

Cross-correlation of instantaneous amplitudes of field potential oscillations: A straightforward method to estimate the directionality and lag between brain areas

Avishek Adhikari^a, Torfi Sigurdsson^b, Mihir A. Topiwala^b, Joshua A. Gordon^{b,c,*}

^a Department of Biological Sciences, Columbia University, New York, NY 10032, United States

^b Department of Psychiatry, Columbia University, New York, NY 10032, United States

^c Division of Integrative Neuroscience, New York State Psychiatric Institute, New York, NY 10032, United States

ARTICLE INFO

Article history:

Received 23 December 2009

Received in revised form 8 June 2010

Accepted 21 June 2010

Keywords:

Cross-correlation

Theta oscillations

Medial prefrontal cortex

Ventral hippocampus

Local field potential

Functional connectivity

Partial directed coherence

ABSTRACT

Researchers performing multi-site recordings are often interested in identifying the directionality of functional connectivity and estimating lags between sites. Current techniques for determining directionality require spike trains or involve multivariate autoregressive modeling. However, it is often difficult to sample large numbers of spikes from multiple areas simultaneously, and modeling can be sensitive to noise. A simple, model-independent method to estimate directionality and lag using local field potentials (LFPs) would be of general interest. Here we describe such a method using the cross-correlation of the instantaneous amplitudes of filtered LFPs. The method involves four steps. First, LFPs are band-pass filtered; second, the instantaneous amplitude of the filtered signals is calculated; third, these amplitudes are cross-correlated and the lag at which the cross-correlation peak occurs is determined; fourth, the distribution of lags obtained is tested to determine if it differs from zero. This method was applied to LFPs recorded from the ventral hippocampus and the medial prefrontal cortex in awake behaving mice. The results demonstrate that the hippocampus leads the mPFC, in good agreement with the time lag calculated from the phase locking of mPFC spikes to vHPC LFP oscillations in the same dataset. We also compare the amplitude cross-correlation method to partial directed coherence, a commonly used multivariate autoregressive model-dependent method, and find that the former is more robust to the effects of noise. These data suggest that the cross-correlation of instantaneous amplitude of filtered LFPs is a valid method to study the direction of flow of information across brain areas.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Recent advances in multi-site recording technology have enabled researchers to sample local field potentials (LFPs) simultaneously from multiple brain regions (DeCoteau et al., 2007). A common interest in such studies is to determine whether one brain region is leading or lagging relative to another, and to estimate the time lag between putatively connected areas. Several groups have estimated the lag across brain areas using recordings of spike trains. Most of these studies estimate directionality by calculating the cross-correlation of spike trains of two areas (Alonso and Martinez, 1998; Holdefer et al., 2000; Lindsey et al., 1992; Snider et al., 1998). Other studies have used related approaches, such as calculating the cross-covariance of spike trains (Siapas et al., 2005), or

different methods, such as the computation of spike-triggered joint histograms (Paz et al., 2009) or the change in phase-locking after shifting the spikes relative to the LFP (Siapas et al., 2005). Although such methods are effective, they are not applicable to studies that record only LFPs. This situation is common, as often spike trains cannot be sampled from multiple areas, or firing rates are too low to determine the directionality of functional connectivity across regions. Recording LFPs in areas with low firing rates can be advantageous, as LFPs can be sampled continuously, while spikes can occur infrequently and irregularly. LFP-based methods may therefore yield higher temporal resolution and greater statistical power than spike-based methods.

Existing methods such as Granger causality (Cadotte et al., 2010; Gregoriou et al., 2009; Popa et al., 2010) and partial directed coherence (PDC) (Astolfi et al., 2006; Baccala and Sameshima, 2001; Taxidis et al., 2010; Winterhalder et al., 2005) are able to estimate the directionality of functional connectivity using only LFPs. However, these methods are mathematically complex (Gourevitch et al., 2006), relying on multivariate models with a large number of free parameters, obscuring the intuitive understanding of what these

* Corresponding author at: Department of Psychiatry, Columbia University, 1051 Riverside Drive, Unit 87, Kolb Annex Room L174, New York, NY 10032, United States. Tel.: +1 212 543 6768; fax: +1 212 543 1174.

E-mail address: jg343@columbia.edu (J.A. Gordon).

methods are actually computing. They can also be sensitive to noise (Taxidis et al., 2010; Winterhalder et al., 2005). Furthermore, such methods generally do not provide estimates of the time lag between brain areas.

Here we report a novel and mathematically straightforward method to estimate the lag between two brain areas that does not require spikes and that can be applied to datasets in which only LFPs have been acquired. The method requires that functional connectivity between the examined structures be accompanied by reasonably coherent activity within a specific frequency range. The method consists of determining the position (or “lag”) of the peak of the cross-correlation of the amplitude envelopes of the LFPs after filtering for the frequency range of interest. Lastly, a non-parametric signed-rank test is performed to verify if the distribution of lags obtained from multiple experiments differs from zero.

To investigate its validity, this method was applied to a dataset in which both spikes and LFPs were recorded from the medial prefrontal cortex (mPFC), while only LFPs were sampled from the ventral hippocampus (vHPC). These areas were chosen because there is a unidirectional projection from the vHPC to the mPFC (Parent et al., 2009; Verwer et al., 1997), suggesting that activity in the vHPC should lead that in the mPFC. Moreover, we have shown theta-frequency (4–12 Hz) coherence between these structures during behavior (Adhikari et al., 2010), suggesting that directionality analysis can be performed in this frequency range with the amplitude cross-correlation measure.

Using the amplitude cross-correlation method, we demonstrate that the vHPC leads the mPFC in the theta-frequency range, with a lag consistent with estimates of the conduction delay of this pathway. Furthermore, there is good agreement between vHPC–mPFC lags calculated with this method and those calculated from phase locking of mPFC spikes to vHPC theta oscillations. Finally, a consistent lag between the vHPC and the mPFC was only found in the theta, but not in the delta and gamma-frequency ranges, in line with studies suggesting that theta-frequency oscillations drive functional connectivity between the hippocampus and the mPFC (Adhikari et al., 2010; Jones and Wilson, 2005; Siapas et al., 2005). The current method was also compared to partial directed coherence (PDC), an existing method to calculate the directionality of functional connectivity with LFPs. This method, similarly to the amplitude cross-correlation method, also demonstrated that the vHPC leads the mPFC in the theta-range. To further compare the two methods, both were applied to simulations in which pink noise was added to biological signals. Strikingly, PDC was more susceptible to errors induced by noise than the amplitude cross-correlation method. These results show that the cross-correlation of the amplitude of filtered field potentials may provide a valid, relatively robust estimation of the lag and the directionality of information flow across brain areas.

2. Materials and methods

2.1. Animals

Three to six month old male wildtype 129Sv/Ev mice were obtained from Taconic (Germantown, NY, USA). Seventeen mice were used for the simultaneous mPFC and vHPC recordings. Sixteen mice were used for the simultaneous vHPC and dorsal hippocampus (dHPC) recordings. An additional cohort of five C57/Bl6 mice bred at Columbia University was used for the simultaneous dHPC and mPFC recordings, from which mPFC single units were isolated. The procedures described here were conducted in accordance with National Institutes of Health regulations and approved by the Columbia

University and New York State Psychiatric Institute Institutional Animal Care and Use Committees.

2.2. Surgery and microdrive construction

Custom microdrives were constructed using interface boards (EIB-16, Neuralynx, Bozeman MT) fastened to a Teflon platform, as described previously (Adhikari et al., 2010). Briefly, animals were anesthetized with ketamine and xylazine (165 and 5.5 mg/kg, in saline) and secured in a stereotactic apparatus (Kopf Instruments, Tujunga, CA). Screws were implanted on the posterior and anterior portions of the skull to serve as ground and reference, respectively. mPFC electrodes were implanted in the deep layers (V/VI) of the prelimbic cortex, at +1.65 mm anterior, 0.5 mm lateral and 1.5 mm depth, relative to bregma. vHPC electrodes were implanted in the CA1 region at 3.16 mm posterior, 3.0 mm lateral and 4.2 mm depth, and dHPC electrodes were targeted to 1.94 mm posterior, 1.5 mm lateral and 1.3 mm depth. Depth was measured relative to brain surface.

