE). In these simulations, we vary the nature of the underlying rhythm of interest - either the spectral bandwidth or the noise/confounding signal present. Through these simulations, we examine the assumptions under which different methods perform best at estimating the phase in real-time. We then illustrate the application of the SSPE method to *in vivo* data in two case studies: a rodent local field potential, and a human electroencephalogram (EEG).

SSPE accurately tracks the phase of a single oscillation

To test the accuracy of causal estimates of phase, we simulate rhythms with different spectral profiles in four different ways. For the first two rhythms, we simulate a 6 Hz sinusoid and add either white noise (constant power spectrum, “Sine in White Noise”, Figure 2A.i) or pink noise (power spectrum decreases with the reciprocal of frequency, “Sine in Pink Noise”, Figure 2A.ii). For the next simulation scenario, to generate a broadband spectral peak, we filter pink noise into

the band of interest using an FIR filter (4-8 Hz, “Filtered Pink Noise”, see *Methods*), amplify this signal, and then add additional (unfiltered) pink noise (Figure 2A.iii). For the last simulation, to generate the signal of interest we sample data from the model underlying the SSPE method. The state space harmonic oscillator model (“State Space Model”) consists of a damped harmonic oscillator driven by noise (the signal) with additive white noise (Figure 2A.iv). We note that, for the first two simulation approaches in which the signal is a pure sinusoid, the true phase is known. For the “Filtered Pink Noise” simulation, we estimate the phase by applying the Hilbert transform to the signal before adding noise (see *Methods*). For the “State Space Model” simulation, we determine phase directly from the model (see *Methods*). Thus, in all simulation scenarios, we have an independent measure of the true phase.

We repeat each simulation 1000 times with different noise instantiations, and calculate error as the circular variance of the difference between the estimated phase and true phase (see *Methods*). We find that for rhythms with narrowband spectral peaks (Sine in White Noise and Sine in Pink Noise, Figure 2B,C), all estimators perform well, and those estimators that presume a narrow band signal (i.e., the filter based approaches) perform best. However, for rhythms with broad spectral peaks (Filtered Pink Noise and State Space Model, Figure 2B,C), the SSPE method outperforms the other causal approaches. Moreover, the SSPE method performs as well as (Filtered Pink Noise), or better than (State Space Model), the acausal FIR method. In what follows, we further explore the implications of this result, and the impact on assessment of phase estimator performance. We conclude that the SSPE method performs well in all four simulation scenarios, and outperforms existing causal phase estimation methods in two instances with broad spectral peaks.

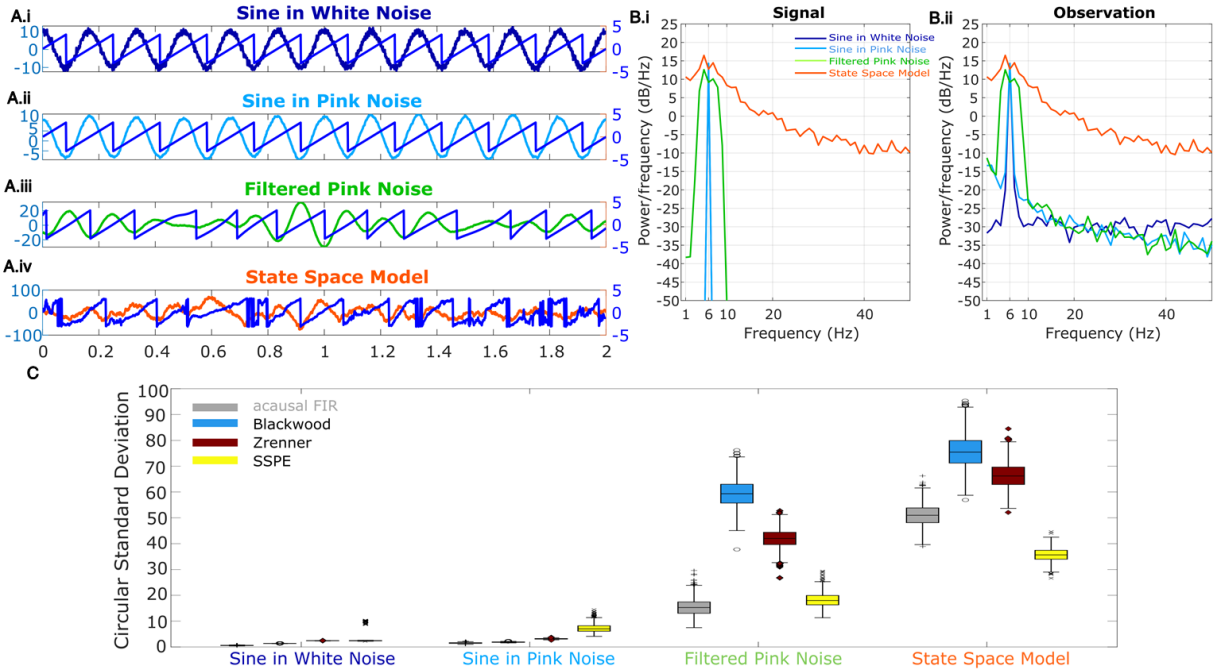

***Figure 2: Wide band rhythm's phase better tracked with SSPE**. **(A)** Example 2 s of simulated observed data (thick curves) and the true phase (thin blue curves) for each scenario. **(B)** For each scenario, example spectra of (B.i) the signal, and (B.ii) the observation (i.e., signal plus noise). **(C)** The phase error for each estimation method (see legend) and simulation scenario. In each box plot, the central mark indicates the median; the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively; the whiskers indicate the most extreme data points not considered outliers.*

SSPE accurately tracks the phase of two simultaneous oscillations.

To replicate a situation that has been documented in EEG (e.g., multiple alpha rhythms coexisting) (VanRullen, 2016) and in LFP (e.g., multiple theta/alpha rhythms in hippocampus) (Tort et al., 2010), we simulate two rhythms with similar frequencies. In this situation, phase estimation requires an approach to identify the presence of two rhythms and target the rhythm of interest. This simulation examines the impact of choosing a specific, previously specified, frequency response to isolate the rhythm of interest, as required by the two comparative causal phase estimation methods implemented here. In some cases, a finely tuned bandpass filter may separate the two rhythms, although such a filter requires sufficient prior knowledge and may be difficult to construct, particularly if the center frequencies are not perfectly stationary. The SSPE method provides an alternative approach to identify the two (or potentially more) rhythms without filtering and then allows tracking of the appropriate target rhythm. Further, while we focus on the phase of the target rhythm here, we note that the SSPE method provides estimates of the amplitude and phase of all modeled rhythms.

To illustrate this scenario, we simulate two slow rhythms at different frequencies. We set the target rhythm as a 6 Hz sinusoid with fixed amplitude, and the confounding rhythm as a 5 Hz sinusoid with 1.5 times the 6 Hz amplitude (Figure 3A). The observed signal consists of the combined target and confounding rhythms, plus white noise, simulated for 15 seconds. We apply each phase estimation method to the observed signal, and estimate the phase of the target rhythm after fitting parameters with 5 seconds of initial data. Visual inspection for this example (Figure 3B) suggests that only the SSPE method accurately tracks the target phase.

Repeating this analysis with the confounding amplitude and frequency varied across a range of values (frequencies 1-11 Hz, and amplitudes 20% to 200% of the target rhythm; see *Methods*) we find consistent results. For all methods, the error increases as the frequency of the confounding rhythm approaches the frequency of the target rhythm (6 Hz, Figure 3C). However, the increased error is restricted to a narrower frequency interval for the SSPE method. For the existing causal phase estimators, the error likely reflects choices in the filtering procedure (i.e, type of filter, features of passband and stopband, window size, see (Sameni & Seraj, 2017)) and modeling procedure (i.e., AR order, number of samples forecasted (L. L. Chen et al., 2013; Zrenner et al., 2020)). Similarly, for the acausal FIR estimator, the error likely reflects the filter passband of 4-8 Hz, with increasing error as the confounding rhythm nears the 6 Hz target rhythm.

For the existing causal and acausal methods, error also increases as the confounding oscillation amplitude increases (Figure 3C).

Previously, we showed that all estimation methods performed well for a signal consisting of a single sinusoid in white noise (see Figure 2, Sine in White Noise). Here, we show that for a signal consisting of two sinusoids in white noise, the SSPE method outperforms the existing methods. Unlike these existing methods, the SSPE method estimates parameters for both the target and confounding rhythms, and tracks each oscillation independently. By doing so, the SSPE method accounts for the confounding effect of the nearby oscillation, and allows more accurate tracking of the target rhythm. We conclude that the SSPE method outperforms these existing methods in the case of two simultaneous rhythms with nearby frequencies.

