A flexible likelihood approach for predicting neural spiking activity from oscillatory phase

Teryn D. Johnson[a], **Todd P. Coleman**[b], **Lara M. Rangel**[a,*]

[a]Department of Cognitive Science, University of California, San Diego, La Jolla, CA 92093, United States

[b]Department of Bioengineering, University of California, San Diego, La Jolla, CA 92093, United States

Abstract

Background: The synchronous ionic currents that give rise to neural oscillations have complex influences on neuronal spiking activity that are challenging to characterize.

New method: Here we present a method to estimate probabilistic relationships between neural spiking activity and the phase of field oscillations using a generalized linear model (GLM) with an overcomplete basis of circular functions. We first use an L1-regularized maximum likelihood procedure to select an active set of regressors from the overcomplete set and perform model fitting using standard maximum likelihood estimation. An information theoretic model selection procedure is then used to identify an optimal subset of regressors and associated coefficients that minimize overfitting. To assess goodness of fit, we apply the time-rescaling theorem and compare model predictions to original data using quantile-quantile plots.

Results: Spike-phase relationships in synthetic data were robustly characterized. When applied to *in vivo* hippocampal data from an awake behaving rat, our method captured a multimodal relationship between the spiking activity of a CA1 interneuron, a theta (5–10 Hz) rhythm, and a nested high gamma (65–135 Hz) rhythm.

Comparison with existing methods: Previous methods for characterizing spike-phase relationships are often only suitable for unimodal relationships, impose specific relationship shapes, or have limited ability to assess the accuracy or fit of their characterizations.

Conclusions: This method advances the way spike-phase relationships are visualized and quantified, and captures multimodal spike-phase relationships, including relationships with multiple nested rhythms. Overall, our method is a powerful tool for revealing a wide range of neural circuit interactions.

*Corresponding author at: Department of Cognitive Science, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0515, United States.

Keywords

Neural oscillations; Local field potential; Spike-phas relationship; Point process modeling; Generalized linear model; L1-regularization; Overcomplete basis; Sparsity; Hippocampus; CA1; Interneuron; Theta rhythm; Gamma rhythm

1. Introduction

Neural oscillations in extracellularly recorded signals arise from repeated, synchronous transmembrane currents across a population of neurons (Buzsáki and Draguhn, 2004; Buzsáki et al., 2012; Herreras, 2016; Jones, 2013; Pesaran et al., 2018). These currents have the ability to create temporal windows in which neurons are more or less depolarized. Consequently, the probability of observing spiking activity in neurons affected by these currents often stereotypically changes according to the phase of the oscillatory signal. However, the relationship between neural spiking activity and the phase of an oscillation is not always straightforward for a number of reasons. In particular, neurons often receive a convergence of different inputs that can have conflicting influences upon their depolarization. For example, a neuron that consistently receives a large number of synchronized excitatory inputs from one source might also receive behaviorally dependent inhibitory inputs from another source that affect the timing and probability of its spiking activity in a complex manner. In this scenario, the spiking activity of the neuron exhibits a dynamic relationship with the phase of the oscillation, which changes depending upon the behavior of the organism. Another source of complexity and unpredictability in determining spike-phase relationships arises from the fact that extracellularly recorded local field potential (LFP) and electroencephalography (EEG) reflect spatially averaged signals that can have both local and distant influences (Buzsáki et al., 2012; Herreras, 2016; Jones, 2013). This means that while some neurons exhibit spiking activity closely related to the phase of a particular oscillation, neighboring neurons might be more or less influenced by the currents generating the oscillation, or may not be affected at all. Thus, the relationship between neural spiking activity and oscillatory phase cannot be assumed *a priori*, and must be characterized experimentally.

By characterizing the relationship between neural spiking activity and the phase of a particular neural oscillation, we gain insight into the manner in which different neurons interact to support information processing. Since the oscillatory frequencies that manifest in a region can be the product of repeated interactions between neurons (Buzsáki and Draguhn, 2004; Kramer et al., 2008; Roopun et al., 2008; Sherman et al., 2016; Whittington et al., 2000), examining the strength of spike-phase relationships across populations of neurons can provide insight into the degree of their mutual engagement in rhythmic circuits. For example, certain gamma frequency oscillations arise from interactions between principal neurons and local inhibitory interneurons (Fries et al., 2007; Tiesinga and Sejnowski, 2009; Whittington et al., 2000). In this scenario, the frequency of the gamma oscillation is dictated by the decay time of inhibition, and both populations exhibit synchronized spiking activity tightly linked to the phase of the gamma oscillation. Identifying populations of principal cells and interneurons in a region that exhibits strong spike-phase relationships to the

same gamma frequency range can thus indicate that a similar interaction is occurring. In addition, strong spike-phase relationships across brain regions can suggest that the rhythmic input from one region influences the spiking activity of another, providing evidence for cross-regional interactions (Engel and Fries, 2010; Fries et al., 2007; Fries, 2009). Thus, characterizing spike-phase relationships can lead to the identification of local and cross-regional interactions that contribute to the dynamic information processing abilities of a region.

There are a number of methods currently implemented in neurophysiological studies to characterize spike-phase relationships. For spike-phase relationships that are thought to be unimodal, researchers often quantify the preferred phase of spiking and the strength of this preference by calculating the circular mean θ_{pref} and the mean resultant length vector $\bar{R}$ as follows (Berens et al., 2009; Lowet et al., 2016; Siapas et al., 2005):

$$\theta_{\text{pref}} = \arctan\left(\frac{\sum_{k=1}^{K} \sin(\theta_k)}{\sum_{k=1}^{K} \cos(\theta_k)}\right) \tag{1}$$

$$\bar{R} = \sqrt{\left(\frac{\sum_{k=1}^{K} \sin(\theta_k)}{K}\right)^2 + \left(\frac{\sum_{k=1}^{K} \cos(\theta_k)}{K}\right)^2} \tag{2}$$

where k indexes spike times, K is the total number of spikes, and θ_k is the phase of an oscillation at the time of each spike. Using the above expressions, it is possible to assess the degree to which spiking activity demonstrates a particular phase preference on average. This quantification method, and other methods that similarly quantify the consistency of spike-phase relationships through averages, can prove useful when assessing the degree to which unimodal phase preferences change across neurons or behavioral conditions (Siapas et al., 2005; Sirota et al., 2008; Siegel et al., 2009). In addition, researchers often implement standard coherence measures to gauge consistency of spike-phase relationships (Chalk et al., 2010; Fries et al., 2008; Pesaran et al., 2002, 2008). A spike spectrum $G_{spike}(f)$ at frequency f is created from a sequence of spike counts over time and compared to a field potential spectrum $G_{field}(f)$ using the following formula:

$$C_{\text{spike-field}}(f) = \frac{G_{\text{spike-field}}(f)}{\sqrt{G_{\text{spike}}(f)G_{\text{field}}(f)}} \tag{3}$$

where $G_{\text{spike-field}}(f)$ is the cross-spectrum and $C_{\text{spike-field}}(f)$ is the coherency between the two spectra. One major pitfall of these commonly used quantification methods, however, is that they are only applicable to unimodal spike-phase relationships. As such, they reduce potentially complex multimodal relationships (e.g. Skaggs et al., 1996) to single values for phase and strength. In addition, these methods do not statistically capture the inherent discrete (e.g. binary) temporal nature of neural spiking. Another concern is that the number of spikes in a data set and the length of the time series can strongly bias these measures (Lepage et al., 2011). Although several research groups have developed methods to quantify these influences and account for them when comparing across trials or conditions (Aoi et

al., 2015; Vinck et al., 2011, 2012), the uncertainty in the estimated relationships due to low spike counts or small data sizes is seldom fully quantified or reported.

Several promising methods for relating neural spiking activity to predictive variables, such as the phase of an oscillation, utilize parametric point process likelihood models (Barbieri et al., 2001; Czanner et al., 2008; Eden et al., 2018; Lewis et al., 2012; Truccolo et al., 2005). Rather than attempting to fully uncover the relationship between spiking and the phase of a neural oscillation, these methods often aim to predict neural spiking using a range of extrinsic covariates, including the phase of an oscillation. By performing model-fitting with point process maximum likelihood estimation, several useful and robust statistical methods may be implemented to both ensure goodness of fit and to quantify uncertainty. Models and methods that do not perform (or do not have the ability to perform) goodness of fit techniques are vulnerable to the imposition of inaccurate relationships onto the data. For example, a function (e.g. sine, polynomial, square) utilized by the model to fit the data imposes a shape that can be quite different from the true relationship (Fig. 1). In addition, when using a single function (and not a mixture of functions), complex relationships are forced into much simpler representations. To avoid these scenarios, the shape of the underlying relationship would need to be known *a priori* so that an appropriate function could be chosen. As described above, however, it is not always possible nor appropriate to assume a specific spike-phase relationship, given that neuronal engagement in larger rhythmic circuits is often unpredictable and dynamic.

In order to capture potentially complex spike-phase relationships without any *a priori* information regarding the underlying relationships, we propose the generation of a point process generalized linear model (GLM) using an overcomplete basis set of circular functions of phase. Specifically, we use Von Mises functions with a range of mean and inverse variance combinations. We propose a two-step process increasingly used in modern high-dimensional sparse statistical modeling (Belloni et al., 2013; Chételat et al., 2017). We first perform regressor selection with a penalized estimation procedure, and then perform maximum likelihood estimation with respect to the active set of regressors previously identified. Specifically, we select an *active set* of regressors from the overcomplete set by implementing an ℓ_1-penalized maximum likelihood estimation procedure and then selecting the regressors with nonzero coefficients. The ability to flexibly incorporate different numbers and combinations of regressors into the active set renders this workflow highly adaptable to any spike-phase relationship. In the next step, we perform standard maximum likelihood model-fitting with regressors from the active set. In order to obviate overfitting, we use model selection criteria, specifically the Akaike information criterion (Akaike, 1974), to identify the optimal active set.