2.3. Behavioral protocol

Animals were permitted to recover for at least one week or until regaining pre-surgery body weight. Mice were then exposed to a small rectangular box in the dark, in which they foraged for pellets for 10 min for the mPFC–vHPC and vHPC–dHPC datasets. Mice performed an alternation task in a T-shaped maze for the dHPC–mPFC dataset as described in (Sigurdsson et al., 2010).

2.4. Data acquisition

Recordings were obtained via a unitary gain head-stage preamplifier (HS-16; Neuralynx) attached to a fine wire cable suspended on a pulley so as not to add any weight to the animal's head. LFPs were recorded against a reference screw located at the anterior portion of the skull. Field potential signals were amplified, bandpass filtered (1–1000 Hz) and acquired at 1893 Hz. Multiunit activity from the mPFC was recorded simultaneously from the same electrodes used to obtain LFPs; multiunit signals were bandpass filtered (600–6000 Hz) and recorded at 32 kHz. Events exceeding a threshold of 40 μ V were selected for analysis of phase-locking to theta (see below). Both LFP and multiunit data were acquired by a Lynx 8 programmable amplifier (Neuralynx) on a personal computer running Cheetah data acquisition software (Neuralynx). The animal's position was obtained by overhead video tracking (30 Hz) of two light-emitting diodes affixed to the head stage.

2.5. Cross-correlation analysis

2.5.1. Band-pass filtering

Data was imported into Matlab for analysis using custom-written software. To calculate the lag between the vHPC and the mPFC, signals were initially band-pass filtered between 7 and 12 Hz. A finite impulse response filter of order n , where n is the sampling frequency, was implemented with a Hamming window, utilizing the MATLAB function *fir1*.

2.5.2. Instantaneous amplitude using the Hilbert transform

The Hilbert transform of each signal was computed with the MATLAB function *hilbert*. The output of the Hilbert transform is a vector containing complex numbers that has the same number of elements as the input signal. The real portion of the complex number is the input itself, while the imaginary part is the input LFP shifted by 90° ($\pi/2$ radians). The absolute magnitude of the complex number at a given time point is the power of the filtered signal at that sample. The magnitude of a complex number is the length of the vector in the complex plane.

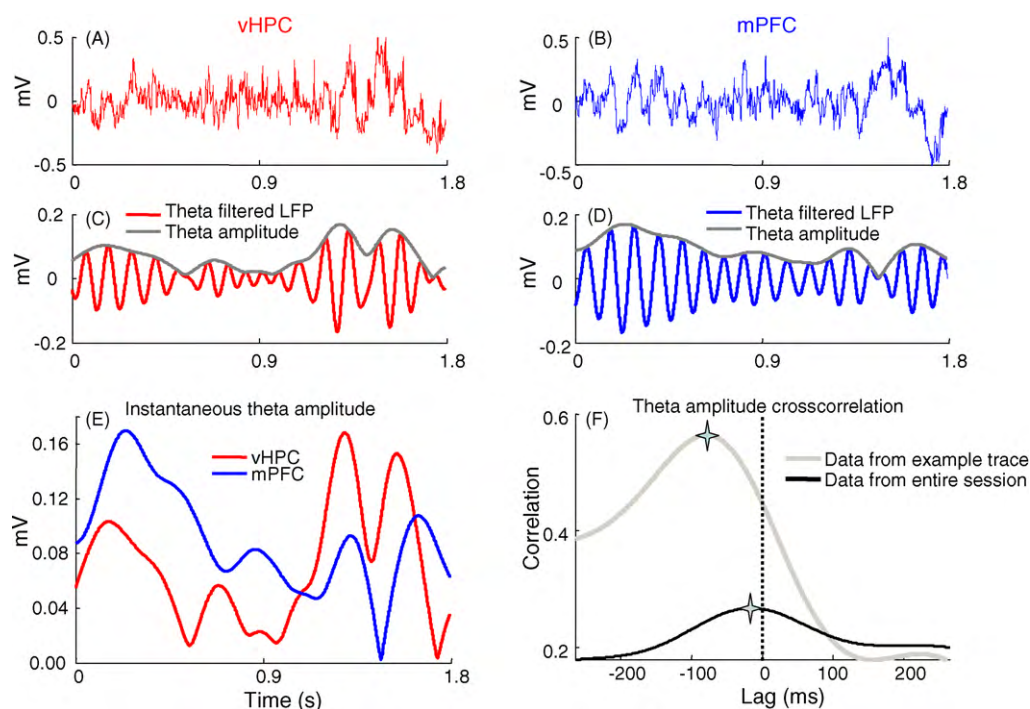


Fig. 1. Calculating lags using the amplitude cross-correlation method. (A and B) Simultaneous local field potential recordings obtained from the vHPC (A) and mPFC (B) of a behaving mouse (C and D) Traces in (A) and (B), filtered for theta-frequency activity (7–12 Hz) (red), overlaid with the instantaneous amplitude obtained from the Hilbert transform (grey). (E) The instantaneous amplitudes of the theta-filtered vHPC (red) and mPFC (blue) signals shown in (C) and (D) overlaid for comparison. (F) Cross-correlation of the example amplitude traces shown in (E) (grey). The amplitude cross-correlation from the entire recording session is overlaid (black). The peaks are indicated by stars.

All the results shown in the main text (Figs. 1–4) were obtained by using the Hilbert transform to calculate the instantaneous amplitudes of the LFPs. Importantly, these results are not dependent on the specific method used to calculate instantaneous amplitudes, as they were reproduced using the Gabor Transform (Supplemental Fig. 1) (Le Van Quyen et al., 2001).

2.5.3. Cross-correlation of the instantaneous amplitudes

After the instantaneous amplitude for all the points in the vHPC and mPFC signals was calculated, the cross-correlation between the amplitudes of the two signals was computed with the MATLAB function *xcorr*, over lags ranging from +0.1 to –0.1 s. The mean amplitude was first subtracted from each vector prior to cross-correlating them, as the DC component of a signal has no relevance for a cross-correlation. The lag at which the cross-correlation peaked was then determined. The significance of each vHPC–mPFC theta amplitude cross-correlation was verified before inclusion in additional analyses using a bootstrap procedure. mPFC and vHPC theta-amplitude envelopes were randomly shifted 5–10 s relative to each other 1000 times. The shifted amplitude envelopes were then cross-correlated, yielding a distribution of cross-correlation peaks expected by chance. The original cross-correlation was considered significant if its peak value was greater than 95% of these randomly generated cross-correlation peaks. The peaks of the mPFC–vHPC theta-filtered amplitude envelope cross-correlations of all 17 animals were significant by this criterion.

2.5.4. Distribution of cross-correlation lags

After cross-correlations of the filtered amplitude vectors were computed, the distribution of the lags at which the cross-correlation peaks occur was obtained. Wilcoxon's non-parametric rank sum test was performed on the sample of lags to verify whether the mean of the distribution was significantly different from zero.

In the convention used in the current work, a negative lag indicates that the vHPC leads the other brain area. Custom written Matlab code to execute this analysis is included in Supplemental text.

2.6. Phase-locking analysis

In order to verify if the above method produces reliable estimates of the lag between two brain areas, the results were compared with an alternative method. In this method, the strength of phase locking of mPFC spikes to vHPC theta oscillations was computed after shifting the spike train by positive or negative shifts. Multiunit spikes represent all the spikes that exceeded 40 μV , whereas spikes for single unit measurements were clustered offline using spike sort 3-D (Neuralynx). Phase locking of spikes to oscillations was assessed by calculating the mean resultant length vector (MRL), a measure from which Rayleigh's z statistic of circular uniformity is derived (Sigurdsson et al., 2010).