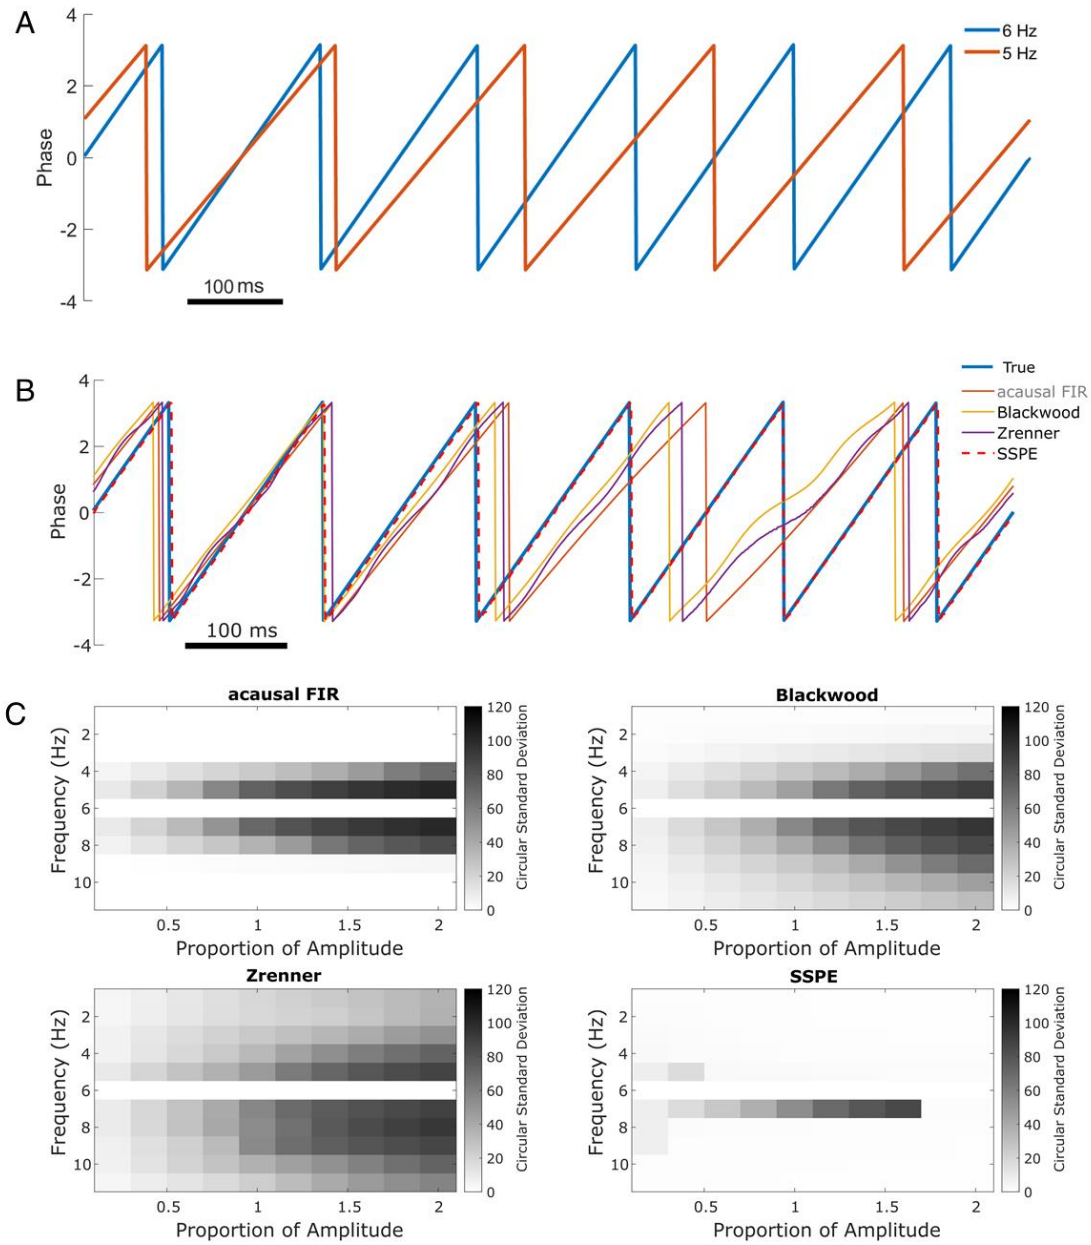

Figure 3. SSPE performs well when two oscillations exist at nearby frequencies. (A) *Example phase for two simulated rhythms, the confounding rhythm at 5 Hz and the target rhythm at 6 Hz.* ***(B)*** *Example phase estimates for the target rhythm in (A) using four different approaches (see legend).* ***(C)*** *Estimation error (circular standard deviation, see scale bars) for each estimation method as a function of frequency and amplitude of the confounding oscillation. In all cases the error increases as the frequency of the confounding oscillation approaches 6 Hz (the frequency of the target oscillation), and as the amplitude of the confounding oscillation increases.*

Phase estimation following phase reset

Phase resets are defined as consecutive samples of data that show a sudden shift in the oscillatory phase, proposed to represent the presence of synchronization of a neural mass in response to a stimulus, for example (Lakatos et al., 2009). Phase resets are common and well documented in electrophysiological recordings (Fiebelkorn et al., 2011; Makeig et al., 2004). Further, accurate tracking of phase resets would provide better evidence for the relevance of phase resetting in generating event related potentials (Sauseng et al., 2007).We examine the performance of each phase estimation method in the presence of simulated phase resets.

To do so, we simulate phase reset as a shift in the phase of a 6 Hz sinusoid by 90 degrees at four points in a 10 s signal (example in Figure 4A). We then compute the error of each phase estimation method in the 167 ms (one period of a 6 Hz rhythm) following the phase reset. Visual inspection of an example phase reset (Figure 4B) suggests that only the SSPE method accurately tracks the phase in the interval following the reset. We note that we fit SSPE parameters using an initial 2 s interval containing no phase resets (Eq. 15). Repeating this analysis for 1000 simulated signals (4000 total phase resets), we find consistent results; the error of the SSPE method is significantly lower than the error of the other methods ($p \sim 0$ in all cases using two-sided t-test, see Table 1).. These results demonstrate that the existing causal and acausal methods implemented here are unable to estimate the phase accurately immediately after a phase reset. We note that these existing methods all require a bandpass filtering step, which limits the fastest rate at which the phase can change (instantaneous frequency (Boashash, 1991)). The SSPE method, however, does not require the phase change within a limited range of instantaneous frequencies, and therefore can track rapid changes in phase. We conclude that the SSPE method is able to most rapidly track the true phase after a phase reset, and thus provides the most accurate estimate.

	SSPE	**Blackwood**	**Zrenner**	**acausal FIR**
Error (Std. Dev)	2.85 (0.89)*******	62.08 (0.38)	44.68 (0.38)	15.04 (0.23)

***** p~0 compared to all other methods**

Table 1: Error following phase reset for different phase estimation methods. *The SSPE method tracks phase resets with error near zero..*