In this paper, we demonstrate a robust ability to characterize a range of spike-phase relationships in both simulated and experimental data. We also present novel results from the rodent hippocampus, characterizing a multimodal interneuron relationship to both a theta (5–10 Hz) and a nested high gamma (65–135 Hz) rhythm. These results highlight a novel use of this method in identifying and quantifying the temporal coordination of spike coupling to multiple, co-occurring rhythms. Our method is thus a powerful tool for visualizing and characterizing neuronal engagement in rhythmic circuits.

2. Methods

2.1. Constructing an overcomplete basis of circular functions

Characterizing the relationship between spiking and neural oscillatory data first requires developing a method for relating two qualitatively different datasets. Neural spiking data is often represented as a point process, with action potentials occurring at discrete time points and irregular intervals. In contrast, phase data describing neural oscillations is continuous and circular, with regularly repeating values ranging from $-\pi$ to π radians. In order to characterize the conditional probability of observing a spike given the phase of a neural oscillation, we regress upon functions that are continuous on a circular plane. The Von Mises function, a circular variant of the normal function, of mean μ and inverse variance κ, is given by

$$V(\theta; \mu, \kappa) = \frac{e^{\kappa \cos(\theta - \mu)}}{2\pi I_0(\kappa)} \tag{4}$$

where I_0 is the Bessel function of order 0. We chose this function to form the overcomplete basis in this study because it is a well-defined function that is commonly used when representing circular datasets (Kim et al., 2013; Masseran et al., 2013; Mooney et al., 2003; Olson et al., 2017). While we have chosen to use Von Mises functions, alternative circular functions (e.g sine, cosine) can also be used to form the overcomplete basis. The phase θ is converted into a set of regressors by defining

$$V_{i,j}(\theta) = V(\theta; \mu_i, \kappa_j), i = 1, \ldots, I, j = 1, \ldots, J \tag{5}$$

where i indicates one of $I = 19$ possible means spanning from $-\pi$ and π in steps of 0.314 radians, and j indicates one of $J = 20$ possible inverse variance values ranging from 0.01 to 30 in steps of 1.5005. For reference, these inverse variance values translate to approximate widths at half-maximum spanning one half to 3/50th of a cycle. A total of 380 continuous functions were generated using (5) from each unique mean and variance pair, forming an overcomplete basis of functions through which spike-phase relationships could be represented.

2.2. Calculating the conditional probability of a spike given the phase of an oscillation

We discretize spiking activity into time bins matching the LFP sampling frequency (e.g. 1 ms), treating the resulting time series as a sequence of Bernoulli random variables which can take values 0 (no spike) or 1 (spike). We use a generalized linear model (GLM) with logit link function (canonical for the Bernoulli probability model) as a discrete time statistical model of neural spiking (Haslinger et al., 2010; Lewis et al., 2012). Specifically, at time bin t, we use the logit link function to associate spiking probability to oscillatory phase θ_t as follows:

$$logit(P(Y_t = 1 | \theta_t)) = \sum_{i=1}^{I} \sum_{j=1}^{J} x_{i,j} V_{i,j}(\theta_t) \tag{6}$$

where $logit(u) = \log\frac{u}{1-u}$, Y_t is a binary random variable at time t which is 1 when a spike occurs and 0 otherwise, and $x = [x_{i,j}: i=1, ..., I, j=1, ..., J]$ is a vector of length $I \times J$ pertaining to the weights applied to each base function $V_{i,j}$.

We can define the regressor vector of length $I \times J$ at time t as

$$R_t = \left[V_{i,j}(\theta_t): i = 1, ..., I, j = 1, ..., J\right] \tag{7}$$

and the inner product as

$$\langle R_t, x \rangle \triangleq \sum_{1=1}^{I} \sum_{j=1}^{J} V_{i,j}(\theta_t)x_{i,j} \tag{8}$$

so that (6) can be rewritten succinctly as

$$P(Y_t = 1 | \theta_t) = \frac{e^{\langle R_t, x \rangle}}{1 + e^{\langle R_t, x \rangle}}. \tag{9}$$

The likelihood of a model is a quantitative measure for determining the plausibility of a model given observed data. Maximizing the likelihood is a well-known procedure for determining x values that optimize the model fit (Poor, 2013). Since $-\log(u)$ is a monotonically decreasing function of u, maximizing the likelihood is equivalent to minimizing the normalized negative log likelihood, given by:

$$l(x) = \frac{1}{N} \left[\sum_{t=1}^{N} - y_t \cdot \langle R_t, x \rangle + \log(1 + e^{\langle R_t, x \rangle}) \right]. \tag{10}$$

Although directly minimizing $l(x)$ provides an estimated weight, $x_{i,j}$ for each basis function $V_{i,j}$, this would lead to overfitting in our setting given the use of this overcomplete basis. As such, the resultant model would be less generalizable, and less likely to represent the underlying relationship because of poor bias-variance tradeoffs (Poor, 2013) that arise when performing maximum likelihood with large numbers of parameters.

2.3. Active regressor selection with an ℓ_1-regularized generalized linear model

In this section, we first discuss challenges in selecting functions for models of different complexity, and then propose a solution for optimizing both the number and identity of basis functions used in the model to represent the data. While some basis functions, such as Zernike polynomials, have a natural ordering based upon increasing complexity and orthogonality (Barbieri et al., 2002; Coleman and Sarma, 2010; Lakshminarayanan and Fleck, 2011), there is no natural ordering for basis functions such as Von Mises functions. One possible method for increasing the number of Von Mises functions in the model adopts a stepwise approach. In this approach, an ordering is established by systematically adding one additional function to the model at a time. Each new function optimizes the fit given the functions already within the model. This stepwise approach increases model complexity in an ordered fashion, but can cause inefficient selection of function identity and possible

inflexibility in the shape of the model. For example, the first Von Mises function selected to characterize a spike-phase relationship with two neighboring peaks may be a function that spans both peaks due to their proximity. Additional functions added using a stepwise approach must improve upon an existing poor fit, rather than allowing functions in the model to flexibly change as complexity is increased. To solve this problem, we propose the use of a term proportional to the ℓ_1 norm of x,

$$\|x\|_1 = \sum_{i\,=\,1}^{I} \sum_{j\,=\,1}^{J} |x_{i,\,j}|, \tag{11}$$

to impose sparsity in the number of regressors used to characterize the data:

$$\min_x \{ l(x) + \lambda_q \|x\|_1 \} \tag{12}$$

where λ_q is a non-negative regularizer coefficient that controls the magnitude of regularization and thus the sparsity in the solution. We define the *active set* $S[q]$ to be the set of (i,j) pairs for which the weights associated with the solution of (12) are nonzero. By systematically varying the magnitude of the regularization λ_q as a function of index q, the active set can be varied, allowing greater or fewer numbers of regressors to be used in the model. Unlike adding an additional regressor to the model in a stepwise manner, this method allows the regressors within the active set to change as λ_q changes, making this method more adaptable than a stepwise method. Thus, ℓ_1-regularization provides a method for both increasing the model complexity and selecting the most appropriate set of Von Mises basis functions for any given model order. Note that while the x is a term within (12), the purpose of this step is to identify the active set $\mathcal{S}[q]$ associated with λ_q, and *not* to identify the optimal weights. This is in part because of a well-known drawback of ℓ_1-regularized estimators is the systematic shrinkage of the large coefficients towards zero (Belloni et al., 2013).

2.4. Model fitting

By defining the inner product with respect to the active set

$$\langle R_t, x \rangle_{\mathcal{S}} = \sum_{(i,\,j)\,\in\,\mathcal{S}} V_{i,\,j}(p) x_{i,\,j}, \tag{13}$$

we can define the likelihood of x with respect to $\mathcal{S}$ as the likelihood in (10), but only operating on the (i,j) components of x and R_t within $\mathcal{S}$:

$$l(x; \mathcal{S}) = \frac{1}{N}\left[\sum_{t\,=\,1}^{N} - y_t \cdot \langle R_t, x \rangle_{\mathcal{S}} + \log(1 + e^{\langle R_t, x \rangle}_{\mathcal{S}}) \right]. \tag{14}$$

Thus, after finding the active set $\mathcal{S}[q]$, weights can be calculated for each Von Mises function using standard maximum likelihood procedures operating upon regressors in the active set:

$$x[q] = \underset{x}{\operatorname{argmin}} \; l(x; \mathcal{S}[q]) \tag{15}$$

where it is assumed that for any (i, j) pair that is not in $\mathcal{S}[q]$, the coefficient $x[q]_{i,j} = 0$. Now that the Von Mises functions have been ordered using a method that best allows them to represent the data, another method is required to select the active set that optimally describes the data while minimizing overfitting.

2.5. Model selection

To minimize overfitting, we implement a model selection procedure that utilizes the Akaike information criterion (AIC) (Akaike, 1974). This method applies an additive penalty term to l according to the number of degrees of freedom for the qth model, which in the case of ℓ_1-regularized problems pertains to the size of the active set (Tibshirani and Taylor, 2012; Zou et al., 2007):

$$A_q = l(x[q]; \mathcal{S}[q]) + \frac{d_q}{N}, \tag{16}$$

where d_q is the number of elements in $\mathcal{S}[q]$. We then select q^* as a local or global minimum of A_q and select

$$\mathcal{S}^* = \mathcal{S}[q^*], x^* = x[q^*]. \tag{17}$$

As such, $\mathcal{S}^*$ is the active set that optimally characterizes the data while minimizing overfitting. This approach thus provides an unbiased method for determining the optimal number of regressors and associated weights to use with the model.

2.6. Calculating the final conditional probability using the selected model

Following the selection of the optimal active set, the optimal conditional probability of a spike given the phase of an oscillation is calculated given the following equation:

$$P(Y_t = 1 | \theta_t) = \frac{e_{\mathcal{S}^*}^{\langle R_t, x^* \rangle}}{1 + e_{\mathcal{S}^*}^{\langle R_t, x^* \rangle}} \tag{18}$$

where $\mathcal{S}^*$ and x^* are the optimal active set and optimal weights, respectively, as described in (17).