To determine whether spikes were significantly phase-locked to theta, theta phases of LFPs were determined through the Hilbert transform, and a phase was assigned to each spike based on the time of the spike's occurrence. The phase is obtained by calculating the angle of the absolute magnitude of the Hilbert transform output. A phase of zero refers to the trough of the theta cycle. To determine the lag between multiunit activity and theta oscillations in each area, phase locking was calculated for 40 different temporal offsets for each multiunit recording, ranging from –100 to +100 ms. Recordings with significant Bonferroni-corrected phase locking in at least one of the shifts were used for the analysis in Fig. 2.

2.7. Partial directed coherence

To compare the current method to an existing method of calculating directionality between LFPs, we analyzed our dataset using partial directed coherence (PDC). PDC is a frequency-domain rep-

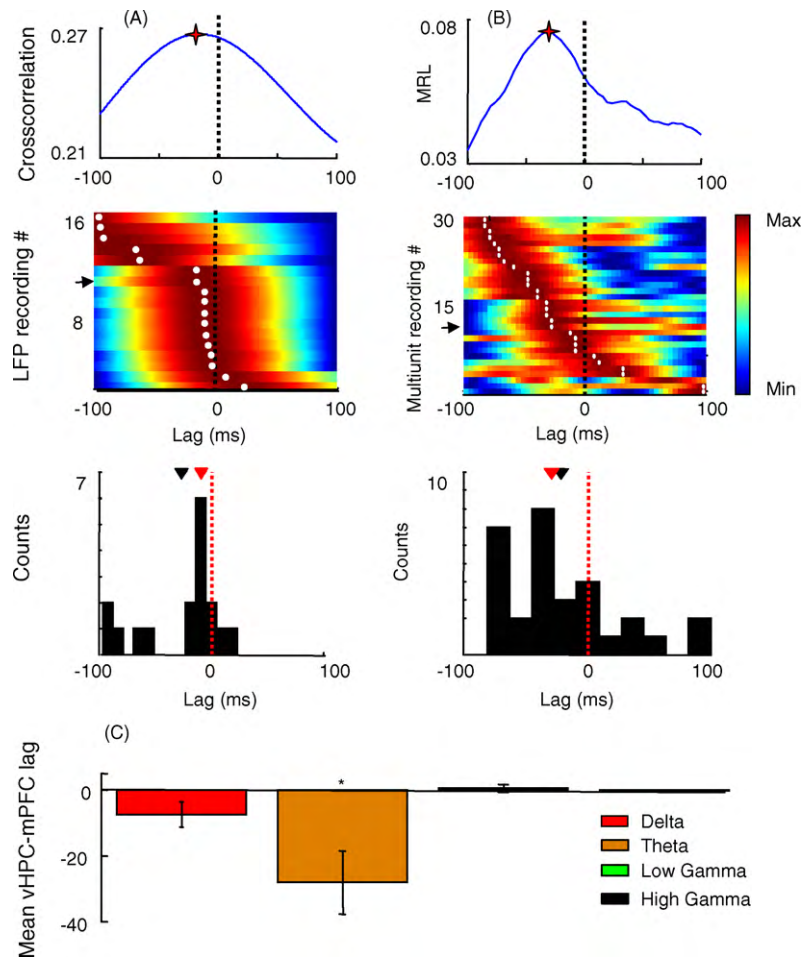


Fig. 2. Estimation of lags between the vHPC and the mPFC using the amplitude cross-correlation (A) and phase locking (B) methods. Upper panels: vHPC-mPFC lag estimate from single recording sessions using the amplitude cross-correlation (A), and the effect of shifting the mPFC spike train on the strength of phase-locking (MRL) to vHPC theta oscillations (B). Stars mark the peaks. Middle panels: normalized color plots of amplitude cross-correlations from 17 recordings (A) and phase-locking shifts from 30 recordings (B). Warmer colors indicate higher cross-correlation peaks or greater phase-locking strength. Each row corresponds to a single LFP (A) or multiunit (B) recording. Rows are arranged according to the peak lag. Arrows mark the rows representing the data shown in the upper panels. The lags at which the cross-correlation (A, middle panel) and phase-locking (B, middle panel) peaks occur are marked with white dots. Lower panels: histograms showing the distribution of peak lags calculated with each method. The distribution of lags is significantly negative for both the amplitude cross-correlation ($p < 0.05$, Wilcoxon rank-sum test, mean lag -28 ± 16.7 ms, median = -9.5 ms, $n = 17$ recordings) and phase-locking ($p < 0.05$ for a paired Wilcoxon's rank sum test, mean lag -24.5 ± 15.7 ms, median = -32 ms, $n = 30$ recordings). Means and medians of the lag distributions are indicated, respectively, by black and red arrowheads. (C) Lag estimates are frequency-specific. Lags were calculated by cross-correlating the amplitudes after filtering for delta (1–4 Hz), theta (7–12 Hz), low gamma (30–50 Hz) and high gamma (50–100 Hz) frequency ranges. Data are presented as means \pm 95% confidence intervals. * $p < 0.01$ for a paired Wilcoxon's rank sum test. In all panels, negative lags indicate that the vHPC leads the mPFC.

resentation of Granger Causality. Given two time series $X_1(t)$ and $X_2(t)$, $X_2(t)$ is said to Granger-cause $X_1(t)$ if knowledge of the past of $X_2(t)$ improves the prediction of $X_1(t)$ beyond how much the past of X_2 can predict itself in the present. $X_2(t)$ may Granger-cause $X_1(t)$ without $X_1(t)$ Granger-causing $X_2(t)$, in which case the predominant directionality would be from X_2 to X_1 ($X_1 \leftarrow X_2$). It is important to note that Granger-causation is directly related only to an increase in predictability and not to causality *per se* or to precedence in time of one signal relative to the other.

In order to compute Granger Causality of an m -variate time-series, a vector autoregressive model was fit to the process, as shown below and reported previously in more detail (Taxidis et al., 2010):

$$\begin{bmatrix} x_1(t) \\ \vdots \\ x_m(t) \end{bmatrix} = \sum_{r=1}^p A_r \begin{bmatrix} x_1(t-r) \\ \vdots \\ x_m(t-r) \end{bmatrix} + \begin{bmatrix} u_1(t) \\ \vdots \\ u_m(t) \end{bmatrix}$$

Here, the model was fit with MATLAB's ARFIT toolbox, with $m = 2$, as only two time-series were considered (vHPC and mPFC raw LFPs).

In the above equation, $u_1(t) \dots u_m(t)$ are the residuals of the model. The residuals were analyzed to verify if they described a white noise process. As expected, the residuals were normally distributed and had the same power at all frequencies. The order of the process, p , determines how many past lags are considered in the model, and was selected using Schwarz's Bayesian Criterion. N most animals the order chosen was approximately 90 samples, corresponding to 47 ms. This model order is appropriate to detect directionality in the vHPC-mPFC circuit, as it represents a period of time larger than the conduction delay of the pathway (approximately 16 ms). To obtain a frequency-domain representation of the data, the Fourier transform of the coefficients of the vector autoregressive model was taken, which for a given frequency f , was computed as shown:

$$A(f) = \sum_{r=1}^p A_r e^{-2\pi i f r}$$

PDC is an estimate of the strength of the directionality between the signals for each frequency. PDC for time series j to time series i