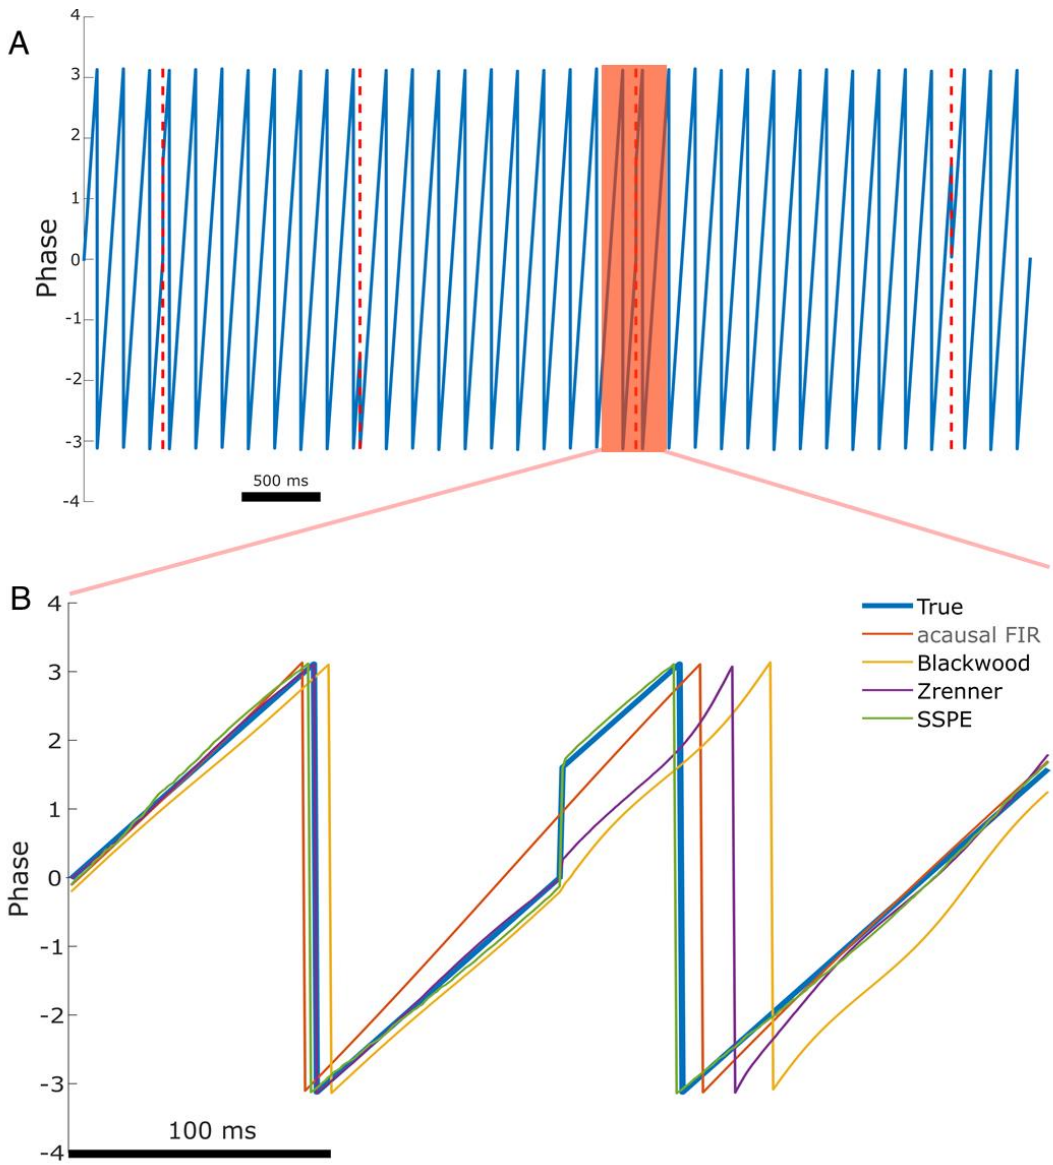

Figure 4. The SSPE accurately tracks phase following a phase reset. (A) *Example phase of a 6 Hz sinusoid (blue) with 4 phase resets (red dashed lines).* ***(B)*** *Example phase estimates (thin curves) and the true phase (thick blue curve) at the indicated phase reset in (A). The SSPE method (in green) tracks the true phase much more closely than other methods following a phase reset.*

Example *in vivo* application: rodent LFP

Having confirmed the SSPE method performs well in simulation, we now illustrate its performance in example *in vivo* recordings. We first apply the SSPE method to an example LFP

recording from rodent infralimbic cortex (see *Methods*), which contains frequent intervals of theta band (4-8 Hz) activity. Unlike the simulated data considered above, in the *in vivo* signal, the target rhythm changes in time (e.g., time intervals exist with, and without, the theta rhythm). Accurate application of the SSPE method requires accurate estimation of the parameters defining the oscillators in the SSPE model. In the simulated data, the stationary target rhythm allowed consistent SSPE model parameter estimation from (any) single interval in time; in those simulations, we estimated the SSPE model parameters from a single, early time interval (see *Methods*), and then used those (fixed) model parameters for future causal phase estimates.

For the *in vivo* data, in which the target rhythm may change in time, model estimates from an initial time may be inappropriate - and produce inaccurate phase estimates - at a later time. However, comparing the estimated phase to the true phase (see *Methods*), we find that the phase error tends to remain consistent in time (blue curve in Figure 5A; circular standard deviation 95 percent interval = [33.11, 75.39]). We find no evidence of a linear trend in error with time (linear fit, slope = 0.01, p = 0.53); i.e., we find no evidence that the error increases in time. Instead, visual inspection suggests that the error increases when the spectral content of the LFP changes (e.g., near 120 s in Figure 5B). We note that the error bounds are within the bounds of the range expected for real-time phase estimates in (Zrenner et al., 2020). Finally, we note that visual inspection of the power spectral density (Figure 5C) demonstrates that the theta band power appears more closely matched by a broad band, rather than a narrow band peak.

An alternative approach is to re-estimate the SSPE model parameters over time. By allowing the parameters of the model oscillators to change in time, this approach may improve the accuracy of the phase estimates. To examine this in the *in vivo* data, we refit the SSPE model parameters in each 10 s interval of the LFP data. More specifically, we refit model parameters in every 10 s interval of the 250 s recording, incrementing the 10 s interval by 1 s, to create 240 models. We then use each of these models to causally estimate the phase for all future times. For example, we use the SSPE model with parameters estimated from the interval 40-50 s to then estimate the phase for data after 50 s. We compute the error in the phase estimates for all 240 models to assess whether refitting the SSPE model parameters impacts future phase estimates. We find no evidence that the distribution of phase errors changes for the refit models compared to the original fixed parameter model for the immediate 30 seconds that followed (p=0.21, Wilcoxon Rank Sum test) or for the entire time period analyzed (red interval in Figure 5A, circular standard deviation 95 percent interval = [30.68, 88.7] degrees; p = 0.35 using Wilcoxon Rank Sum test). The original fixed parameter SSPE model (fit using the first 10 s of data; blue curve in Figure 5A) performs as well as the models refit with future data. We conclude that, under the assumption that parameters for the frequency band of interest are stable, we are able to estimate the SSPE model parameters whenever there exists a consistent rhythm of interest and then apply this model to estimate phase at future times (here on the scale of 100s of seconds). In other words, for these data, any 10 s interval with sufficiently strong theta rhythm can be used to fit parameters before the model is applied causally. This result has an important practical implication; the SSPE model parameters can be fit using an initial segment of data, and then applied to derive causal estimates of future phase values..In other words, the model has robust estimation

properties; parameters estimated early in the recording remain appropriate for future observations.

In addition, these results allow a novel insight into the rhythmic activity of the data, beyond those available using traditional methods. Typical analysis of neural rhythms either presumes the presence of a rhythm (for subsequent filtering) or operates in the frequency domain using the power spectral density and statistical testing to identify significant rhythms. Here, through direct fitting of the state space model to the time series data, we find consistent model parameters that accurately track the observed theta rhythm. This consistency is maintained through variations in the rhythm's amplitude. These observations suggest that the theta rhythm in infralimbic LFP maintains a consistent underlying rhythmic structure over the entire duration of the observation.

Confidence in the Phase Estimate

Past approaches to causally estimate phase have not attempted to quantify confidence in the phase estimate beyond utilizing the amplitude as a surrogate (Zrenner et al., 2020). However, using the amplitude assumes a monotonic linear relationship between the phase confidence bounds and the amplitude, which may not be true (e.g., for an AR(2) process with complex roots, Spyropoulos et al., 2019). Unlike existing methods, the SSPE method provides real-time estimates of the credible interval estimates for the phase (resulting from estimating the parameters of the SSPE model and tracking the phase using the Kalman filter; see *Methods*). To illustrate this, we show the phase and credible interval estimates of the SSPE method for an example 3 s interval of the LFP data (Figure 6). We note that, during a prominent theta rhythm, the credible intervals tightly track the mean phase estimate; at other times - when the theta rhythm is less obvious - the credible intervals expand. Previous studies have proposed a relationship between the theta amplitude envelope, estimated using the traditional approach of narrow band filtering followed by a Hilbert transform, and confidence in the phase (Zrenner et al., 2020); when the theta signal is strong (i.e., the amplitude is large), the theta phase estimate improves. To examine this relationship here, we plot the theta amplitude envelope versus the phase credible intervals derived from the SSPE method (Figure 6C.i). Visual inspection reveals that, when the amplitude is large, small decreases in amplitude coincide with increased credible intervals, as expected. However, if we choose to approximate certainty in the phase estimate through thresholding the amplitude, we find that fixed amplitude thresholds produce wide ranges of credible intervals (Figure 6C.ii). This is especially apparent at lower amplitude thresholds; choosing a threshold of 65% results in credible intervals that range from 0.04 (minimum) to 3.05 (maximum). Therefore, a fixed amplitude threshold does not guarantee an accurate phase estimate. Instead, we propose that inspection of the credible intervals better guides certainty in the phase estimate. These credible intervals are directly derived from the modeling procedure, and more interpretable than amplitude thresholds.