2.7. Goodness of fit

To evaluate whether or not a proposed statistical model of neural spiking appropriately characterizes the spike train time series, a goodness-of-fit measure is needed to quantitatively assess the agreement between the model and neural spiking activity. The time-rescaling theorem (Brown et al., 2002) indicates that with an accurate model, a transformation of inter-spike-intervals (ISIs) can be performed to construct a new spike train with independent, identically distributed ISIs. Standard goodness-of-fit with quantile-quantile (Q-Q) plots can be applied to these ISIs. Altogether, performing this procedure allows for an assessment of whether the statistical model indeed characterizes the statistical nature of the measured spike train (Brown et al., 2002). Confidence intervals (e.g. 95%)

alongside the Q-Q plot can be used to assess the extent to which the model accurately represents the data. In addition, a χ^2 test can be performed in order to assess the extent to which the estimated conditional probability deviates from a theoretically equiprobable distribution, with $p < 0.05$ used as a significance threshold. The entire workflow is illustrated in Fig. 2.

2.8. Simulated neural data

To simulate neural oscillations in the local field potential, we generated 60 s of continuous, circular phase data that repeatedly spanned $-\pi$ to π radians in uniform steps. To emulate theta oscillations typical of the rodent hippocampus, the phase data completed a cycle every 125 data points, corresponding to an 8 Hz oscillation recorded at a sampling rate of 1000 Hz.

In order to evaluate the effectiveness of the model in characterizing different spike-phase relationships, simulated unimodal and multimodal relationships were created. To generate these relationships, the same Von Mises distributions used to form the overcomplete basis described in Section 2.1 were chosen to create ground truth probability values obeying the following relationship:

$$\hat{P}(Y_t = 1 \mid \theta_t) = \langle R_t, w \rangle \tag{19}$$

where $\hat{P}$ is the simulated conditional probability of observing a spike at time t for a specific phase θ_t, R_t is the vector of Von Mises functions evaluated at θ_t, given by (7), and w is a length $I \times J$ vector of coefficients. Note that this method of generating the simulated data does not contain the logit function and thus differs from (6) used in the model. Since the simulated dataset falls outside the class of parametric statistical models for which we will perform inference, this allows us to test the ability to approximately recover spike-phase statistical relationships that lie outside the model class.

To create unimodal spike-phase relationships, an (i, j) pair was chosen at random to select a single Von Mises distribution $f_{i,j}$. This process was repeated for every cycle of the phase data set. The associated w vector contained all zeros except in the (i, j) position, where w_{ij} was randomly selected to scale the Von Mises function while ensuring that $\langle R_t\ w \rangle$, at each t remained between 0 and 1. At each time point t, spikes were generated according to the probability dictated by (19). To generate multimodal spike-phase relationships, multiple (i, j) pairs were chosen at random, and each coefficient w_{ij} was scaled to ensure that $\langle R_t\ w \rangle$, at each time point remained between 0 and 1.

2.9. Identification of data limitations

In order to simulate similar spike-phase relationships from neurons exhibiting a range of different firing rates, w_{ij} was adjusted to create peak probability values ranging from 0.001 to 0.9 for the unimodal condition. In a data set in which phase data completes a cycle every 50 data points (corresponding to a 20 Hz oscillation recorded at a sampling rate of 1000 Hz) these peak probability values correspond to a range of firing rates from approximately 0.3 Hz to 314.2 Hz. This method substantially adjusts the probability of observing a spike

on any given cycle, while maintaining the degree of phase modulation within a cycle. To increase the size of the dataset, these cycles were repeated for larger or smaller intervals depending on the desired data length. Spikes were generated for data sets ranging from 1 to 300 s.

Since the workflow requires at least one spike to attempt characterization of the data, there is a limited ability to test the effectiveness of the model in low firing rate and small data size conditions. In these circumstances, our methodology will generally overestimate spiking probabilities. Given the tendency for inaccurate spike probability estimates in small data sizes, the approximate mean firing rates for simulated data sets (Fig. 4) were calculated from the total number of spikes in a 300 s interval. To compare the conditional intensity functions estimated by our method to the ground truth functions used to simulate the spiking data, two-sample Kolmogorov–Smirnov (K-S) tests were performed between $\hat{P}$ and P. In addition, the normalized root mean squared error (NRMSE) was calculated as:

$$\text{NRMSE}_{N,g} = \log\left(\sqrt{\frac{\sum_{t=1}^{T}\left(g^{-1}P(Y=1\,\middle|\,\theta_t) - \hat{P}(Y=1\,\middle|\,\theta_t)\right)^2}{T}}\right) \tag{20}$$

where $P(Y=1|\theta)$ is given by (18), $\hat{P}(Y=1\,\middle|\,\theta)$ is given by (19), g is the maximum probability of the simulated conditional intensity function, and T is the duration of the simulated data. Note that in (20), the predicted conditional intensity function is normalized using g to ensure that the calculated error is consistent across all neuron firing rates.

As neural oscillations are often non-sinusoidal, it is important to assess the extent to which model estimates are influenced by an uneven distribution of phase data. To this end, we generated 60 s of a sinusoidal and a non-sinusoidal 8 Hz signal (Fig. 5A). The non-sinusoidal signal was generated from a bursty oscillator, with approximately 20% of the signal dedicated to the rise time of the oscillator and 80% dedicated to the decay time. The Hilbert transform was then applied to extract phase data from both signals. We then generated simulated spiking activity such that the probability of observing a spike at any given phase was the same across both sinusoidal and non-sinusoidal signal types. Using our method, we estimated the probability of observing a spike given the phase of an oscillation under three different conditions: (1) a unimodal spike-phase relationship centered at 0 radians where there is a higher sampling of phases in our non-sinusoidal signal, (2) a unimodal spike-phase relationship centered at π radians where there is a lower sampling of phases in our non-sinusoidal signal, and (3) a condition in which spiking is statistically independent of phase such that there is an equal probability (approximately 0.035) of observing a spike at every phase, based upon the average firing rate of the neuron in conditions 1 and 2.

2.10. Local field potential and single neuron recordings in rat hippocampus

Single neuron spiking activity and local field potential signals from rat hippocampus were acquired as described previously (Rangel et al., 2016). Briefly, tetrodes of 12 μm nickel-chrome wire were lowered into the CA1 region of the right dorsal hippocampus. Final tetrode locations were verified via a Nissl stain in 40 μm coronal sections. Neural

recordings were acquired using a 96-channel recording system (OmniPlex, Plexon, Dallas, TX), bandpass filtered between 400 and 8000 Hz to digitally isolate spikes, and bandpass filtered between 1 and 400 Hz to digitally isolate local field potentials. The spikes from a single neuron were isolated by comparing waveform features across tetrode wires (Offline Sorter, Plexon, Dallas, TX), including peak and valley voltage amplitudes, total peak-to-valley distance, and principal components analysis. Principal cells and interneurons were differentiated according to both firing rate and waveform characteristics. Specifically, neurons were clustered according to mean firing rate, mean width at half-maximum amplitude, and mean temporal offset from the peak to the trough of the waveform. This resulted in the identification of interneurons with mean firing rates of at least 5 Hz, a mean width half-maximum amplitude of less than 150 μs and a mean peak to trough waveform width of less than 350 μs. A 3rd-order Butterworth filter was used to bandpass filter the LFP for theta (5–10 Hz) or high gamma (65–135 Hz), and the instantaneous phase was calculated by taking the arctangent of the complex Hilbert transform of the filtered signal (Rangel et al., 2016). For this study, local field potential activity and the spiking activity from a single CA1 interneuron were simultaneously acquired from the same tetrode during 300 s when the rat was allowed to freely explore an approximately 4 ft^2 platform for randomly spaced cereal pieces.

3. Results

3.1. Simulated data

To generate unimodal spike-phase relationships, a single Von Mises distribution was selected at random to be used as a ground truth conditional probability function, as described in Section 2.8. For our example unimodal spike-phase relationship, our workflow selected six basis functions as the optimal number for characterizing the spike-phase relationship (Fig. 3A left, χ^2 =9960, $p < 0.0001$, d.f. =59,993).

To generate a multimodal spike-phase relationship, five Von Mises distributions were selected at random, averaged, and scaled to create a ground truth conditional probability function, also as described in Section 2.8. For our example multimodal spike-phase relationship, our workflow selected fourteen basis functions as the optimal number for characterizing the spike-phase relationship (Fig. 3B left, $\chi^2 = 8820$, $p < 0.0001$, d.f.= 59,985).

For both simulated unimodal and multimodal spike-phase relationships, the predicted distributions of spikes across phase closely matched the ground truth conditional probability functions that were utilized (Fig. 3A and B, middle). This is corroborated by the fact that QQ plots associated with the time-rescaled inter-spike intervals for both data sets remained within the 95% confidence interval (Fig. 3A and B, right).

3.2. Simulated data with a range of spike probabilities and data sizes

In order to assess the effectiveness of our method in capturing spike-phase relationships given neurons of different firing rates and data sets of various sizes, both the firing rate and the data size were systematically adjusted, as described in Section 2.9. The NRMSE

between the actual (i.e. ground truth) and model predicted conditional probability functions was calculated using (20) (Fig. 4A). This measure assesses the degree to which conditional probability functions estimated by our workflow are different from the ground truth function used to generate the spiking data. In addition, two-sample Kolmogorov–Smirnov tests were performed to assess whether ground truth and model predicted conditional probability functions were statistically different (Fig. 4B). In general, models generated from our workflow captured ground truth best given higher firing rates and longer data sets. Notably, Q-Q plots of actual and theoretical (i.e. model predicted) spiking data remained within the 95% confidence interval for all firing rate and data size combinations.