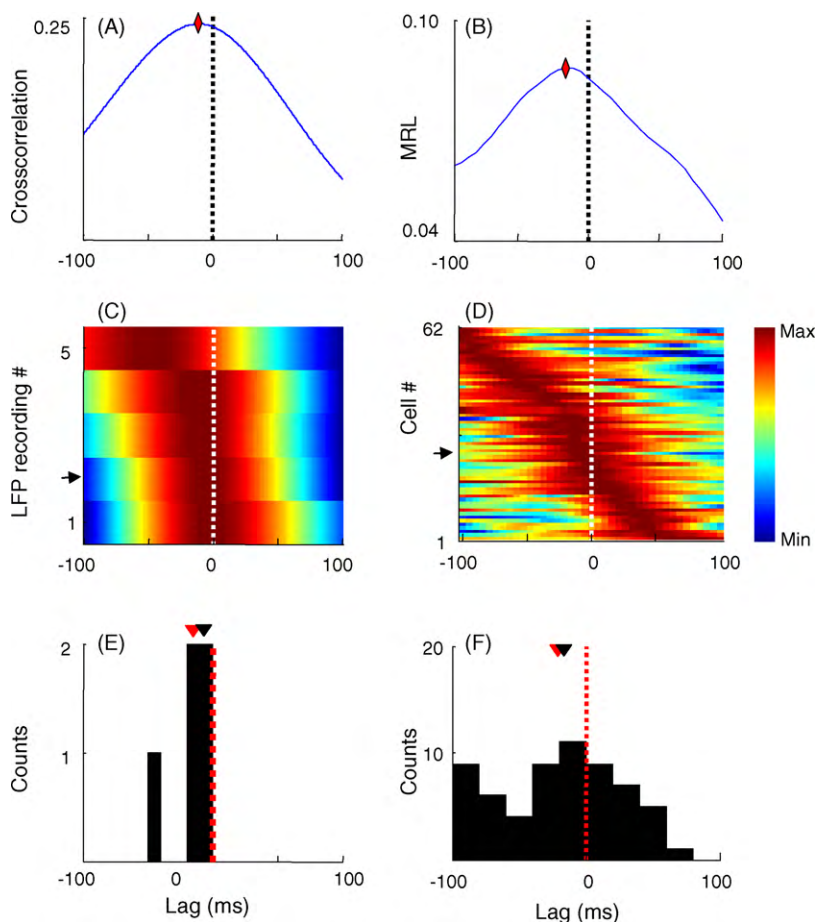


Fig. 3. Estimation of lags between dHPC and mPFC using the amplitude cross-correlation and phase locking methods. (A) dHPC–mPFC lag estimated from a single recording session by crosscorrelating the amplitudes of theta-filtered traces. Note that the peak occurs at a negative lag, indicating that the dHPC leads the mPFC in the theta-range. (B) Effect of shifting an mPFC spike train from a single unit on the strength of phase-locking (MRL). The plot shows that this mPFC single unit phase locks best to dHPC theta of the past, in agreement with the directionality shown in (A). Diamonds denote the peaks in both (A) and (B). (C) Normalized color plots of amplitude cross-correlations from 5 recordings and phase-locking shifts from 62 mPFC single units (D). Warmer colors indicate higher cross-correlation peaks or greater phase-locking strength. Each row corresponds to a single LFP (C) or single unit (D) recording. Rows are arranged according to the peak lag. Arrows mark the rows representing the data shown in the upper panels. (E and F) Histograms showing the distribution of peak lags calculated with each method. The distribution of lags is significantly negative for both the amplitude cross-correlation (E) ($p < 0.05$, Wilcoxon rank sum test, mean lag -15.4 ± 7.9 ms, $n = 5$ recordings) and phase-locking method (F) ($p < 0.003$ for a paired Wilcoxon's rank sum test, mean lag -20.2 ± 5.8 ms, $n = 30$ recordings). Means and medians of the lag distributions are indicated, respectively, by black and red arrowheads.

at frequency f (represented as $|\pi i \leftarrow j(f)|$), was calculated as shown below, using custom MATLAB routines provided by B. Lau and A. Saez (Columbia University).

$$|\pi i \leftarrow j(f)| = \frac{|\bar{A}_{ij}(f)|}{\sqrt{\sum_k |\bar{A}_{kj}(f)|^2}},$$

where k varies from 1 to m (the number of time-series being modeled), $\bar{A}(f) = I - A(f)$, and I is the identity matrix. The denominator normalizes PDC for each frequency, such that PDC values fall between 0 and 1. PDC values at a given frequency attempt to estimate the strength of the directionality (or improvement in prediction) between the time-series $j(t)$ and $i(t)$ at that frequency.

2.8. Simulations with pink noise

2.8.1. Generation of pink noise

In order to verify if the amplitude cross-correlation method is more robust than PDC to artifacts induced by noise in the signal, simulations were used to test both methods by adding noise to the signals. Pink noise (power falls with $1/f$, where f is the frequency) was used instead of white noise (same power for all f), because noise in LFPs generally has a $1/f$ spectrum. Pink noise was generated by passing Gaussian zero-mean white noise through a $1/f$

filter. As expected, the power spectra $P(f)$ of the resulting process was very well fit by the formula $P(f) = P(f_0) * (1/f)$, where $P(f_0)$ is the power at the lowest frequency. To compare the resilience of the current method with PDC to pink noise, two types of simulations were performed.

2.8.2. Addition of pink noise of the same amplitude to both LFPs

In the first simulation, two copies of the same theta-filtered trace were shifted relative to each other by 28 ms to induce a well-defined lag. 28 ms was chosen because it is the mean vHPC–mPFC lag calculated by the amplitude cross-correlation method. Pink noise was then generated independently and added to each signal in 500 simulations at each of ten different noise levels, corresponding to signal-to-noise power ratios of 1 to 0.2. The amplitude cross-correlation method and PDC were then used to calculate directionality between the signals for each simulation.

2.8.3. Addition of pink noise of different amplitudes to the leading LFP

For the second simulation experiment, two theta-filtered segments of simultaneously recorded vHPC and mPFC signals were used. As indicated by the arrows in Fig. 7A, in these traces, the vHPC

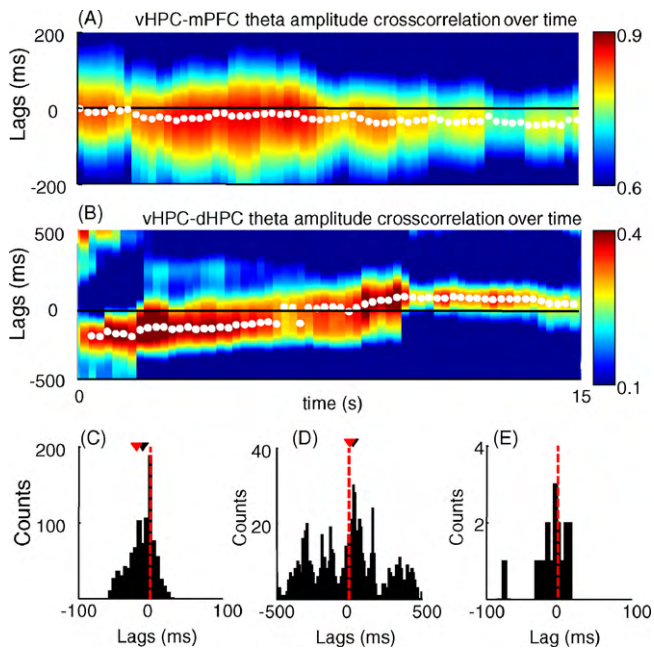


Fig. 4. Application of the amplitude cross-correlation method to bidirectionally connected areas. (A) A representative 15-second theta amplitude cross-correlation over time for the vHPC–mPFC is shown. Warmer colors correspond to higher cross-correlation values and white points mark the peak of the cross-correlation for each time window. Note that consistent with the existence of a unidirectional monosynaptic projection from the vHPC to the mPFC, the peaks of the cross-correlation fall primarily on negative lags for the vHPC–mPFC cross-correlation, indicating that the vHPC leads the mPFC. (B) Same as (A), but for simultaneously recorded vHPC and dHPC traces. In agreement with the existence of bidirectional monosynaptic projections between the vHPC and dHPC, at different time points the cross-correlation peaks at positive or negative lags, presumably reflecting periods in which the vHPC is leading or lagging relative to the dHPC, respectively. (A and B) Cross-correlations were calculated in 8 s windows with 97% overlap between successive windows. (C) Histogram of the lags at which the cross-correlation peaks occur for the entire 10 min recording from which the data plotted in (A) was obtained show that the distribution of vHPC–mPFC lags has a negative mean ($p < 0.005$, Wilcoxon's test), while the distribution of vHPC–dHPC lags (D) is not significantly different from zero ($p < 0.72$, Wilcoxon's test). (C and D) Each count of the histogram refers to one 8 s window in which the cross-correlation was computed. The distribution of the mean vHPC–dHPC lag for each animal is shown in (E), and it is not significantly different from zero. Each count in the histogram corresponds to the mean lag of one animal. In (C–E), red and black arrowheads indicate the means and medians of the distribution, respectively.

clearly leads the mPFC. In each simulation, pink noise of small but fixed amplitude was added to the mPFC signal, while noise traces of varying amplitudes were added to the vHPC. The power of the noise added to the vHPC was varied from 0.1 to 4-fold the power of the noise added to the mPFC. For each of the six noise levels added to the vHPC, 500 simulations were run, in which both vHPC and mPFC noise was newly generated. After adding noise to the signals, the directionality between the two signals was calculated using the amplitude cross-correlation method and PDC.