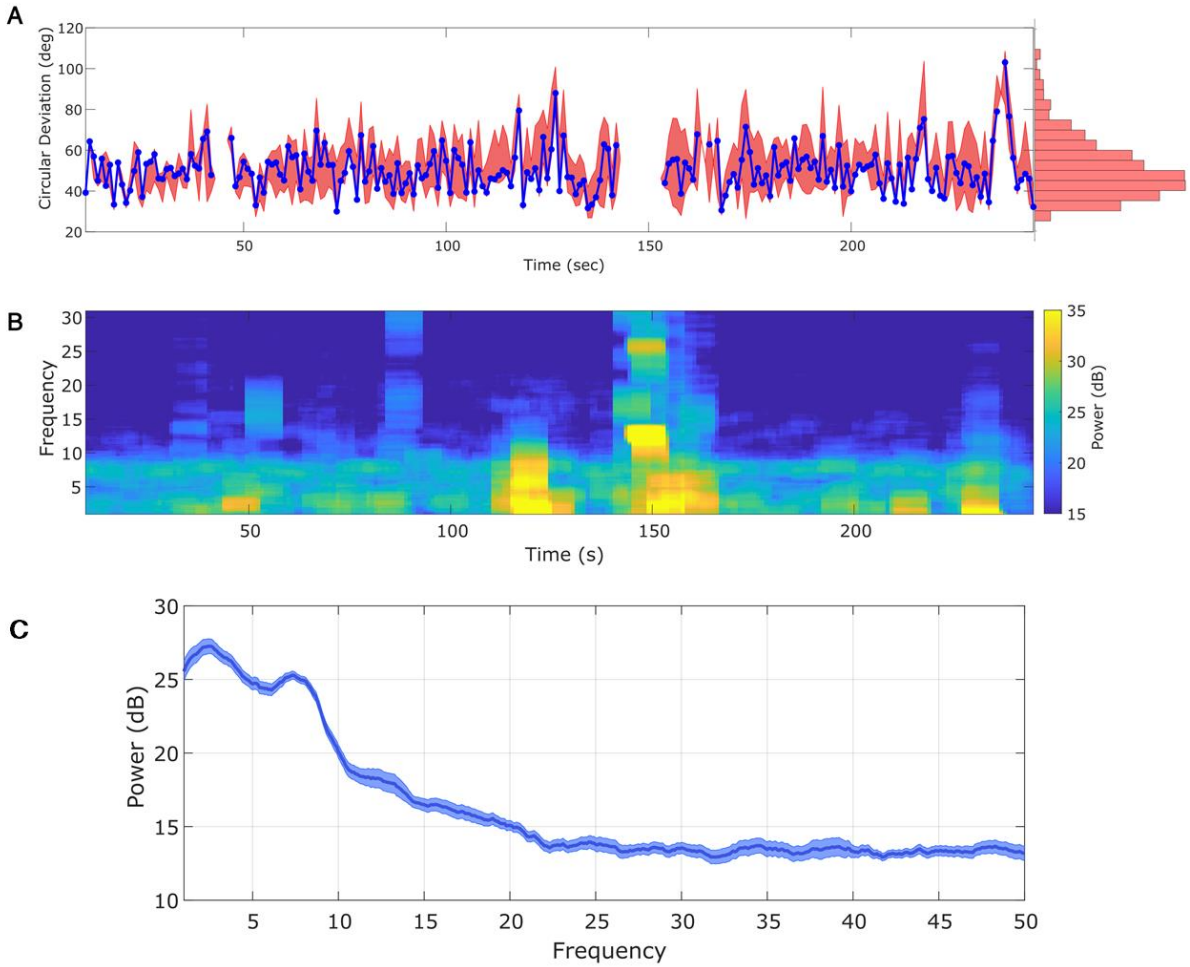

Figure 5: Any interval with a prominent rhythm of interest (here: theta) can be used to fit SSPE parameters. (A) *The phase error for a single instance of the model fit (using data at times 10-20 s, blue curve), and the 90 percent interval for phase error (red bands) at time t derived using parameter estimates from all models estimated with data prior to t. The moments where no error is reported are moments when an acausal approach failed to detect the theta band rhythm. Right: a histogram of circular deviation across all time demonstrating that error remains below 60 degrees in the majority of cases.* ***(B)*** *Corresponding spectrogram of the rodent LFP data (multi-taper method with window size 10 s, window overlap 9 s, frequency resolution 2 Hz, and 19 tapers).* ***(C)*** *Average power spectral density (PSD) of (B) over all times with theta band oscillator identified in (A). The mean PSD (dark blue) and 95% confidence intervals (light blue) reveal a broadband peak in the theta band.*

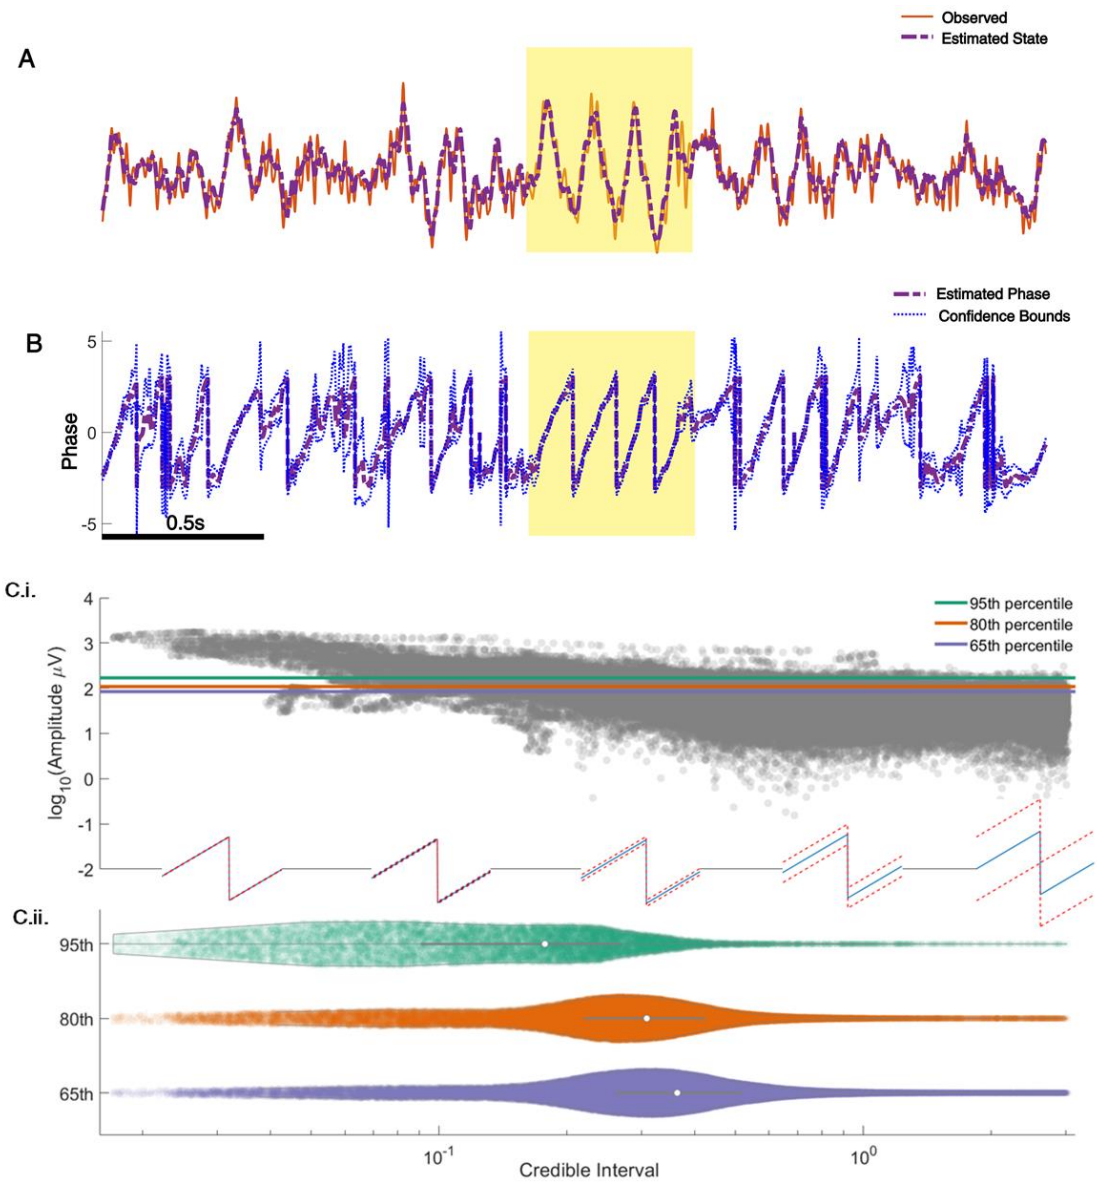

Figure 6. SSPE tracks in vivo credible intervals for the phase. (A) *Example rodent LFP data (red, solid) with a consistent broadband peak in the theta band (4-8 Hz). The estimated state of the SSPE method (purple, dashed) tracks the observed LFP.* ***(B)*** *Phase estimates (purple, dashed) and credible intervals (blue, dotted) for the example LFP data in (A). When the rhythm appears (yellow shaded region) the confidence bounds approach the mean phase; when the rhythm drops out the confidence bounds expand.* ***(C.i)*** *Phase credible intervals (versus theta rhythm amplitude for each time point (gray dot). On the x-axis we plot an example cycle of a rhythm with credible intervals, see x-axis of (C.ii) for numerical values of credible intervals.* ***(C.ii)*** *Violin plot of credible intervals for thresholds set at the 65th, 80th and 95th percentile of the distribution of amplitude. The distribution of credible intervals increases*

with reduced amplitude threshold, and all amplitude thresholds include times with large credible intervals.