3.3. Simulated spike-phase relationships to sinusoidal and non-sinusoidal signals

While the simulated sinusoidal signal contained a uniform distribution of phases (Fig. 5B), the simulated non-sinusoidal signal contained a larger number of phases centered around 0 radians (Fig. 5C). To assess the extent to which this non-uniform phase distribution may influence model estimates, spiking data was generated with either (1) a statistical dependence to phase with a higher probability of firing centered at 0 radians, (2) statistical dependence to phase with a higher probability of firing centered at π radians, or (3) statistical independence from phase. These spike-phase conditions were created for both sinusoidal and non-sinusoidal signals. For both sinusoidal and non-sinusoidal signals, our workflow selected two basis functions as the optimal number for characterizing a unimodal spike-phase relationship centered at 0 radians (Fig. 5D; sinusoidal: χ^2=1980, $p < 0.0001$, d.f.=60,029; non-sinusoidal: χ^2=1760, $p < 0.0001$, d.f.=59,997). Our workflow selected three basis functions to characterize a unimodal spike-phase relationship to the sinusoidal signal centered at π radians (Fig. 5E, χ^2=2020, $p < 0.0001$, d.f.=60,028), and two basis functions to describe the same relationship to the non-sinusoidal signal (Fig. 5E, χ^2=1690, $p < 0.0001$, d.f.=59,997). Lastly, our workflow selected two regressors and a single regressor to characterize statistically independent spiking activity with respect to the phase of the sinusoidal and non-sinusoidal signal, respectively (Fig. 5F; sinusoidal: χ^2=3.34, $p < 0.188$ n.s., d.f.=60,029; non-sinusoidal: χ^2=3.39, $p < 0.0656$ n.s, d.f.=59,998). In all conditions, our method accurately estimated the ground truth probability functions used to generate the spiking data (bottom, Fig. 5D–F).

3.4. Rat hippocampal data

Our methodology predicted one basis function as the global minimum of (16) for characterizing the relationship between the spiking activity of a CA1 interneuron and the phase of a simultaneously recorded hippocampal theta (5–10 Hz) oscillation (Fig. 6D, χ^2 =497, $p < 0.0001$, d.f.=299,998). The Q-Q plot comparing the actual spiking distribution to the probability distribution described by the model reveals aspects of the data that deviate outside of the 95% confidence interval (Fig. 6G). Comparing actual and theoretical spiking data for isolated 10 s intervals within the larger 300 s dataset revealed epochs in which Q-Q plots remained within the 95% confidence interval (Fig. 6H). Notably, when examining A_q in (16) across model orders, a local minimum was also observed at a model order containing 22 basis functions (Fig. 6D, χ^2=517, $p < 0.0001$, d.f.=299,977). The associated probability distribution revealed an additional multimodal rhythmic modulation occurring at intervals within the theta cycle when spiking probability is decreased (Fig. 6E). As the period for

the theta (5–10 Hz) ranges from 100 to 200 ms, we estimated the period of the nested rhythmic modulation to correspond to approximately 65 Hz to 135 Hz. To further explore this activity, a high gamma (65–135 Hz) filter was applied to the local field potential and our workflow was reapplied to examine spike-phase relationships to this new frequency range. Our methodology selected seven basis functions as the optimal number for characterizing the spike-phase relationship to gamma (Fig. 6I, χ^2=110, $p < 0.0001$, d.f.=299,992).

4. Discussion

In this paper, we propose an objective method for creating a point process statistical model to characterize spike-phase relationships. Contrary to most other approaches, our method does not impose a unimodal spike-phase relationship shape upon the data. Rather, we utilize an overcomplete basis of Von Mises basis functions with varying mean and inverse variance combinations as potential regressors for the model. By establishing an overcomplete basis, our methodology is better able to characterize any possible spike-phase relationship shape. We then utilize a modern high-dimensional statistics approach to perform active set selection with ℓ_1-regularized maximum likelihood estimation, followed by maximum likelihood on the active regressors to calculate model weights. As ℓ_1-regularization imposes sparsity, our methodology is able to flexibly select the active set of functions that best characterizes the data at different levels of model complexity, thereby minimizing the imposition of biases regarding the shape and modality of spike-phase relationships upon the data. We then utilize the Akaike information criterion (AIC) to select the active set and corresponding model weights that best describe the data while minimizing overfitting. This method robustly characterizes a wide range of spike-phase relationships, and reveals coordinated, rhythmic relationships that were previously difficult to identify.

In simulated data, the conditional probabilities estimated by our workflow for both unimodal and multimodal spike-phase relationships were in high agreement with the ground truth relationships, as demonstrated by every Q-Q plot remaining within the 95% confidence interval (Figs. 3 and 4). This is particularly noteworthy given that the ground truth distribution was created using an identity link function in (19), whereas the model was fit according to a logit link function in (18). These results demonstrate a capacity of the model to recover the underlying shape of the spike-phase relationships even when the parametrization in our model is different than the parametrization used to generate the data.

The performance of the workflow given different neuronal firing rates and data sizes illustrates several potential issues that can arise from attempting to characterize spike-phase relationships with minimal data. For example, although the Q-Q plots suggest a good fit to the data, the ground truth and predicted spike-phase relationships are often quite different when low firing rates are combined with small data sizes (Fig. 4A and B). Thus, even though the model accurately captures spike-phase relationships in the data available, the data is not sufficient to appropriately represent the ground truth probability. In addition, as at least one spike is required to attempt characterization of spike-phase relationships, our methodology will generally overestimate spiking probability for low firing rate and small data size combinations. These issues, however, are not unique to our method. Instead, this demonstrates the limitations of this type of data and serves as a valuable reference

for spike-phase relationship characterization methods in general. Researchers may refer to Fig. 4 to assess the feasibility of characterizing spike-phase relationships given the firing rates of the neurons in the data sizes available. Researchers can further use Fig. 4 to guide future experiments in determining appropriate data sizes for a given range of neuronal firing rates to ensure that the data being modeled is representative of the underlying relationship. In future work, theoretical approaches to bound errors in the model fit (e.g. with Fisher Information (Lewi et al., 2007)) could be applied to analytically assess the uncertainty in our estimates as a function of data size or firing rate.

While our proposed workflow is agnostic to the method used to extract phase from the local field potential, it is important to address potential issues that may arise when attempting to characterize spiking relationships to non-sinusoidal waveforms. For example, there is often a concern that spurious spike-phase relationships can arise from a nonuniform distribution of phases in non-sinusoidal signals (Fig. 5C). Under conditions in which there is a statistical dependence between the simulated neuronal spiking and local field potential phase, our method accurately estimates spike probabilities to both sinusoidal and non-sinusoidal signals (Fig. 5D and E). Notably, under conditions in which the probability of observing a spike is independent of phase, conditional probability estimates did not significantly deviate from a theoretically equiprobable distribution. This indicates that our method does not introduce spurious spike-phase relationships to non-sinusoidal signals (Fig. 5F). Since our method estimates the probability of observing a spike at a given phase, a larger number of spikes observed at a particular phase due to a larger sampling of that phase is balanced by having a greater number of instances of that phase. Small deviations in phase value representation (e.g. slightly non-uniform phase distributions) such as those observed in our non-sinusoidal examples will therefore only affect the rate at which the workflow converges to the true value of that relationship. Our method therefore minimizes biases that may be introduced due to non-uniform distributions of phase, and accurately estimates spike-phase relationships to both sinusoidal and non-sinusoidal oscillations.

In addition to being able to characterize simulated data, this methodology captures spike-phase relationships in rodent hippocampal data. Interneurons in the rat hippocampus have well-established relationships to a theta (5–10 Hz) rhythm that commonly manifests in the region during periods of free exploration (Buzsáki, 2002; Colgin, 2013; Czurkó et al., 2011). Our model selection procedure identified a global minimum when using a single Von Mises function to characterize the spike-phase relationship of a CA1 interneuron to the local theta oscillation during a 300s period of random foraging. Although the Q-Q plot of actual and theoretical spiking data deviated from the 95% confidence interval (Fig. 6G), Q-Q plots of isolated 10 s intervals within the 300 s period revealed multiple intervals in which the estimated conditional probability was in high agreement with the data (Fig. 6H). Thus, the estimated conditional probability likely captures a spike-phase relationship that changes throughout the 300 s interval. This is not surprising, as spike-phase relationships in the hippocampus are tightly linked to behavior and the random foraging epoch contains intervals of free exploration, grooming, sitting, and consummatory behaviors. These results highlight the need to tightly control for behavior when characterizing spike-phase relationships, while also demonstrating the capability of the methodology at discerning relationships that evolve over time.

Our analysis of hippocampal data showcases the use of our approach as a powerful data visualization and quantification tool. The simultaneous assessment of spike-phase relationships to multiple rhythms is challenging to characterize without any *a priori* knowledge of typically co-occurring oscillations and their frequency ranges. Our model selection procedure identified a local minimum when using twenty-two Von Mises functions, revealing a relationship between a CA1 interneuron and a 65–135 Hz high gamma frequency range nested within a larger 5–10 Hz theta oscillation. Although it is well-established that the hippocampal 5–10 Hz theta rhythm exhibits coupling to a 65–135 Hz high gamma rhythm during behavioral periods of active exploration (Colgin et al., 2009; Colgin and Moser, 2010), here the high gamma frequency range was empirically defined by estimating the period of multiple peaks within the conditional probability function for a spike-phase relationship to theta. The 65–135 Hz range was thus identified through a regularly occurring spiking relationship nested within theta cycles rather than through spectral decomposition and interpretation of the LFP. Ultimately, these results prompted the development of new regressors to estimate interneuron spike-phase relationships to a 65–135 Hz high gamma rhythm. Notably, the spike-phase relationship to the gamma frequency range appears to occur at a phase of theta when the probability of observing a spike is lowest, revealing an engagement in potentially distinct rhythmic circuits. Combining regressors that characterize spike-phase relationships to multiple rhythms in a single model will ultimately allow researchers to uncover the manner in which neurons dynamically participate in multiple rhythmic circuits during information processing.