2.9. Statistics

Wilcoxon's two-tailed signed-rank non-parametric test was used throughout. $p < 0.05$ was considered statistically significant.

2.10. Histology

Upon the completion of recording, animals were deeply anesthetized and electrolytic lesions were made to determine the position of the electrode tips. Lastly, animals were perfused with formalin. Brain sections were mounted on slides to visualize and photograph lesions.

3. Results

3.1. Computation of the lag between two LFPs by amplitude cross-correlation

Each step of the method used to estimate the lag between two brain regions from LFPs is illustrated in Fig. 1. Examples of simultaneously recorded traces from the vHPC and mPFC are shown in Fig. 1A and B. The middle panels (Fig. 1C and D) show the same traces after filtering the signals in the theta-range (7–12 Hz). This frequency range was chosen because it is the predominant hippocampal oscillation in the awake moving rodent (Buzsáki, 2002), and previous reports have suggested that hippocampal activity in the theta-range may be propagated to the mPFC (Siapas et al., 2005). Note that the amplitude of the theta-filtered traces varies considerably across time. In order to obtain a reliable measure of the signal's amplitude with high temporal resolution, the Hilbert transform of the LFPs were calculated. The absolute magnitude of the Hilbert transform provides the amplitude of the signal for each time point at which the LFP was recorded. The amplitude traces (Fig. 1C and D, grey traces) reveal the existence of a temporal microstructure that reflects increasing and decreasing theta oscillation power and is not readily apparent in the raw LFPs. Careful visual inspection of Fig. 1E suggests that the peaks in vHPC theta amplitude precede those of the mPFC. Accordingly, the cross-correlation of the two amplitude traces (Fig. 1F, grey trace) has a peak at a negative lag, indicating that theta activity in the vHPC leads that in the mPFC.

Cross-correlations of the amplitudes of the theta-filtered LFPs from 17 mice are shown in Fig. 2. Note that most of the curves peak at a negative lag (dark red bands in Fig. 2A). The mean of the distribution was negative (mean lag = -28.1 ± 16.7 ms, median = -9.5 ms) and significantly different from zero ($p < 0.01$, Wilcoxon's rank sum test). The narrow peak of the vHPC–mPFC lag distribution (Fig. 2A, lower panel) suggests that the estimate of the lag calculated through this method is consistent across animals. Although the distribution of lags appears to be bimodal, it is still significantly negative if the points with large negative lags (< 50 ms) are excluded. Finally, there was no correlation of the computed lag with the layer from which vHPC signals were recorded (data not shown).

3.2. Test of the amplitude cross-correlation method

To further test the validity of this method, the lag estimated with the amplitude cross-correlation method was compared with the lag computed through a previously published method (Siapas et al., 2005; Sigurdsson et al., 2010) that uses spikes and LFPs. In this method, the strength of phase locking of mPFC spikes to vHPC theta oscillations is calculated after shifting the spike train both forwards and backwards in time. If neural activity in the vHPC leads that in the mPFC, phase locking of mPFC spikes to vHPC field potentials would be maximal when the spike train is shifted slightly to the past. Indeed, in line with previous reports (Siapas et al., 2005; Sigurdsson et al., 2010), mPFC spikes phase locked to hippocampal theta oscillations more robustly when shifted to the past ($n = 30$ multiunit recordings from 12 mice, mean shift = -24.5 ± 15.7 ms, $p < 0.05$, Wilcoxon's rank sum test). Furthermore, the vHPC–mPFC lag calculated with the phase locking and the cross-correlation methods are not significantly different from each other or from the conduction delay (-16 ms) of the vHPC–mPFC pathway (Thierry et al., 2000) ($p > 0.05$, Wilcoxon's signed-rank test) indicating that both methods may provide valid estimates of the lag between the vHPC and the mPFC.

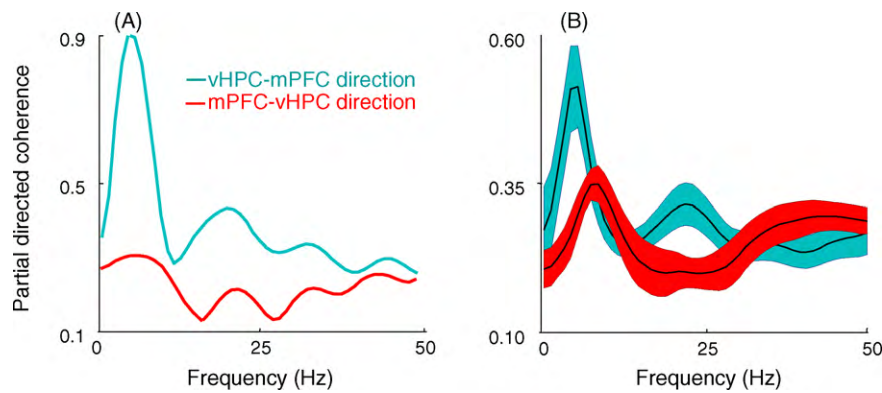


Fig. 5. Partial directed coherence indicates that vHPC is leading the mPFC in the theta-range. (A) Representative example of PDC on simultaneously recorded vHPC and mPFC LFPs from a single session. Note the prominent peak in the theta-range of the PDC in the vHPC to mPFC direction. (B) Average PDC across animals is plotted. Note that in the theta-range the predominant direction of flow, as indicated by higher PDC values, is in the vHPC to mPFC direction. Shaded areas indicate SEMs.

3.3. The amplitude cross-correlation method is frequency-specific

The above results indicate that the amplitude cross-correlation method can estimate lags that are in good agreement with the phase locking method and with the conduction delay of the vHPC–mPFC pathway. To further evaluate the amplitude cross-correlation method, vHPC–mPFC lags were calculated across multiple frequency ranges. We applied the amplitude cross-correlation method after band-pass filtering the LFPs in the delta (1–4 Hz), theta (7–12 Hz), low gamma (30–50 Hz) and high gamma (50–100 Hz) frequency bands (Fig. 2C). Strikingly, a consistent and significant lag between the vHPC and the mPFC was only found at the theta-range.

3.4. Applying the amplitude cross-correlation to areas that are bidirectionally connected

The above results suggest that the amplitude cross-correlation method can be applied to areas that have a unidirectional monosynaptic connection. However, often researchers wish to study the directionality of functional connectivity in areas that have indirect or bidirectional connections. To explore whether the method can be generalized to these situations, we applied it to simultaneous recordings from the dHPC and mPFC, which are indirectly connected, and from the vHPC and dHPC, which are directly and bidirectionally connected.

The dHPC and mPFC are bidirectionally and indirectly connected through the rhinal cortices, the vHPC and the *nucleus reuniens* of the thalamus (Burwell and Witter, 2002; Hoover and Vertes, 2007). In the dHPC–mPFC dataset, LFPs were sampled from the dHPC and single units and LFPs were recorded through stereotrodes from the mPFC. An example dHPC–mPFC theta amplitude cross-correlation (Fig. 3A) shows that the peak occurs at a negative lag, suggesting that dHPC activity in the theta-range leads the mPFC. Similar profiles were found in each of five animals (Fig. 3B), and the population of lags obtained by this method (Fig. 3C) has a significantly negative mean ($p < 0.05$, Wilcoxon's test, mean lag = -15.4 ± 7.9 ms). As in the vHPC–mPFC dataset, we also calculated lags using the spike-shift method in the same mice. A representative example of the effect of shifting the spike train relative to the dHPC LFP on the phase-locking strength (Fig. 3B) shows that this unit phase locks more robustly to dHPC theta the past (optimal lag = -17 ms). The distribution of optimal lags of for all the single units recorded (Fig. 3D and F) reveals that the mean of this distribution is significantly negative ($p < 0.01$, Wilcoxon's test, mean lag = -20.2 ± 5.8 ms), in rough agreement with the lag calculated by the amplitude cross-correlation method and with previous

reports (Siapas et al., 2005; Sigurdsson et al., 2010). These data suggest that the amplitude cross-correlation method can be used to estimate the directionality of functional connectivity and the lag across brain areas that have indirect, bidirectional connections.