Example *in-vivo* application: human EEG

To illustrate application of the SSPE method in another data modality, we analyze an example EEG recording of the mu rhythm (8-13 Hz) from (Zrenner et al., 2020). To do so, we apply the SSPE method to track three oscillators (1, 10, 22 Hz) in 10 s intervals (see *Methods*), and examine the phase estimates of the 8-12 Hz oscillator. Consistent with the rodent LFP data, we find that re-fitting of the SSPE model parameters is not required. Instead, the SSPE model parameters estimated from an initial 10 s segment of the EEG produce phase estimation errors consistent with SSPE models refit at future times (Figure 7A). Like the theta rhythm in the LFP rodent data, the mu rhythm in the human EEG data maintains a consistent spectral profile in time (Figure 7B) and appears to be broadband (Figure 7C). Therefore, estimating SSPE model parameters early in the data accurately captures the mu rhythm tracked later in the data, and there is no need to perform the (computationally expensive) refitting of the model.

We note that the phase estimates and credible intervals for the mu rhythm reveal how this rhythm waxes and wanes in time (example in Figure 7D,E). When a strong mu rhythm emerges, tight credible intervals surround the mean phase estimate tracking the rhythm of interest; as the mu rhythm wanes, the credible intervals expand. As with the LFP, we are able to assess certainty in our phase estimates using the credible intervals.

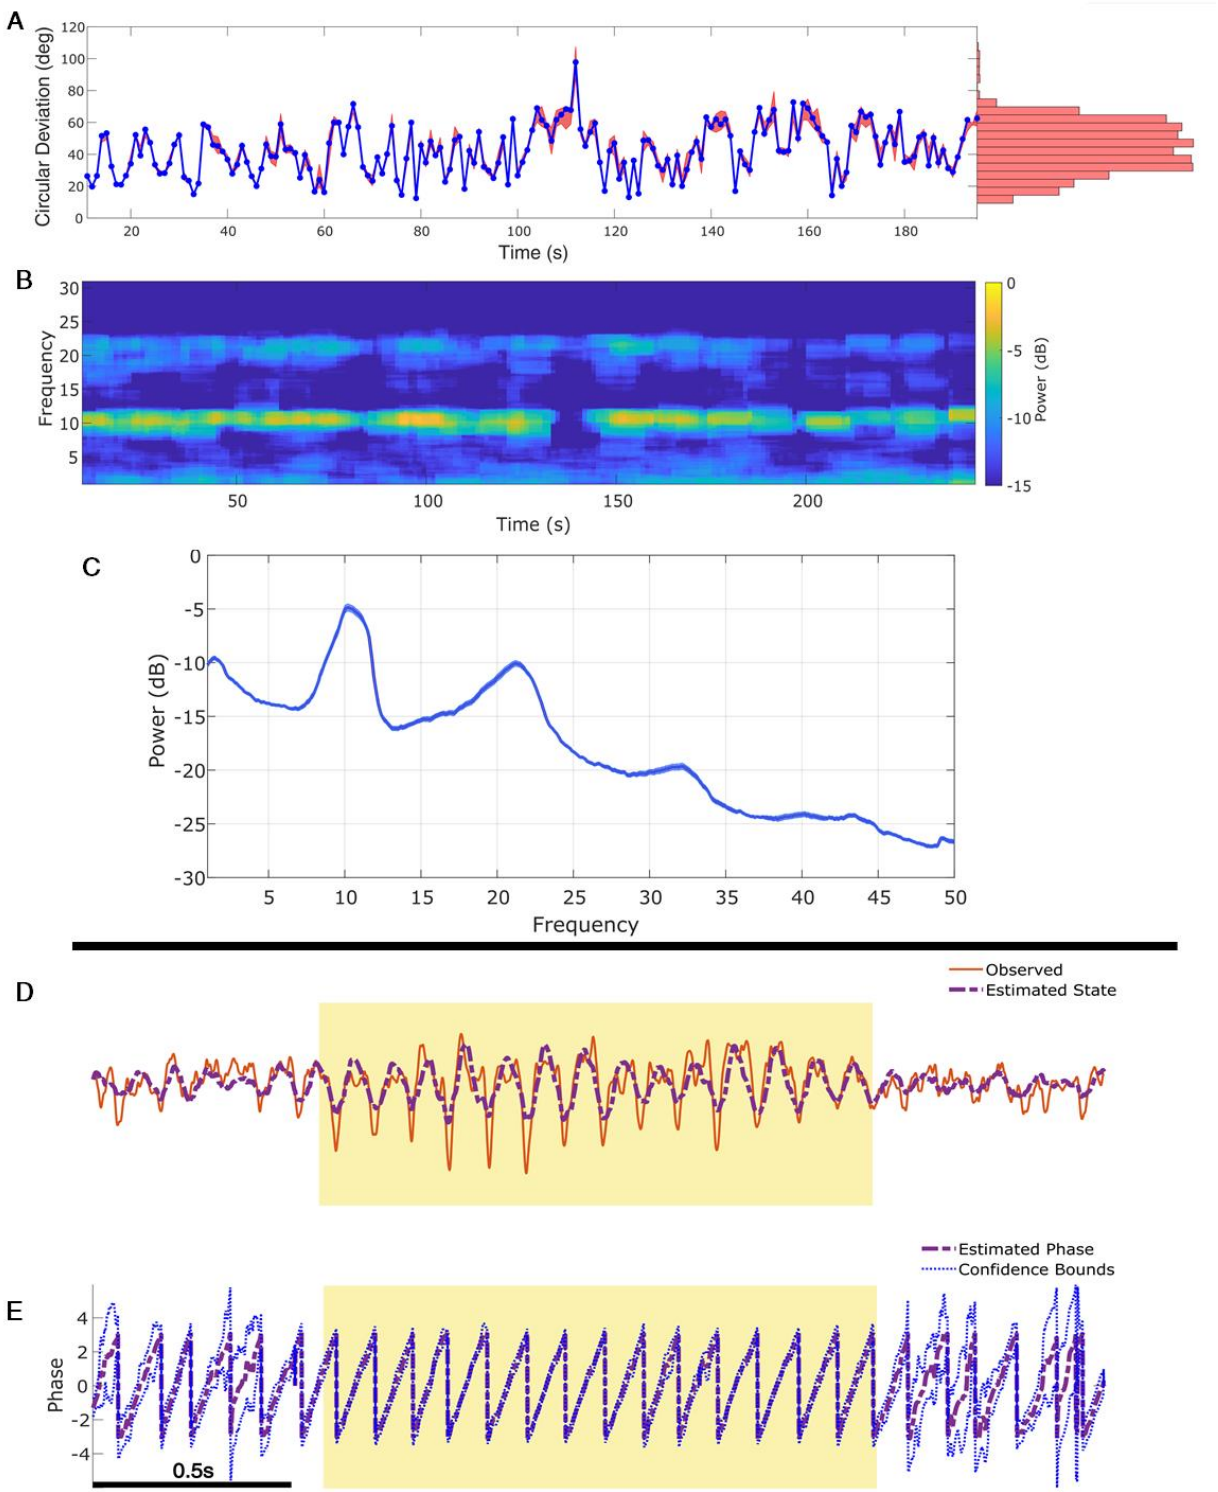

Figure 7:SSPE parameters can be estimated from any interval of EEG mu rhythm and credible intervals for mu rhythm phase indicate periods of confidence. (A) *Phase error for a model fit using data from the first 10 seconds (blue curve), and the 95% confidence*

intervals (red) in the error for phase estimates at time t using models fit at times prior to time t. Right: histogram of error showing that error remains below 60 degrees in most cases. ***(B)*** *Corresponding spectrogram for the EEG data (multi-taper method with window size 10 s, window overlap 9 s, frequency resolution 1 Hz, and 19 tapers).* ***(C)*** *Average power spectral density (PSD) of (B) over all times in (A). The mean PSD (dark blue) and 95% confidence intervals (light blue) reveal a broadband peak in the mu-rhythm near 11 Hz.* ***(D)*** *Example human EEG data (red, solid) with a consistent broadband peak in the mu rhythm (8-13 Hz). The estimated state of the SSPE method (purple, dashed) tracks the observed EEG.* ***(E)*** *Phase estimates (purple, dashed) and credible intervals (blue, dotted) for the example LFP data in C. When the rhythm appears (yellow shaded region) confidence bounds approach the mean phase; when the rhythm drops out the confidence bounds expand.*

Speed and Accuracy of TORTE Real-time Implementation

To assess the validity of the integration of SSPE into TORTE, we generated a 5 Hz sine wave with no additive noise and estimated the phase using the MATLAB and C++ implementations. The parameter estimates for the fitted model were computed on the first ten seconds of the data. The computed parameters were nearly identical between the two implementations and the time to estimate the parameters (which, as above, would need to be done only once per experiment) ranged from 2 s to 100 s. The phase estimates of the two implementations have a mean circular difference of 2.06 deg (Figure 8B) with a circular standard deviation of 0.65 deg; this variation in phase estimates likely arises from differences in filtering in the TORTE real-time (causal filtering on individual buffers) implementation and the offline MATLAB real-time (acausal filtering across all data) implementation. Further, given that it is a near constant shift in phase, the online learning algorithm in TORTE will correct for the biased phase estimates. We also examined the speed of evaluating different sized buffers with SSPE within TORTE. Across five buffer sizes from 420 to 1700 samples (encompassing 9 to 37 ms of data, see Figure 8C), estimating phase with SSPE took less than 0.35 ms in the worst case. We note that the system latency of TORTE is in the range of a few milliseconds (driven mainly by the USB-based implementation of the current Open Ephys system), whereas the buffers were evaluated on the scale of microseconds. That is, the real-time bottleneck lies in signal amplification and digitization, not this processing. In conclusion, we present a real-time implementation of SSPE that could be easily deployed in various experimental setups.