Previous studies examining spike-phase relationships have often utilized a combination of circular histograms to visualize spiking activity across the phases of an oscillation, coherence measurements to quantify relationships, and circular statistics to test for non-uniformity (e.g. Rayleigh statistic). Some studies employ statistical modeling techniques, utilizing a single regressor function (e.g. cosine) to capture gross, unimodal spike-phase relationships (Barbieri et al., 2001; Canolty et al., 2010; Dragoi and Buzsáki, 2006; Siapas et al., 2005; Siegel et al., 2009; Whittingstall and Logothetis, 2009). Using a single shape, however, such as a sine wave, imposes a constraint onto the modeled spike-phase relationship that does not necessarily represent the true relationship within the data. Some studies have attempted to capture potentially complex, multimodal relationships by using smoothing or binning techniques. One such study utilizes kernel polynomials in concert with kernel smoothing techniques to increase the adaptability of the generated model (Kim et al., 2013). These smoothing techniques, however, have the potential drawback of arbitrarily smoothing out aspects of spike-phase relationships that are potentially meaningful, if at relatively lower scale. Overcomplete basis functions have shown promise in better approximating underlying probabilistic relationships, particularly with sparse data representations (Lewicki and Sejnowski, 2000; Nakanishi-Ohno et al., 2016; Yang et al., 2011). We demonstrate in this paper that using ℓ_1-regularization coupled with an overcomplete basis set enables an unbiased characterization of a large range of potentially complex spike-phase relationships beyond what was previously possible.

Our proposed workflow can be adjusted as desired at several important stages. For example, while we have chosen Von Mises functions as regressors for our models, additional circular functions (e.g. sine waves) can also be used. In addition, while we have proposed

a particular set of Von Mises functions for the purpose of characterizing spike-phase relationships in this study, one can easily adjust the range of means and inverse variances to select the desired number of potential regressors as needed. A variety of model selection procedures may also be implemented in our workflow as an alternative to AIC that use different approaches to minimize overfitting (e.g. Bayesian information criterion, cross-validation, etc). Notably, our hippocampal data highlights the fact that interesting characteristics, such as nested rhythms, may be found at local minima that are not captured at the global minimum determined by AIC. This suggests that it may be important to consider more models than the optimal model selected by our method. Lastly, to improve the goodness of fit, additional intrinsic and extrinsic covariates can be added to the methodology akin to those previously proven to enhance spiking predictability (Truccolo et al., 2005). For example, regressors can be added pertaining to previous spiking activity for capturing refractory effects, and ensemble spiking activity for capturing local network activity effects. In addition, relating each spike-phase pair at time t with a distinct behavioral state during our sample 300 s foraging interval could drastically improve the model fit. While we acknowledge that the inclusion of additional intrinsic and extrinsic covariates will improve the goodness of fit of a model, the focus of this paper is improving the method in which a single covariate is constructed.

Overall, our method provides researchers with a robust tool for investigating rhythmic influences on neural spiking activity and the complex manner in which these influences can be further temporally coordinated. The inherent flexibility in our method allows researchers to explore a wide range of potential spike-phase relationships, thereby enabling the discovery of previously unknown rhythmic interactions. This new information can then be used to develop novel hypotheses regarding the temporally coordinated interactions underlying the processing abilities of neural networks.

Acknowledgements

References

Akaike H, 1974. A new look at the statistical model identification. IEEE Trans. Autom. Control 19 (6), 716–723.

Aoi MC, Lepage KQ, Kramer MA, Eden UT, 2015. Rate-adjusted spike—LFP coherence comparisons from spike-train statistics. J. Neurosci. Methods 240, 141–153. [PubMed: 25460189]

Barbieri R, Frank LM, Quirk MC, Solo V, Wilson MA, Brown EN, 2002. Construction and analysis of non-Gaussian spatial models of neural spiking activity. Neurocomputing 44, 309–314.

Barbieri R, Quirk MC, Frank LM, Wilson MA, Brown EN, 2001. Construction and analysis of non-Poisson stimulus-response models of neural spiking activity. J. Neurosci. Methods 105 (1), 25–37. [PubMed: 11166363]

Belloni A, Chernozhukov V, et al. , 2013. Least squares after model selection in high-dimensional sparse models. Bernoulli 19 (2), 521–547.

Berens P, et al. , 2009. CircStat: a MATLAB toolbox for circular statistics. J. Stat. Softw 31 (10), 1–21.

Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM, 2002. The time-rescaling theorem and its application to neural spike train data analysis. Neural Comput. 14 (2), 325–346. [PubMed: 11802915]

Buzsáki G, 2002. Theta oscillations in the hippocampus. Neuron 33 (3), 325–340. [PubMed: 11832222]

Buzsáki G, Anastassiou CA, Koch C, 2012. The origin of extracellular fields and currents – EEG, ECoG, LFP and spikes. Nat. Rev. Neurosci 13 (6), 407–420. [PubMed: 22595786]

Buzsáki G, Draguhn A, 2004. Neuronal oscillations in cortical networks. Science 304 (5679), 1926–1929. [PubMed: 15218136]

Canolty RT, Ganguly K, Kennerley SW, Cadieu CF, Koepsell K, Wallis JD, Carmena JM, 2010. Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies. Proc. Natl. Acad. Sci 201008306.

Chalk M, Herrero JL, Gieselmann MA, Delicato LS, Gotthardt S, Thiele A, 2010. Attention reduces stimulus-driven gamma frequency oscillations and spike field coherence in V1. Neuron 66 (1), 114–125. [PubMed: 20399733]

Chételat D, Lederer J, Salmon J, et al. , 2017. Optimal two-step prediction in regression. Electron. J. Stat 11 (1), 2519–2546.

Coleman TP, Sarma SS, 2010. A computationally efficient method for nonparametric modeling of neural spiking activity with point processes. Neural Comput. 22 (8), 2002–2030. [PubMed: 20438334]

Colgin LL, 2013. Mechanisms and functions of theta rhythms. Annu. Rev. Neurosci 36, 295–312. [PubMed: 23724998]

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser M-B, Moser EI, 2009. Frequency of gamma oscillations routes flow of information in the hippocampus. Nature 462 (7271), 353. [PubMed: 19924214]

Colgin LL, Moser EI, 2010. Gamma oscillations in the hippocampus. Physiology 25 (5), 319–329. [PubMed: 20940437]

Czanner G, Eden UT, Wirth S, Yanike M, Suzuki WA, Brown EN, 2008. Analysis of between-trial and within-trial neural spiking dynamics. J. Neurophysiol 99 (5), 2672–2693. [PubMed: 18216233]

Czurkó A, Huxter J, Li Y, Hangya B, Muller RU, 2011. Theta phase classification of interneurons in the hippocampal formation of freely moving rats. J. Neurosci 31 (8), 2938–2947. [PubMed: 21414915]

Dragoi G, Buzsáki G, 2006. Temporal encoding of place sequences by hippocampal cell assemblies. Neuron 50 (1), 145–157. [PubMed: 16600862]

Eden UT, Frank LM, Tao L, 2018. Characterizing complex, multi-scale neural phenomena using state-space models. Dynamic Neuroscience. Springer, pp. 29–52.

Engel AK, Fries P, 2010. Beta-band oscillations – signalling the status quo? Curr. Opin. Neurobiol 20 (2), 156–165. [PubMed: 20359884]

Fries P, 2009. Neuronal gamma-band synchronization as a fundamental process in cortical computation. Annu. Rev. Neurosci 32, 209–224. [PubMed: 19400723]

Fries P, Nikoli D, Singer W, 2007. The gamma cycle. Trends Neurosci. 30 (7), 309–316. [PubMed: 17555828]

Fries P, Womelsdorf T, Oostenveld R, Desimone R, 2008. The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. J. Neurosci 28 (18), 4823–4835. [PubMed: 18448659]

Haslinger R, Pipa G, Brown E, 2010. Discrete time rescaling theorem: determining goodness of fit for discrete time statistical models of neural spiking. Neural Comput. 22 (10), 2477–2506. [PubMed: 20608868]

Herreras O, 2016. Local field potentials: myths and misunderstandings. Front. Neural Circuits 10.

Jones SR, 2013. Local field potential, relationship to electroencephalography (EEG) and magnetoencephalography (MEG). Encycl. Comput. Neurosci 1–6.

Kim S, Ganguli S, Frank L, 2013. Spatial information outflow from the hippocampal circuit: distributed spatial coding and phase precession in the subiculum. J. Neurosci 32.

Kramer MA, Roopun AK, Carracedo LM, Traub RD, Whittington MA, Kopell NJ, 2008. Rhythm generation through period concatenation in rat somatosensory cortex. PLoS Comput. Biol 4 (9), e1000169. [PubMed: 18773075]

Lakshminarayanan V, Fleck A, 2011. Zernike polynomials: a guide. J. Mod. Opt 58 (7), 545–561.

Lepage KQ, Kramer MA, Eden UT, 2011. The dependence of spike field coherence on expected intensity. Neural Comput. 23 (9), 2209–2241. [PubMed: 21671792]

Lewi J, Butera R, Paninski L, 2007. Real-time adaptive information-theoretic optimization of neurophysiology experiments. Advances in Neural Information Processing Systems 857–864.

Lewicki MS, Sejnowski TJ, 2000. Learning overcomplete representations. Neural Comput. 12 (2), 337–365. [PubMed: 10636946]

Lewis LD, Weiner VS, Mukamel EA, Donoghue JA, Eskandar EN, Madsen JR, Anderson WS, Hochberg LR, Cash SS, Brown EN, et al. , 2012. Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness. Proc. Natl. Acad. Sci 109 (49), E3377–E3386. [PubMed: 23129622]

Lowet E, Roberts MJ, Bonizzi P, Karel J, De Weerd P, 2016. Quantifying neural oscillatory synchronization: a comparison between spectral coherence and phase-locking value approaches. PLoS One 11 (1), e0146443. [PubMed: 26745498]

Masseran N, Razali A, Ibrahim K, Latif M, 2013. Fitting a mixture of von Mises distributions in order to model data on wind direction in peninsular Malaysia. Energy Convers. Manage. 72, 94–102.