In an awake-behaving rodent, the predominant direction of information flow is thought to be from the hippocampus to the cortex (Sigurdsson et al., 2010). Accordingly, we found that in both vHPC–mPFC and in the dHPC–mPFC datasets, hippocampal activity leads the mPFC. However, in many cases the predominant direction of information flow between the brain areas recorded may change across time (Gregoriou et al., 2009). To study the results of applying the method in such cases, we recorded LFPs simultaneously from the vHPC, dHPC and mPFC, and calculated lags from successive windows across time. In line with the findings described above for the vHPC–mPFC recordings, a negative lag in the theta-filtered amplitude cross-correlation was present across time (Fig. 4A), with peaks occurring consistently at negative lags (mean lag -12.5 ms, $p < 0.001$, Wilcoxon's test; Fig. 4A and C). Conversely, the lag between vHPC and dHPC varied across time, between positive and negative lags (Fig. 4B). The mean lag between these areas connected monosynaptically and bidirectionally did not differ from zero across time (Fig. 4D, mean lag -3.3 ms in this example animal; histogram has one count per time window) and across animals (Fig. 4E, mean lag -8.8 ± 10.2 ms, histogram has one count per animal). This result suggests that when applied to bidirectionally connected areas, the amplitude cross-correlation method can reveal the temporal dynamics of the directionality of the circuit, by indicating which area is leading in a given time window.

In summary, these results demonstrate the utility of the amplitude cross-correlation method in examining bidirectionally and indirectly connected brain regions, and further demonstrate its utility in examining how directionality may change across time.

3.5. Comparison of amplitude cross-correlation method with partial directed coherence

In order to compare the results of the amplitude cross-correlation with an existing method that calculates directionality with LFPs, we applied partial directed coherence (PDC) to the raw signals of the vHPC–mPFC dataset. Consistent with the results obtained using amplitude cross-correlation and phase-locking methods, PDC values averaged across animals had peaks in the theta-frequency range, and these peaks were significantly greater for the vHPC to mPFC (Fig. 5B) direction than vice-versa. This result is also in line with a previous analysis of hippocampal–cortical connectivity with PDC in anesthetized rats (Taxidis et al., 2010).

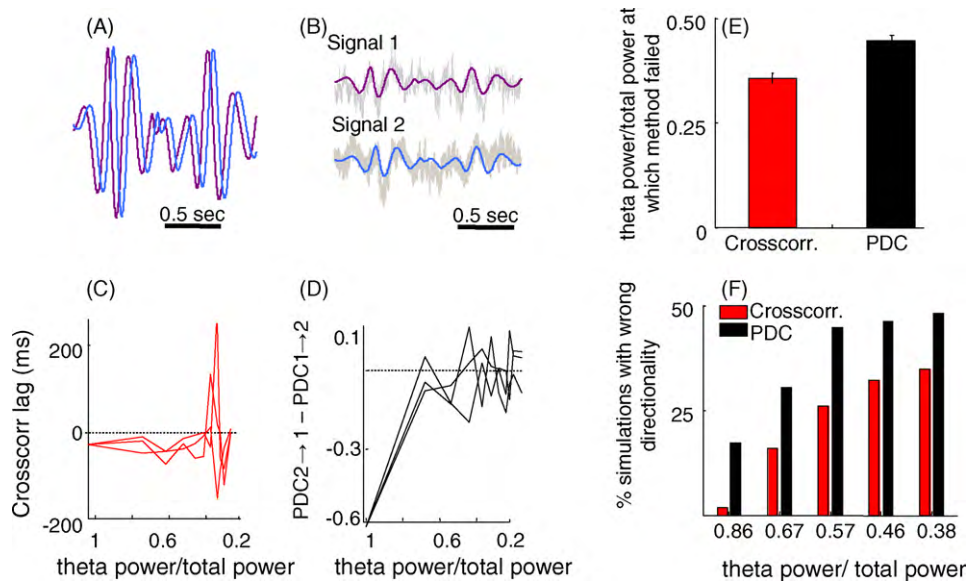


Fig. 6. Partial directed coherence is more sensitive to noise than the amplitude cross-correlation method. (A) Two signals were created from the same two-second segment of vHPC theta-filtered trace. One signal was shifted relative to the other by 28 ms (this is the mean delay between the vHPC and the mPFC calculated by the amplitude cross-correlation method). Thus, the purple trace leads the blue trace with a lag of 28 ms. (B) Pink noise was generated randomly and added to both signals. Ten different levels of noise were added, such that the fraction of theta power relative to total power in the signals after adding noise was varied from 1 to 0.2. In the example shown, theta power/total power = 0.67. (C) The amplitude cross-correlation method was applied to verify the directionality after adding pink noise to both signals, in 500 simulations in which noise was newly generated, at 10 different amplitudes. Calculation of the lag by the cross-correlation method for three representative simulations is shown. Points with negative lags have the expected directionality. The cross-correlation method fails only when high levels of noise are added, as shown by the points with positive lags with low theta power to total power ratios. Note that the x-axis is reversed, such that higher values (high signal to noise ratios) are on the left. (D) Same as in (C), but for PDC calculated from the identical simulations. Correct directionality is reflected as negative values on the y-axis. PDC1 → 2 indicates PDC in the purple trace in A causing the blue trace. Note that this method does not consistently indicate that the purple trace leads the blue trace even after adding only moderate amounts of noise to the signal. (E) Mean noise level at which each method first failed, averaged across 500 simulations. The PDC method on average fails at lower noise levels than the amplitude cross-correlation method ($p < 0.0001$, rank sum test). (F) For five noise levels, the percentage of simulations in which the wrong directionality was calculated is shown. At every noise level PDC had a significantly higher failure rate than the amplitude cross-correlation method ($p < 0.05$, Fisher's exact test). All PDC values shown are averages across the theta-range.

To compare the two methods, we used simulations to analyze how sensitive each method is to noise. We first studied the amount of noise that must be added to a signal to prevent each method from identifying the directionality between two signals where the underlying directionality is known. To this end, a 2-s segment of a randomly chosen vHPC recording was filtered in the theta-frequency range. A second signal was then created by shifting the original signal by 28 ms, resulting in two identical signals separated by a defined lag (Fig. 6A). A 28 ms shift was chosen because it is the average lag calculated for the vHPC–mPFC pathway with the cross-correlation method. Varying levels of pink noise were then generated (10 levels of noise) and added to both signals in 500 independently generated simulations for each noise level (Fig. 6B). For each simulation, the directionality was then calculated using both PDC and the amplitude cross-correlation methods. Fig. 6C and D shows the results of three example simulations at each of the 10 noise levels (the same three simulations are shown in both panels). In these examples, the amplitude cross-correlation method correctly quantified directionality and lag at noise levels at which PDC failed to identify directionality. In general, across the population of simulations, PDC failed to identify the correct directionality at lower levels of noise than did the amplitude cross-correlation method ($p < 0.05$, Fisher's exact test; Fig. 6E and F).