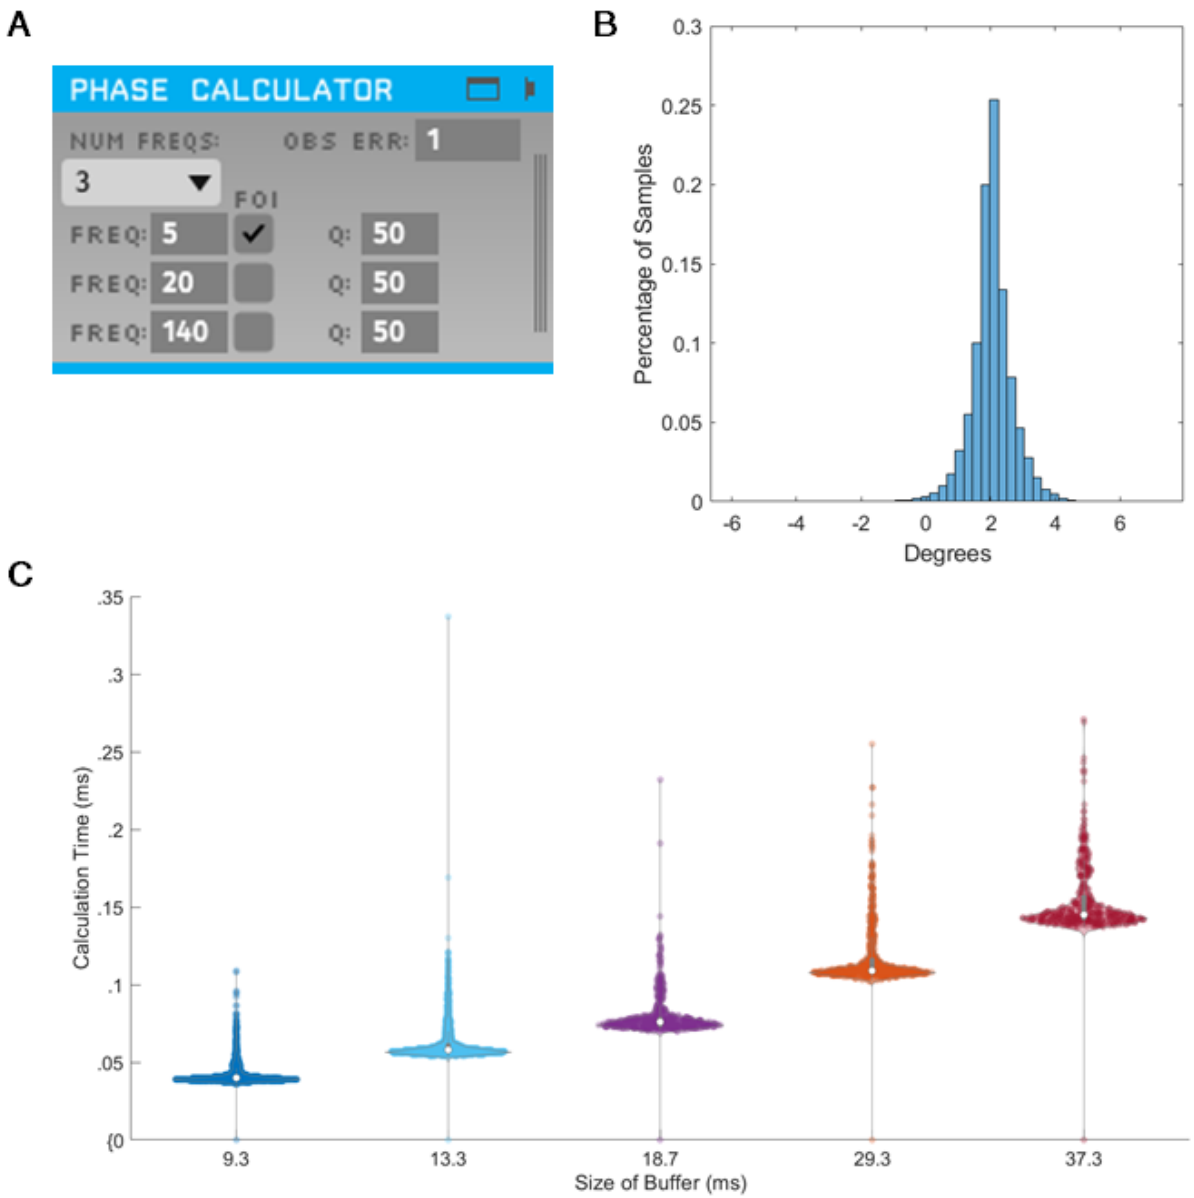

***Figure 8: SSPE implementation in TORTE is accurate, with negligible latency**. **(A)** OpenEphys GUI for using SSPE. The user specifies the number of frequencies to track, the center frequencies to track, the frequencies of interest for phase calculation and output (FOI), variance for the FOI and the observation error. Observation error (Q) determines the effective bandwidth for each frequency. **(B)** Histogram of the circular standard deviation between MATLAB (offline) and TORTE (real-time) implementations of the SSPE. Small variation results from causal low-pass filtering in TORTE and acausal filtering in the offline phase estimates. **(C)** Time to evaluate phase versus buffer size. Longer buffer sizes from TORTE require longer calculation time for application of SSPE. However, the calculation time is approximately two orders of magnitude smaller than the buffer size.*

Discussion

We have introduced a real-time phase estimator derived from a time series state space model. This approach - the SSPE method - addresses limitations of existing methods (namely, the requirements associated with bandpass filtering in phase estimation) and provides a real-time estimate of confidence in the phase estimate. We demonstrated the benefits of the SSPE method in simulations with varying breadth of the peak rhythm frequency, a confounding rhythm with a nearby frequency, and phase resets. In two case studies of *in vivo* data, we showed that re-estimation of model parameters is unnecessary and that the SSPE method provides clear estimates of confidence in phase estimates. We propose that the SSPE serves as a useful tool for real-time phase applications, e.g. real-time phase-locked stimulation or real-time phase-amplitude synchrony estimation.

Narrow-band vs Broadband Signals

While brain rhythms tend to be broadband (Buzsaki, 2004; Roopun et al., 2008), the most common methods in neuroscience to track a rhythm's phase generally expect a narrowband oscillation. However, real-time methods that require a narrowband oscillation struggle to capture the phase accurately for broadband rhythms (e.g., Figure 2). Moreover, when the broadband rhythm has non-negligible power across all frequencies, the standard non-real-time approach (acausal FIR and Hilbert transform) tracks the phase with limited accuracy. For forecasting needed in real-time analyses, implementing the same assumption (of a broadband peak) in the model may improve phase estimation accuracy. In the SSPE method, the assumption of a narrowband oscillation is not required, as the covariance allows flexibility in modelling peaks of different bandwidths.

Noise in the Signal

By applying a bandpass filter directly on the observed signal, existing real-time phase estimators do not disambiguate the signal and the noise, thus folding the noise into the phase estimate. The impact of structured noise (e.g., a competing rhythm or sudden phase shift) is particularly visible in Figure 3 and in Figure 4. To address this, the SSPE approach explicitly separates the signal and noise, and models two separate components of noise: state noise and observation noise. The SSPE state consists of harmonic oscillators driven by (white) noise. The resulting spectrum consists of broadband peaks at the oscillator frequencies upon a brown aperiodic component, similar to spectra observed in neural field activity. Confounding rhythms (a particular type of noise) can be explicitly modeled and accounted for in the SSPE approach. We consider the observation noise as additive white noise (flat spectrum) due to unknown noise in the measurement process. This model structure is more flexible than methods reliant on bandpass filters, allowing different expectations for the noise (e.g., modeling a confounding

rhythm as a pure harmonic oscillator) or adding additional model elements to account for power at higher frequencies (e.g., adding moving average terms c.f. (Matsuda & Komaki, 2017)).

Detecting a Rhythm for Phase Estimation

To estimate phase for a target rhythm, there should be evidence that the rhythm exists in the data. A few contemporary real-time phase estimators tackle this problem by using an amplitude threshold to decide when a rhythm is present (Rivero & Ditterich, 2021; Rutishauser et al., 2013). An amplitude criterion depends on accurate knowledge of the distribution of amplitudes possible for the data, relative to the background noise. The SSPE method, in contrast, directly represents the presence/absence of a rhythm in the credible intervals for the phase, as demonstrated for the *in vivo* data. The result is directly interpretable, and permits a more intuitive decision metric (e.g., a threshold depending on the credible interval width) than an arbitrary amplitude criterion. Further, while we expect a central frequency for a rhythm estimated from the data, the Gaussian assumptions of the harmonic oscillator driven by noise allow for the instantaneous frequency to be stochastic. This permits greater flexibility in tracking a rhythm's central frequency compared to prior phase estimators, and supports analysis of broadband rhythms with drifting central peaks, and abrupt changes in phase of a fixed rhythm (e.g., Figure 4). However, the SSPE method does require identification of target and confounding rhythms for accurate estimation. If details about the rhythms (central frequencies and bandwidth) are unknown, or the target rhythm is known to exist only for short intervals in a long recording, then alternative methods that allow real-time visualization and user-directed intervention in estimation may be pursued (Rivero & Ditterich, 2021; Rutishauser et al., 2013).