Mooney J, Helms P, Jolliffe I, 2003. Fitting mixtures of von Mises distributions: a case study involving sudden infant death syndrome. Comput. Stat. Data Anal 41 (3), 505–513.

Nakanishi-Ohno Y, Obuchi T, Okada M, Kabashima Y, 2016. Sparse approximation based on a random overcomplete basis. J. Stat. Mech.: Theory Exp 2016 (6), 063302.

Olson J, Tongprasearth K, Nitz D, 2017. Subiculum neurons map the current axis of travel. Nat. Neurosci 20 (2), 170–172. [PubMed: 27991899]

Pesaran B, Nelson MJ, Andersen RA, 2008. Free choice activates a decision circuit between frontal and parietal cortex. Nature 453 (7193), 406. [PubMed: 18418380]

Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA, 2002. Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat. Neurosci 5 (8), 805. [PubMed: 12134152]

Pesaran B, Vinck M, Einevoll GT, Sirota A, Fries P, Siegel M, Truccolo W, Schroeder CE, Srinivasan R, 2018. Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. Nat. Neurosci

Poor HV, 2013. An Introduction to Signal Detection and Estimation. Springer Science & Business Media.

Rangel LM, Rueckemann JW, Rivière PD, Keefe KR, Porter BS, Heimbuch IS, Budlong CH, Eichenbaum H, 2016. Rhythmic coordination of hippocampal neurons during associative memory processing. Elife 5, e09849. [PubMed: 26751780]

Roopun AK, Kramer MA, Carracedo LM, Kaiser M, Davies CH, Traub RD, Kopell NJ, Whittington MA, 2008. Period concatenation underlies interactions between gamma and beta rhythms in neocortex. Front. Cell. Neurosci 2, 1. [PubMed: 18946516]

Sherman MA, Lee S, Law R, Haegens S, Thorn CA, Hämäläinen MS, Moore CI, Jones SR, 2016. Neural mechanisms of transient neocortical beta rhythms: converging evidence from humans, computational modeling, monkeys, and mice. Proc. Natl. Acad. Sci 113 (33), E4885–E4894. [PubMed: 27469163]

Siapas AG, Lubenov EV, Wilson MA, 2005. Prefrontal phase locking to hippocampal theta oscillations. Neuron 46 (1), 141–151. [PubMed: 15820700]

Siegel M, Warden MR, Miller EK, 2009. Phase-dependent neuronal coding of objects in short-term memory. Proc. Natl. Acad. Sci 106 (50), 21341–21346. [PubMed: 19926847]

Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsáki G, 2008. Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. Neuron 60 (4), 683–697. [PubMed: 19038224]

Skaggs WE, McNaughton BL, Wilson MA, Barnes CA, 1996. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6 (2), 149–172. [PubMed: 8797016]

Tibshirani RJ, Taylor J, et al. , 2012. Degrees of freedom in lasso problems. Ann. Stat 40 (2), 1198–1232.

Tiesinga P, Sejnowski TJ, 2009. Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? Neuron 63 (6), 727–732. [PubMed: 19778503]

Truccolo W, Eden U, Fellows M, Donoghue J, Brown E, 2005. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. J. Neurophysiol 93 (2), 1074–1089. [PubMed: 15356183]

Vinck M, Battaglia FP, Womelsdorf T, Pennartz C, 2012. Improved measures of phasecoupling between spikes and the local field potential. J. Comput. Neurosci 33 (1), 53–75. [PubMed: 22187161]

Vinck M, Oostenveld R, Van Wingerden M, Battaglia F, Pennartz CM, 2011. An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample-size bias. Neuroimage 55 (4), 1548–1565. [PubMed: 21276857]

Whittingstall K, Logothetis NK, 2009. Frequency-band coupling in surface EEG reflects spiking activity in monkey visual cortex. Neuron 64 (2), 281–289. [PubMed: 19874794]

Whittington MA, Traub R, Kopell N, Ermentrout B, Buhl E, 2000. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. Int. J. Psychophysiol 38 (3), 315–336. [PubMed: 11102670]

Yang P, Tang G, Nehorai A, 2011. Sparsity-enforced regression based on over-complete dictionary. In: 2011 4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). IEEE. pp. 261–264.

Zou H, Hastie T, Tibshirani R, et al. , 2007. On the "degrees of freedom" of the lasso. Ann. Stat 35 (5), 2173–2192.
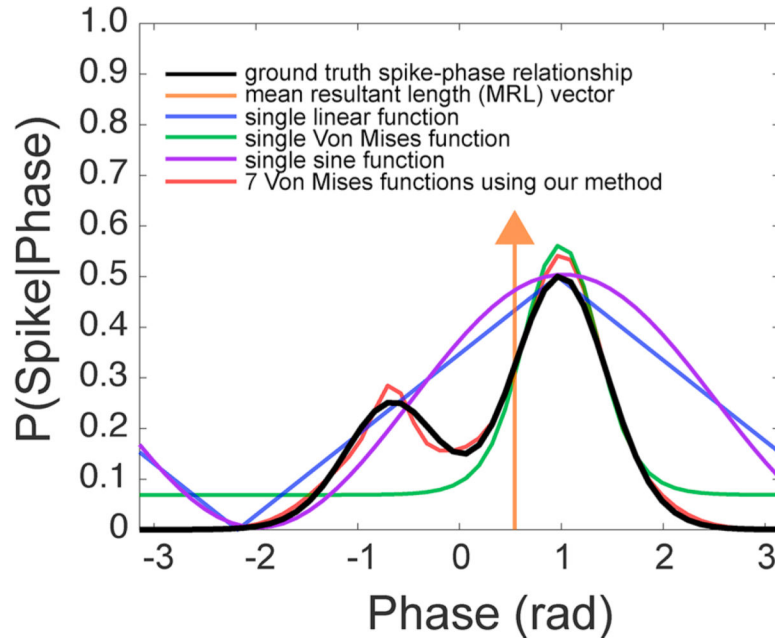
Fig. 1.

Illustration of select methods for visualizing and quantifying spike-phase relationships. A multimodal spike-phase relationship simulated from a mixture of Von Mises functions is shown in black. The estimated conditional probability of observing a spike given the phase of an oscillation using a single line function (blue), a single Von Mises function (green), a single sine function (purple), and 7 Von Mises functions assembled with our proposed workflow (red) is shown. In addition, we calculate the mean resultant length vector for the data, and show the estimated strength and preferred phase using this method (orange arrow). Our method is the only method capable of capturing the multimodal nature of this example spike-phase relationship. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
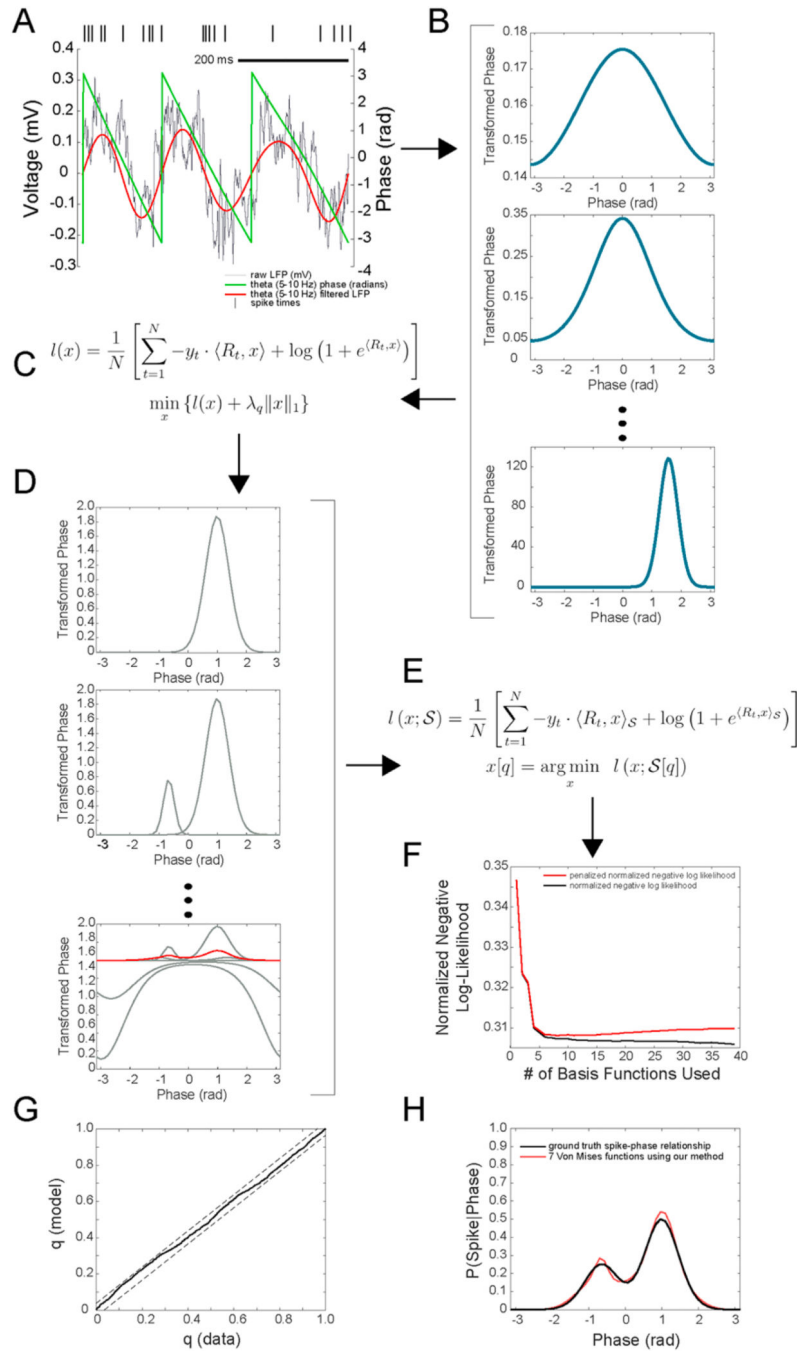
Fig. 2.
Illustration describing the methodological pipeline for predicting the probability of a spike given the phase of an oscillation. (A) Raw local field potential signal (gray), filtered signal in the theta (5–10 Hz) frequency range (red), and instantaneous phase of the theta oscillation (green) with corresponding spike times above (black tick marks). (B) A sample set of Von Mises basis functions used to construct model regressors. (C) Equation for calculating basis function weights using a ℓ_1-regularized generalized linear model, where the regularizer coefficient (λ_q) is systematically varied in order to selectively order the number of basis