Previous reports have demonstrated that PDC can incorrectly find directionality between signals when the two signals have different variances (Taxidis et al., 2010; Winterhalder et al., 2005). We therefore asked whether the cross-correlation method is similarly sensitive to the presence of different levels of noise across two signals. To do so, we first selected a 1-min sample of simultaneously recorded vHPC and mPFC signals. Both signals were filtered for the theta-range. Short segments of the samples are shown in Fig. 7A (Black traces). Note that the vHPC clearly leads the mPFC. Next, we

added a constant and modest amount of pink noise to the mPFC signal, while adding varying amounts of noise to the vHPC signal (grey traces). The directionality between the two signals was then calculated using both PDC and amplitude cross-correlation methods. Noise was independently generated 500 times for each of the six levels of noise shown. Regardless of the ratio of noise added to the two signals, the amplitude cross-correlation method consistently indicated that the mPFC follows the vHPC. While increasing vHPC noise increased the variability in the detected lags, there was no effect of increasing vHPC noise on the mean estimated lag (Fig. 7B, $p < 0.72$, one-way ANOVA). In dramatic contrast to the amplitude cross-correlation method, the directionality indicated by PDC was dependent on the amount of noise added to the vHPC signal (Fig. 7C). When the ratio of noise added to the vHPC signal was smaller or equal to that added to the mPFC, PDC correctly identified the directionality. However, when the noise added to the vHPC signal was larger than that added to the mPFC, PDC incorrectly calculated the reverse directionality. These results indicate that PDC, but not the amplitude cross-correlation, is greatly biased if different signals have different amounts of noise in them. In such cases, PDC results suggest that the cleaner signal leads the noisier signal, independently of the underlying directionality, in line with previous reports (Taxidis et al., 2010; Winterhalder et al., 2005). The amplitude cross-correlation is less precise but equally accurate with the addition of differential amounts of noise.

While the PDC and amplitude cross-correlation methods both show appropriate directionality in our dataset, these simulations suggest that PDC may be less robust to noise than the amplitude cross-correlation. Interestingly, a similar susceptibility to noise was found with Granger causality, another multivariate autoregressive model-dependent measure of directionality, in both types of simulations (Supplemental Fig. 2).

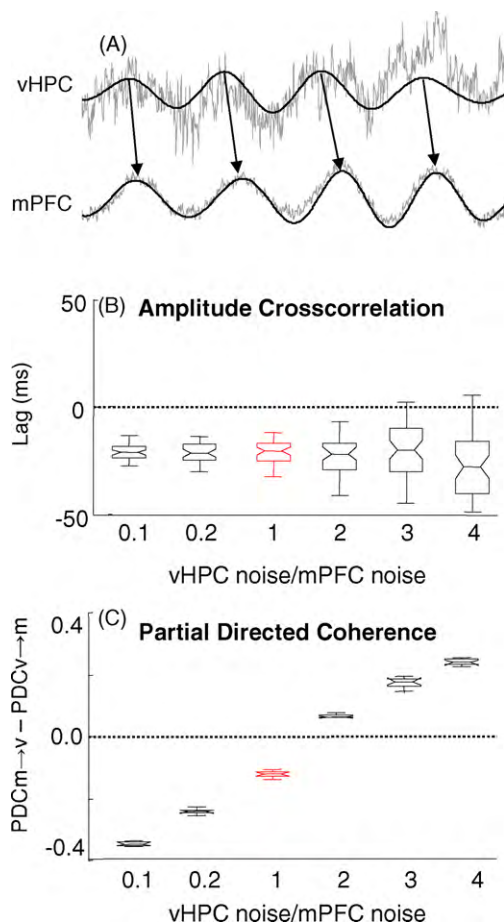


Fig. 7. Partial directed coherence, but not the cross-correlation method, is biased if one of the signals has different noise levels than the other signal. (A) Half-second segments of 1-min long signals are shown. Black traces are theta-filtered traces from the vHPC and mPFC. Note that the vHPC leads the mPFC in these traces, as indicated by the black arrows. Grey traces were obtained after adding different levels of pink noise to each of the filtered signals. In this example, the noise added to the vHPC is four-fold greater than the mPFC noise. Across the simulations, the noise added to the mPFC remained constant, while the noise added to the vHPC was varied from 0.1 to 4-fold of the noise added to the mPFC. Noise was generated in six increasing amplitudes in 500 simulations and added to the vHPC. The signals were then analyzed by the amplitude cross-correlation method (B) and PDC (C). (B) Boxplot shows the median lags calculated by the amplitude cross-correlation method after different amounts of noise were added to the vHPC signal while keeping constant the amplitude of the noise added to the mPFC signal. The lag calculated by the amplitude cross-correlation method remains negative, indicating that mPFC follows the vHPC, even when the amount of noise added to the vHPC is greater than the mPFC noise. Boxplots show the mean lag and the 25th and 75th percentiles of the distributions. Whiskers indicate the range. (C) PDC was calculated for the same simulations used in (B). Correct directionality (vHPC leading) is represented as negative values on the y-axis (PDC_{v→m} greater than PDC_{m→v}, where v and m stand for vHPC and mPFC, respectively). PDC indicates that the vHPC leads the vHPC only when the vHPC is equally or less noisy than the mPFC. Note that PDC consistently indicates that the less noisy signal leads the noisier signal, regardless of the underlying directionality. In (B) and (C), the condition of equivalent noise levels in the two signals (vHPC noise/mPFC noise = 1) is shown in red. All PDC values shown are averages across the theta-range.

4. Discussion

We report a novel, mathematically straightforward method to calculate the lag in neural activity between two brain areas utilizing multi-site LFP recordings. The method does not require sampling of spikes. The present data indicate that the amplitude cross-correlation method can estimate the principal direction of functional connectivity between two brain regions and provide a reliable and consistent estimate of the lag between them. We

further demonstrate that in the hippocampal–prefrontal circuit, the lag calculated with the amplitude cross-correlation method is consistent with lags calculated through phase locking (Siapas et al., 2005), and with the experimentally determined conduction delay between the two areas (Thierry et al., 2000). Each method indicates that the direction of functional connectivity is from the vHPC to the mPFC. We further demonstrate that the amplitude cross-correlation method can detect frequency-specific connectivity, and can identify indirect and bidirectional connectivity as well as monosynaptic, unidirectional connectivity. Finally, we demonstrate that the amplitude cross-correlation method is relatively robust to added noise, when compared to existing methods of determining directionality from LFPs.

The principal advantage of the method lies in its practicality and theoretical simplicity. Estimating directionality and lag using exclusively LFPs is of great value as LFP recordings are much more straightforward to obtain than single unit spike data, especially simultaneously from multiple areas. While widely used and straightforward methods for calculating directionality with LFPs exist, they are limited. Direct cross-correlation of filtered or raw signals, for example, can fail to produce reliable lag estimates even in the presence of large, synchronous oscillations. This failure results from the potential for variable phase offsets in the rhythms between the two LFPs (Supplemental Fig. 3), which does not contaminate the amplitude cross-correlation. Instantaneous phase difference is another method that has been used to calculate directionality (Gregoriou et al., 2009), although it can be unclear how to convert a phase difference in degrees into a lag in milliseconds, as each area may oscillate in a different peak frequency. For example, in the current data set, the mean phase difference between mPFC and vHPC is 0.22 radians. However, as mPFC and vHPC have different mean theta frequencies, converting this phase difference to a time lag will produce different results depending on whether mPFC or vHPC mean theta frequency is used to calculate the lag, and there is no justification for choosing one of the peak frequencies over the other. Furthermore, as phase is a circular measure, phase difference calculations cannot determine if a difference of 60° indicates that signal 1 is leading signal 2 by 60° or if it is lagging signal 2 by 300°.