Computational Speed

For real-time phase applications, speed of estimation is critical. We focus here on algorithmic speed as opposed to hardware speed, which is a function of the computing device. Algorithmic speed can be measured at two different points: parameter estimation and application. In the SSPE method and other real-time phase estimators (L. L. Chen et al., 2013; Zrenner et al., 2020), the parameter estimation stage is the slower aspect of phase estimation. Past work has required 10s of minutes for parameter estimation (L. L. Chen et al., 2013). Here, for the SSPE method using a window size of 10 s and a sampling rate of 1 kHz, we find that parameter estimation time rarely exceeds a minute or two. In the application stage, we find that speed is orders of magnitude faster than the sampling rate (Figure 8), considerably improving on the algorithmic delays present in existing algorithms (Rivero & Ditterich, 2021; Rutishauser et al., 2013).

Limitations and Future Directions

Like all real-time phase estimators, the proposed method possesses specific limitations and opportunities for future improvement. First, the SSPE method treats the generative model for any rhythm as a harmonic oscillator driven by noise. A pure sinusoid is outside this model class, and as such the SSPE method performs slightly worse than filter-based algorithms at tracking the phase of a pure sine. Nonetheless, the error remains small for the SSPE tracking a pure sinusoid in noise (Figure 2). We note that a simple alteration in the SSPE model (removing the state variance Q) would allow more accurate tracking of a sinusoid, but not the broadband rhythms more commonly observed *in vivo*. Second, slow drifts (< 0.1 Hz) in the baseline of the data can impair the ability of the SSPE method to track the phase. Slow drifts could be modeled as an additional confounding oscillator, however, estimating its presence would require long segments of time (possibly exceeding 40 s). Alternatively, the SSPE model could be updated to include a model of the slow drift dynamics (perhaps as a piecewise linear term). Third, in applying the SSPE method, we assume that the model we estimate for a rhythm remains appropriate for the extent of the experimental duration. While we showed in two *in vivo* case studies that this assumption may be reasonable, in other cases this assumption may fail. Rhythms, observed over time, may be better represented by models with changing parameters. To address this, the SSPE method could implement a switching model that utilizes multiple sets of parameters and switches between parameter sets as necessary, or by refitting the SSPE model as time evolves. Finally, we note that data sampled at a high rate can reduce computational efficiency when fitting parameters for SSPE. To address this, a quasi-Newton optimization - instead of expectation maximization, as implemented here - may improve performance, as suggested in Matsuda and Komaki (2017). Determining the optimization framework best suited for practical use remains an important direction for future work.

Conclusion

We introduced a new algorithm - the SSPE method - to track phase in real-time that addresses several limitations of current methods. We demonstrated in simulation that SSPE performs better than current state-of-the art techniques for real-time phase estimation and comparably, or better, than a commonly used acausal FIR approach. We showed in *in vivo* data that we can estimate SSPE model parameters and apply them for hundreds of seconds without re-estimation. Finally, we showed how credible intervals from the SSPE method track confidence in the real-time phase estimate, and may provide an intuitive, user-determined decision metric for stimulation at a particular phase. We have provided reference implementations suitable for immediate incorporation into phase-aware *in vivo* experiments. Future work using the SSPE method in practice may help provide critical evidence on the potential functional importance of phase in neural dynamics.

References

Blackwood, E., Lo, M., & Widge, A. S. (2018). Continuous Phase Estimation for Phase-Locked Neural Stimulation Using an Autoregressive Model for Signal Prediction*. *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 4736–4739. https://doi.org/10.1109/EMBC.2018.8513232

Boashash, B. (1991). *Estimating and Interpreting The Instantaneous Frequency of a Signal-Part 1: Fundamentals*. 19.

Brockwell, A. E., Rojas, A. L., & Kass, R. E. (2004). Recursive bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology*, *91*(4), 1899–1907. https://doi.org/10.1152/jn.00438.2003

Busch, N. A., Dubois, J., & VanRullen, R. (2009). The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal of Neuroscience*, *29*(24), 7869–7876. https://doi.org/10.1523/JNEUROSCI.0113-09.2009

Buzsaki, G. (2004). Neuronal Oscillations in Cortical Networks. *Science*, *304*(5679), 1926–1929. https://doi.org/10.1126/science.1099745

Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195301069.001.0001

Cagnan, H., Pedrosa, D., Little, S., Pogosyan, A., Cheeran, B., Aziz, T., Green, A., Fitzgerald, J., Foltynie, T., Limousin, P., Zrinzo, L., Hariz, M., Friston, K. J., Denison, T., & Brown, P. (2017). Stimulating at the right time: Phase-specific deep brain stimulation. *Brain*, *140*(1), 132–145. https://doi.org/10.1093/brain/aww286

Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–515. https://doi.org/10.1016/j.tics.2010.09.001

Chen, L. L., Madhavan, R., Rapoport, B. I., & Anderson, W. S. (2013). Real-Time Brain Oscillation Detection and Phase-Locked Stimulation Using Autoregressive Spectral

Estimation and Time-Series Forward Prediction. *IEEE Transactions on Biomedical Engineering*, *60*(3), 753–762. https://doi.org/10.1109/TBME.2011.2109715

Cole, S. R., & Voytek, B. (2017). Brain Oscillations and the Importance of Waveform Shape. *Trends in Cognitive Sciences*, *21*(2), 137–149. https://doi.org/10.1016/j.tics.2016.12.008

Desideri, D., Zrenner, C., Ziemann, U., & Belardinelli, P. (2019). Phase of sensorimotor μ-oscillation modulates cortical responses to transcranial magnetic stimulation of the human motor cortex. *The Journal of Physiology*. https://doi.org/10.1113/JP278638

Eden, U. T., Frank, L. M., & Tao, L. (2018). Characterizing Complex, Multi-Scale Neural Phenomena Using State-Space Models. In Z. Chen & S. V. Sarma (Eds.), *Dynamic Neuroscience* (pp. 29–52). Springer International Publishing. https://doi.org/10.1007/978-3-319-71976-4_2

Ethan Blackwood. (2019). *tne-lab/phase-calculator: Phase calculator with fixed visualization and structs for channel info*. Zenodo. https://doi.org/10.5281/zenodo.2633295

Fiebelkorn, I. C., Foxe, J. J., Butler, J. S., Mercier, M. R., Snyder, A. C., & Molholm, S. (2011). Ready, Set, Reset: Stimulus-Locked Periodicity in Behavioral Performance Demonstrates the Consequences of Cross-Sensory Phase Reset. *Journal of Neuroscience*, *31*(27), 9971–9981. https://doi.org/10.1523/JNEUROSCI.1338-11.2011

Fries, P. (2015). Rhythms for Cognition: Communication through Coherence. *Neuron*, *88*(1), 220–235. https://doi.org/10.1016/j.neuron.2015.09.034

Gaillard, C., Ben Hadj Hassen, S., Di Bello, F., Bihan-Poudec, Y., VanRullen, R., & Ben Hamed, S. (2020). Prefrontal attentional saccades explore space rhythmically. *Nature Communications*, *11*(1), 925. https://doi.org/10.1038/s41467-020-14649-7

Galka, A., Yamashita, O., Ozaki, T., Biscay, R., & Valdés-Sosa, P. (2004). A solution to the dynamical inverse problem of EEG generation using spatiotemporal Kalman filtering. *NeuroImage*, *23*(2), 435–453. https://doi.org/10.1016/j.neuroimage.2004.02.022

Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009). High-frequency, long-range

coupling between prefrontal and visual cortex during attention. *Science (New York, N.Y.)*, *324*(5931), 1207–1210. https://doi.org/10.1126/science.1171402

Haegens, S., & Zion Golumbic, E. (2018). Rhythmic facilitation of sensory processing: A critical review. *Neuroscience & Biobehavioral Reviews*, *86*, 150–165. https://doi.org/10.1016/j.neubiorev.2017.12.002

Helfrich, R. F., Fiebelkorn, I. C., Szczepanski, S. M., Lin, J. J., Parvizi, J., Knight, R. T., & Kastner, S. (2018). Neural Mechanisms of Sustained Attention Are Rhythmic. *Neuron*, *99*(4), 854-865.e5. https://doi.org/10.1016/j.neuron.2018.07.032

Hemptinne, C. de, Ryapolova-Webb, E. S., Air, E. L., Garcia, P. A., Miller, K. J., Ojemann, J. G., Ostrem, J. L., Galifianakis, N. B., & Starr, P. A. (2013). Exaggerated phase–amplitude coupling in the primary motor cortex in Parkinson disease. *Proceedings of the National Academy of Sciences*, *110*(12), 4780–4785. https://doi.org/10.1073/pnas.1214546110