functions used. (D) Sample ordering of Von Mises basis functions used to characterize the data in Fig. 1, showing the active set of Von Mises functions with increasing model order complexity (gray). Our method selected seven functions as the optimal number of functions to represent the data (bottom), and we show the weighted combination of these functions for comparison (see also red line in Fig. 1). (E) Equation for calculating the weights for each basis function based upon the ordering determined in step C. (F) Normalized negative log likelihood (black) and penalized normalized negative log likelihood (red). (G) Sample Q-Q plot of actual and theoretical spiking data with associated 95% confidence interval. (H) Conditional probability estimated by the model that best characterizes the data while minimizing overfitting according to F. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
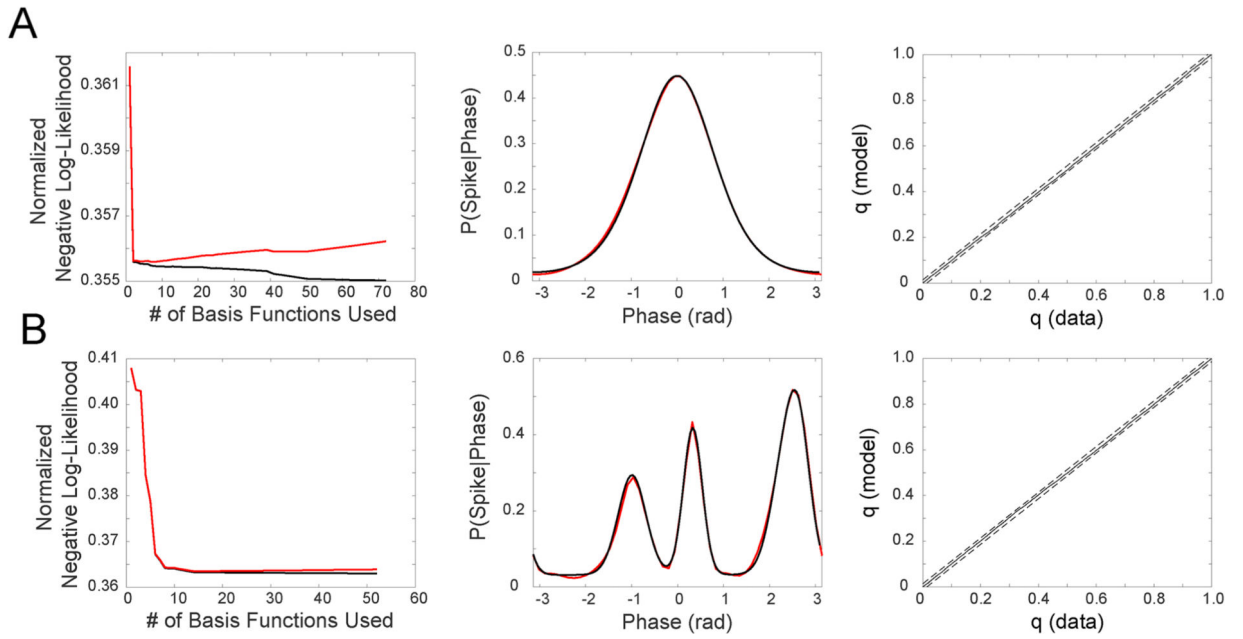
Fig. 3.
Model results for simulated unimodal (A) and multimodal (B) spike-phase relationships. Left: Normalized negative log likelihood (black) and penalized normalized negative log likelihood (red). Middle: Ground truth conditional probability function (black) and predicted conditional probability function (red). Right: Q-Q plots (solid line) with associated 95% confidence intervals (dashed lines). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
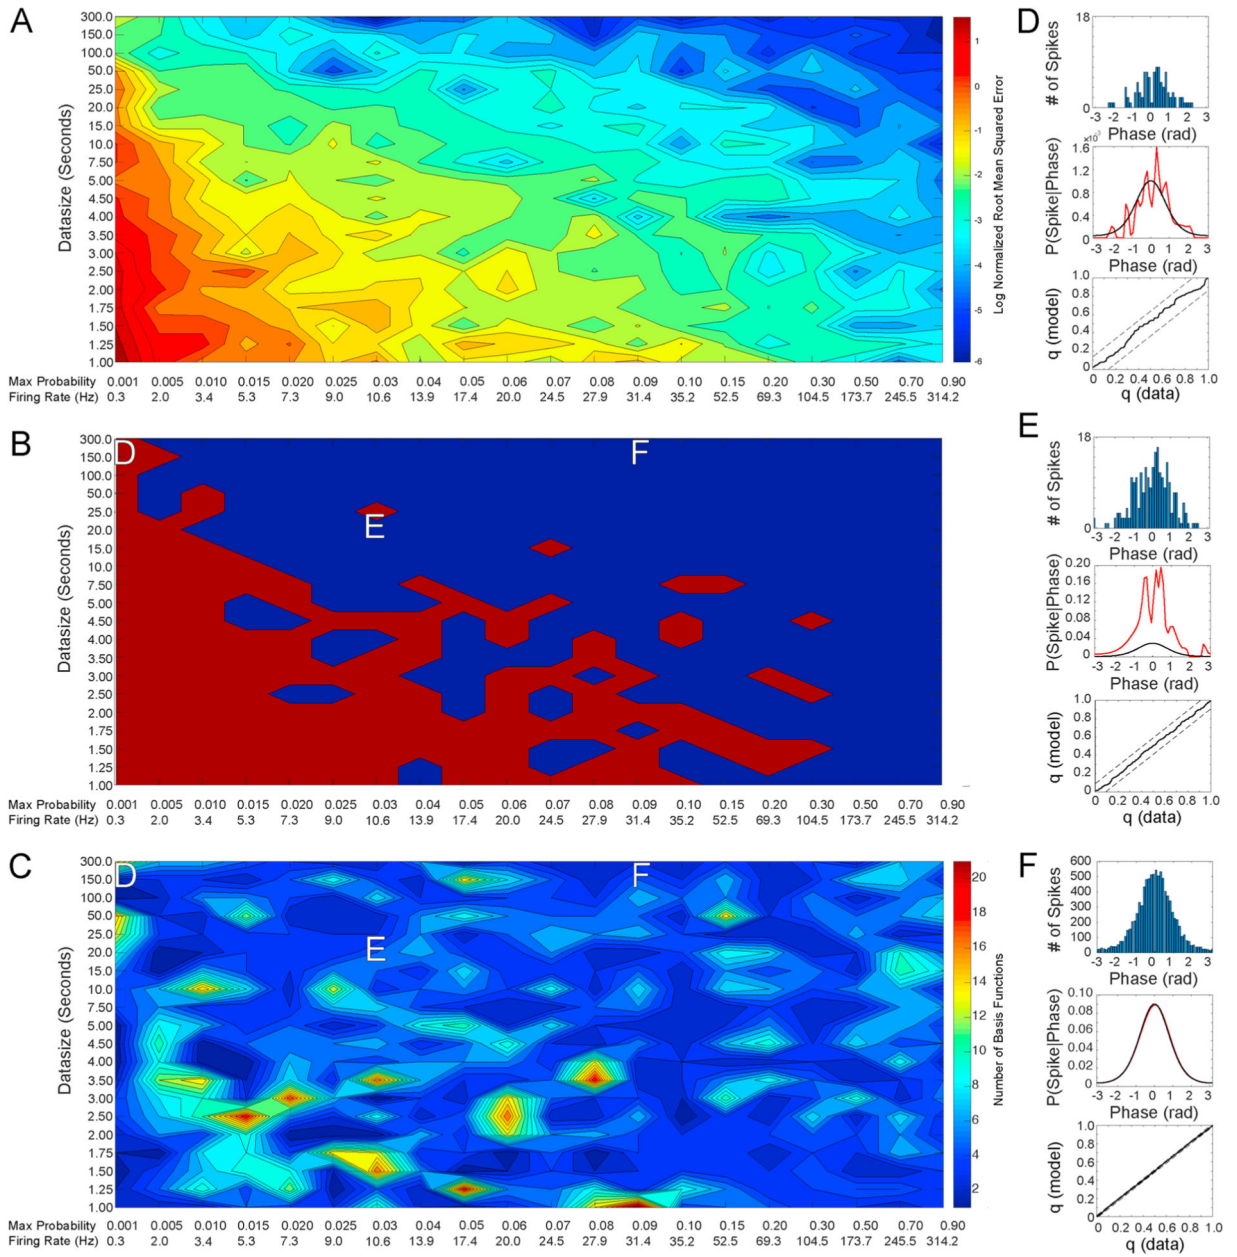
Fig. 4.
Model results given data sets with a range of different spiking probabilities and data sizes. (A) Log of the normalized root mean squared error (NRMSE) between conditional probability functions estimated by our method and the ground truth function used to generate the spiking data (see Eq. (20)). Maximum spiking probabilities (*x*-axis) range from 0.001 to 0.9 and data sizes (*y*-axis) range from 1 to 300 s (for method, see Section 2.9). Approximate firing rates, indicated below spike probabilities on the *x*-axis, are calculated from the total number of spikes in a 300 s interval. (B) Color map indicating whether ground truth spike-phase distributions and model predicted distributions are significantly different (red), assessed using a Kolmogorov–Smirnov test with a significance threshold of *p*=0.05. (C) The number of basis functions utilized after implementing model selection procedures.