Hyafil, A., Giraud, A.-L., Fontolan, L., & Gutkin, B. (2015). Neural Cross-Frequency Coupling: Connecting Architectures, Mechanisms, and Functions. *Trends in Neurosciences*, *38*(11), 725–740. https://doi.org/10.1016/j.tins.2015.09.001

Hyman, J. M., Wyble, B. P., Goyal, V., Rossi, C. A., & Hasselmo, M. E. (2003). Stimulation in Hippocampal Region CA1 in Behaving Rats Yields Long-Term Potentiation when Delivered to the Peak of Theta and Long-Term Depression when Delivered to the Trough. *Journal of Neuroscience*, *23*(37), 11725–11731. https://doi.org/10.1523/JNEUROSCI.23-37-11725.2003

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, *82*(1), 35–45. https://doi.org/10.1115/1.3662552

Kim, S.-E., Behr, M. K., Ba, D., & Brown, E. N. (2018). State-space multitaper time-frequency analysis. *Proceedings of the National Academy of Sciences*, *115*(1), E5–E14. https://doi.org/10.1073/pnas.1702877115

Kramer, M. A., & Eden, U. T. (2016). *Case Studies in Neural Data Analysis: A Guide for the*

Practicing Neuroscientist. MIT Press.

Kundu, B., Johnson, J. S., & Postle, B. R. (2014). Prestimulation phase predicts the TMS-evoked response. *Journal of Neurophysiology*, *112*(8), 1885–1893. https://doi.org/10.1152/jn.00390.2013

Lakatos, P., O'Connell, M. N., Barczak, A., Mills, A., Javitt, D. C., & Schroeder, C. E. (2009). The Leading Sense: Supramodal Control of Neurophysiological Context by Attention. *Neuron*, *64*(3), 419–430. https://doi.org/10.1016/j.neuron.2009.10.014

Lepage, K. Q., Kramer, M. A., & Eden, U. T. (2013). Some Sampling Properties of Common Phase Estimators. *Neural Computation*, *25*(4), 901–921. https://doi.org/10.1162/NECO_a_00422

Lo, M.-C., Younk, R., & Widge, A. S. (2020). Paired Electrical Pulse Trains for Controlling Connectivity in Emotion-Related Brain Circuitry. *IEEE Transactions on Neural Systems and Rehabilitation Engineering: A Publication of the IEEE Engineering in Medicine and Biology Society*, *28*(12), 2721–2730. https://doi.org/10.1109/TNSRE.2020.3030714

Makeig, S., Debener, S., Onton, J., & Delorme, A. (2004). Mining event-related brain dynamics. *Trends in Cognitive Sciences*, *8*(5), 204–210. https://doi.org/10.1016/j.tics.2004.03.008

Mansouri, F., Dunlop, K., Giacobbe, P., Downar, J., & Zariffa, J. (2017). A Fast EEG Forecasting Algorithm for Phase-Locked Transcranial Electrical Stimulation of the Human Brain. *Frontiers in Neuroscience*, *11*. https://doi.org/10.3389/fnins.2017.00401

Maris, E., Fries, P., & van Ede, F. (2016). Diverse Phase Relations among Neuronal Rhythms and Their Potential Function. *Trends in Neurosciences*, *39*(2), 86–99. https://doi.org/10.1016/j.tins.2015.12.004

Matsuda, T., & Komaki, F. (2017). Time Series Decomposition into Oscillation Components and Phase Estimation. *Neural Computation*, *29*(2), 332–367. https://doi.org/10.1162/NECO_a_00916

Rivero, C. R., & Ditterich, J. (2021). A user-friendly algorithm for adaptive closed-loop phase-

locked stimulation. *Journal of Neuroscience Methods*, 13.

Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R. D., Kopell, N. J., & Whittington, M. A. (2008). Temporal interactions between cortical rhythms. *Frontiers in Neuroscience*, *2*. https://doi.org/10.3389/neuro.01.034.2008

Rosenblum, M. G., Pikovsky, A. S., & Kurths, J. (1997). Phase synchronization in driven and coupled chaotic oscillators. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, *44*(10), 874–881. https://doi.org/10.1109/81.633876

Rutishauser, U., Kotowicz, A., & Laurent, G. (2013). A method for closed-loop presentation of sensory stimuli conditional on the internal brain-state of awake animals. *Journal of Neuroscience Methods*, *215*(1), 139–155. https://doi.org/10.1016/j.jneumeth.2013.02.020

Sameni, R., & Seraj, E. (2017). A robust statistical framework for instantaneous electroencephalogram phase and frequency estimation and analysis. *Physiological Measurement*, *38*(12), 2141–2163. https://doi.org/10.1088/1361-6579/aa93a1

Sauseng, P., Klimesch, W., Gruber, W. R., Hanslmayr, S., Freunberger, R., & Doppelmayr, M. (2007). Are event-related potential components generated by phase resetting of brain oscillations? A critical discussion. *Neuroscience*, *146*(4), 1435–1444. https://doi.org/10.1016/j.neuroscience.2007.03.014

Schaworonkow, N., Caldana Gordon, P., Belardinelli, P., Ziemann, U., Bergmann, T. O., & Zrenner, C. (2018). μ-Rhythm Extracted With Personalized EEG Filters Correlates With Corticospinal Excitability in Real-Time Phase-Triggered EEG-TMS. *Frontiers in Neuroscience*, *12*. https://doi.org/10.3389/fnins.2018.00954

Siegle, J. H., López, A. C., Patel, Y. A., Abramov, K., Ohayon, S., & Voigts, J. (2017). Open Ephys: An open-source, plugin-based platform for multichannel electrophysiology. *Journal of Neural Engineering*, *14*(4), 045003. https://doi.org/10.1088/1741-2552/aa5eea

Siegle, J. H., & Wilson, M. A. (2014). Enhancement of encoding and retrieval functions through

theta phase-specific manipulation of hippocampus. *ELife*, *3*, e03061. https://doi.org/10.7554/eLife.03061

Soulat, H., Stephen, E. P., Beck, A. M., & Purdon, P. L. (2019). *State Space Methods for Phase Amplitude Coupling Analysis* [Preprint]. Neuroscience. https://doi.org/10.1101/772145

Spyropoulos, G., Dowdall, J. R., Schölvinck, M. L., Bosman, C. A., Lima, B., Peter, A., Onorato, I., Klon-Lipok, J., Roese, R., Neuenschwander, S., Singer, W., Vinck, M., & Fries, P. (2019). *Spontaneous variability in gamma dynamics described by a linear harmonic oscillator driven by noise* [Preprint]. Neuroscience. https://doi.org/10.1101/793729

Tort, A. B. L., Fontanini, A., Kramer, M. A., Jones-Lush, L. M., Kopell, N. J., & Katz, D. B. (2010). Cortical Networks Produce Three Distinct 7–12 Hz Rhythms during Single Sensory Responses in the Awake Rat. *The Journal of Neuroscience*, *30*(12), 4315–4324. https://doi.org/10.1523/JNEUROSCI.6051-09.2010

VanRullen, R. (2016). Perceptual Cycles. *Trends in Cognitive Sciences*, *20*(10), 723–735. https://doi.org/10.1016/j.tics.2016.07.006

Voloh, B., & Womelsdorf, T. (2016). A Role of Phase-Resetting in Coordinating Large Scale Neural Networks During Attention and Goal-Directed Behavior. *Frontiers in Systems Neuroscience*, *10*, 18. https://doi.org/10.3389/fnsys.2016.00018

Widge, A. S., Heilbronner, S. R., & Hayden, B. Y. (2019). Prefrontal cortex and cognitive control: New insights from human electrophysiology. *F1000Research*, *8*. https://doi.org/10.12688/f1000research.20044.1

Withers, C. S., & Nadarajah, S. (2013). The distribution of the amplitude and phase of the mean of a sample of complex random variables. *Journal of Multivariate Analysis*, *113*, 128–152. https://doi.org/10.1016/j.jmva.2012.05.017

Yousefi, A., Basu, I., Paulk, A. C., Peled, N., Eskandar, E. N., Dougherty, D. D., Cash, S. S., Widge, A. S., & Eden, U. T. (2019). Decoding Hidden Cognitive States From Behavior and Physiology Using a Bayesian Approach. *Neural Computation*, *31*(9), 1751–1788.

https://doi.org/10.1162/neco_a_01196

Zrenner, C., Desideri, D., Belardinelli, P., & Ziemann, U. (2018). Real-time EEG-defined excitability states determine efficacy of TMS-induced plasticity in human motor cortex. *Brain Stimulation*, *11*(2), 374–389. https://doi.org/10.1016/j.brs.2017.11.016

Zrenner, C., Dragana Galevska, Jaakko O. Nieminen, David Baur, MariaIoanna Stefanou, & Ulf Ziemann. (2020). *The shaky ground truth of real-time phase estimation*. 21.