White labels D, E, and F within figures A–C indicate the spike probabilities and data sizes for examples in D–F. (D–F) Top: Distribution of simulated spikes across the phases of a 20 Hz oscillation. Middle: Ground truth conditional probability function (black) and predicted conditional probability function (red). Bottom: Q-Q plots (solid line) with associated 95% confidence intervals (dashed lines). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
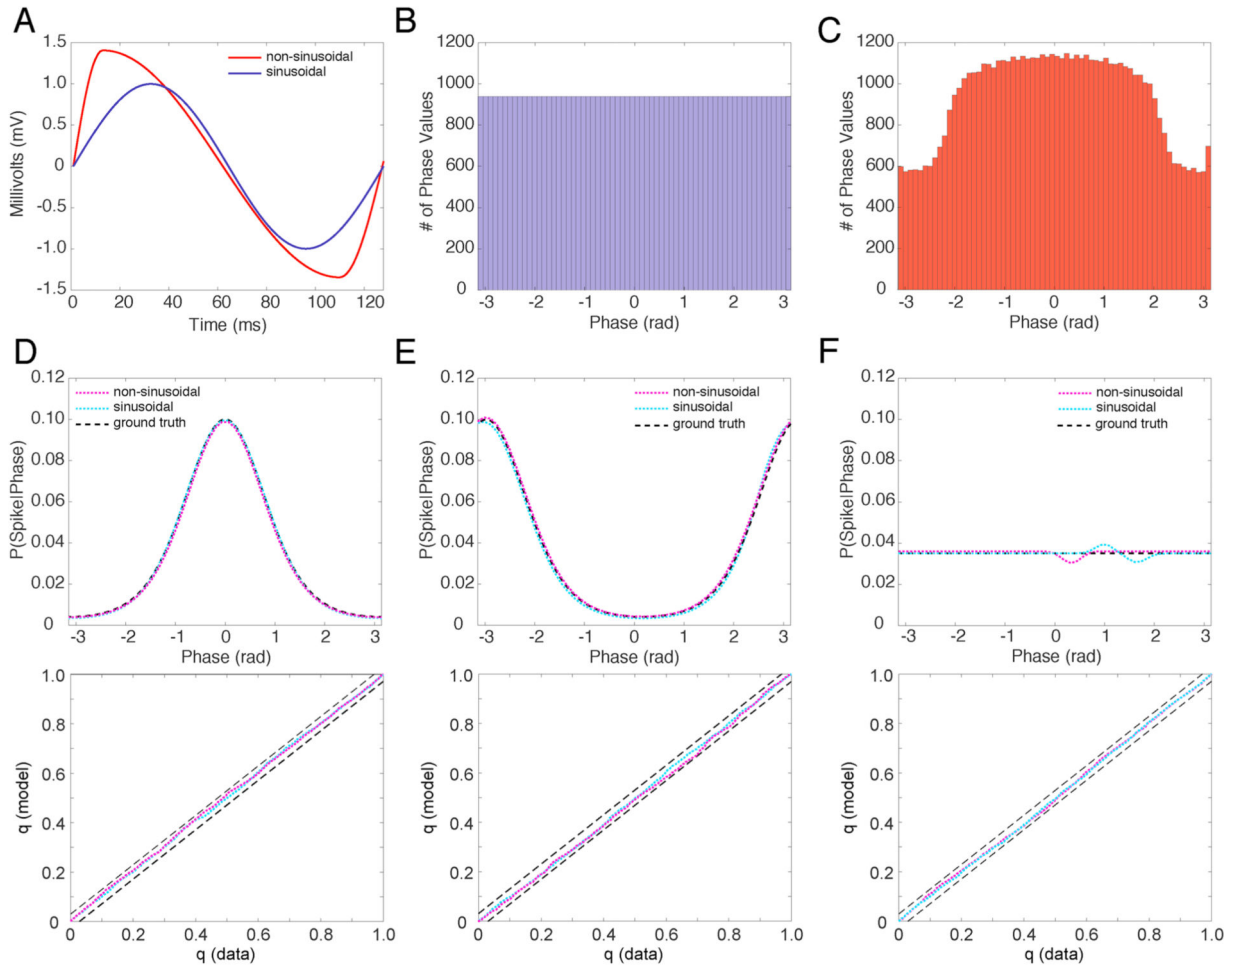
Fig. 5.
Model results for spike-phase relationships to sinusoidal and non-sinusoidal signals. (A) Single cycle of a simulated sinusoidal and non-sinusoidal signal. Distribution of phases for 60 s of a (B) sinusoidal (blue) and (C) non-sinusoidal (red) 8 Hz signal. While the sinusoidal signal contains a uniform distribution of phases, the non-sinusoidal condition contains a larger number of phases centered around 0 radians. (D–F) The estimated probability of observing a spike given the phase of a sinusoidal (dashed cyan) and non-sinusoidal (dashed magenta) oscillation for three different conditions: (D) a unimodal spike-phase relationship centered at 0 radians where there is a higher sampling of phases in our non-sinusoidal signal, (E) a unimodal spike-phase relationship centered at π radians where there is a lower sampling of phases in our non-sinusoidal signal, and (F) a condition in which there is an equal probability (0.035) of observing a spike at every phase that is based upon the average firing rate in D and E. Ground truth conditional probability functions for D–F are shown in black dashed lines. Corresponding Q-Q plots (colored dashed lines) with associated 95% confidence intervals (black dashed lines) are shown directly below (D–F). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
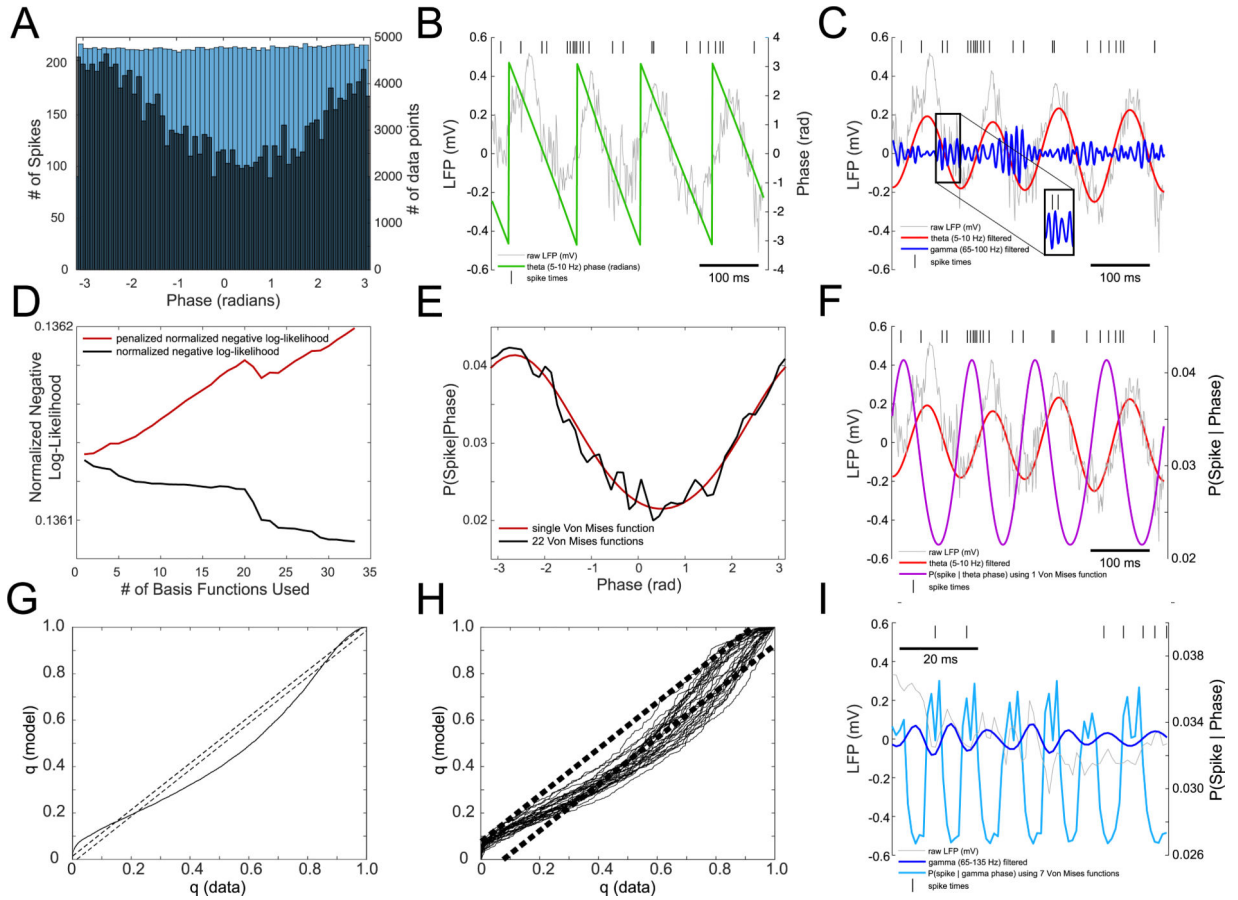
Fig. 6.

Rat hippocampal neuron data. (A) Histogram of interneuron spikes across the phases of a 5–10 Hz theta oscillation (black, left axis) and a histogram presenting the distribution of theta phases from the filtered LFP (blue, right axis). (B) Representative segment of raw local field potential data (gray), theta phase (green), and spiking times (black tick marks). (C) Raw local field potential data (gray), 5–10 Hz theta filtered local field potential (red), gamma filtered local field potential (blue), and spike times (black tick marks). (D) Normalized negative log likelihood (black) and AIC penalized normalized negative log likelihood (red). (E) Estimated conditional probability of observing a spike given the phase of theta using one basis function (red) and twenty-two basis functions (black). (F) Raw local field potential data (gray), theta filtered local field potential (red), single basis function theta conditional intensity function (magenta), and spike times (black tick marks). (G and H) Q-Q plots of actual and theoretical spiking data for (G) 300 s and (H) isolated 10 s intervals using the model selected by the workflow. (I) Raw local field potential data (gray), gamma filtered local field potential (blue), gamma conditional intensity function (cyan) and spike times (black tick marks). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)