Mathematically more sophisticated methods to calculate the directionality of functional connectivity across brain areas with LFPs such as Granger causality and PDC also exist. These methods, while useful, are conceptually complex and not yet widely accepted in the literature. Moreover, Granger causality and PDC are not typically used to obtain lag estimates. As PDC and Granger are based on fitting multivariate regression models to a dataset, the data must conform to several constraints, such as having stationarity and a Gaussian distribution. Furthermore, the residuals of the model must describe a white noise process (Cadotte et al., 2010; Gregoriou et al., 2009). In contrast, the data does not need to fulfill these conditions for the cross-correlation method to be applicable. In fact, the non-stationarity in power over time is precisely what is used to estimate directionality in the cross-correlation method. There are, however, three main advantages of PDC and Granger causality relative to our method. First, these methods can be applied to the entire frequency range at once, while the amplitude cross-correlation needs to be separately calculated for each frequency range and is therefore somewhat dependent on the choice of filter boundaries and filter type. Second, PDC and Granger causality provide a measure of the degree of functional connectivity in both directions, whereas the amplitude cross-correlation method only provides the overall directionality for a pair of areas. Third, if PDC and Granger causality are applied to datasets in which three or more areas were recorded, they will provide measures of functional connectivity for all possible brain area pairs in a single step. The cross-correlation method

would have to be applied in separate steps for each pair of brain areas.

In order to compare PDC and the cross-correlation method we applied both to our dataset. Importantly, both methods were in good agreement with each other and with the spike-shift phase-locking method, indicating that on average, across animals, the vHPC leads the mPFC. However, simulations with noise indicate that the amplitude cross-correlation method may be more robust to certain types of noise, compared to PDC. First, the cross-correlation method is less sensitive when pink noise of equal amplitude is added to both signals. This result suggests that PDC is more likely to not find directionality if the frequency band of interest has low power in one of the signals, as is often the case. Second, if one of the signals has higher levels of noise than the other, PDC, but not the cross-correlation method, tends to indicate that the less noisy signal leads the noisier signal, as reported elsewhere (Taxidis et al., 2010; Winterhalder et al., 2005). This is a relevant point to consider, as the ratio of power in the frequency band of interest to total power can be widely different across brain areas. Although these simulations suggest the amplitude cross-correlation method is more resilient to digitally added pink noise, further studies are needed to verify if this is the case in biological signals recorded in noisy conditions.

Although the amplitude cross-correlation method may be applicable to many datasets, it does have limitations that will potentially restrict its utility. For example, it may not always be clear on which frequency band the method should be applied. In these cases, several steps can be taken to find the relevant frequency band at which directionality is present and is detectable. The most straightforward one is to use bands defined by prominent peaks in the coherence or power spectra in the brain areas of interest. Knowledge of previous literature may also be of help. For instance, the observations reported here rely on theta-frequency synchrony, which is known to be prominent between the hippocampus and its targets. However, gamma-frequency synchrony can occur across functionally connected cortical regions (Hermer-Vazquez et al., 2007), raising the possibility that the method reported here may be of use in evaluating cortico-cortical functional connectivity if applied to this frequency range. Lastly, in the absence of candidate frequency bands it is possible to apply the amplitude cross-correlation method in successive frequency windows over a broad range, for example, from 1 to 100 Hz. Indeed, we have used such an unbiased approach to show that the theta-range is the only frequency band with a consistent and significant lag between the vHPC and the mPFC across animals in our dataset (Supplemental Fig. 4).

It is important to note that like any method to estimate lag or functional connectivity between areas, the amplitude cross-correlation method does not conclusively demonstrate a causal relationship. These analyses indicate only that mPFC activity follows vHPC activity, and does not prove that mPFC activity is driven by vHPC activity. Neither the current method nor the others mentioned here rule out the possibility that a third area drives both the vHPC and mPFC with different lags. Nonetheless, evidence of monosynaptic, unidirectional connectivity between the vHPC and mPFC has been demonstrated both anatomically (Parent et al., 2009; Verwer et al., 1997) and physiologically, (Thierry et al., 2000), consistent with the notion, suggested by the amplitude cross-correlation method, that the vHPC drives mPFC activity.

Acknowledgements

We thank members of the Gordon and Hen labs for helpful discussions of the experimental analysis. We thank A. Wiltschko for providing analysis scripts to compute the Gabor Transform. We also thank B. Lau and A. Saez for providing the scripts

to compute PDC and advice on how to implement it. This work was supported by grants to J.A.G. from the NIMH (K08 MH098623 and R01 MH081958). J.A.G. is also the recipient of a NARSAD/Bowman Family Foundation Young Investigator Award and an APIRE/GlaxoSmithKline Young Faculty Award.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jneumeth.2010.06.019.

References

- Adhikari A, Topiwala MA, Joshua AG. Synchronized activity between the ventral hippocampus and the medial prefrontal cortex during anxiety. *Neuron* 2010;65(2):257–69.
- Alonso JM, Martinez LM. Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat Neurosci* 1998;1:395–403.
- Astolfi L, Cincotti F, Mattia D, Marciani MG, Baccala LA, de Vico Fallani F, et al. Assessing cortical functional connectivity by partial directed coherence: simulations and application to real data. *IEEE Trans Biomed Eng* 2006;53:1802–12.
- Baccala LA, Sameshima K. Partial directed coherence: a new concept in neural structure determination. *Biol Cybern* 2001;84:463–74.
- Burwell RD, Witter MP. Basic anatomy of the parahippocampal region in monkeys and rats. New York: Oxford University Press; 2002.
- Buzsaki G. Theta oscillations in the hippocampus. *Neuron* 2002;33:325–40.
- Cadotte AJ, Demarse TB, Mareci TH, Parekh M, Talathi SS, Hwang DU, et al. Granger causality relationships between local field potentials in an animal model of temporal lobe epilepsy. *J Neurosci Methods* 2010.
- DeCoteau WE, Thorn C, Gibson DJ, Courtemanche R, Mitra P, Kubota Y, et al. Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. *Proc Natl Acad Sci USA* 2007;104:5644–9.
- Gourevitch B, Bouquin-Jeannes RL, Faucon G. Linear and nonlinear causality between signals: methods, examples and neurophysiological applications. *Biol Cybern* 2006;95:349–69.
- Gregoriou GG, Gotts SJ, Zhou H, Desimone R. High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 2009;324:1207–10.
- Hermer-Vazquez R, Hermer-Vazquez L, Srinivasan S, Chapin JK. Beta- and gamma-frequency coupling between olfactory and motor brain regions prior to skilled, olfactory-driven reaching. *Exp Brain Res* 2007;180:217–35.
- Holdefer RN, Miller LE, Chen LL, Houk JC. Functional connectivity between cerebellum and primary motor cortex in the awake monkey. *J Neurophysiol* 2000;84:585–90.
- Hoover WB, Vertes RP. Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Struct Funct* 2007;212:149–79.
- Jones MW, Wilson MA. Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol* 2005;3:e402.
- Le Van Quyen M, Foucher J, Lachaux J, Rodriguez E, Lutz A, Martinerie J, et al. Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *J Neurosci Methods* 2001;111:83–98.
- Lindsey BG, Hernandez YM, Morris KF, Shannon R. Functional connectivity between brain stem midline neurons with respiratory-modulated firing rates. *J Neurophysiol* 1992;67:890–904.
- Parent MA, Wang L, Su J, Netoff T, Yuan LL. Identification of the hippocampal input to medial prefrontal cortex in vitro. *Cereb Cortex* 2009.
- Paz R, Bauer EP, Pare D. Measuring correlations and interactions among four simultaneously recorded brain regions during learning. *J Neurophysiol* 2009;101:2507–15.
- Popa D, Duvarci S, Popescu AT, Lena C, Pare D. Coherent amygdalocortical theta promotes fear memory consolidation during paradoxical sleep. *Proc Natl Acad Sci USA* 2010;107:6516–9.
- Siapas AG, Lubenov EV, Wilson MA. Prefrontal phase locking to hippocampal theta oscillations. *Neuron* 2005;46:141–51.
- Sigurdsson T, Stark KL, Karayiorgou M, Gogos JA, Gordon JA. Impaired hippocampal-prefrontal synchrony in a genetic mouse model of schizophrenia. *Nature* 2010;464:763–7.
- Snider RK, Kabara JF, Roig BR, Bonds AB. Burst firing and modulation of functional connectivity in cat striate cortex. *J Neurophysiol* 1998;80:730–44.
- Taxidis J, Coomber B, Mason R, Owen M. Assessing cortico-hippocampal functional connectivity under anesthesia and kainic acid using generalized partial directed coherence. *Biol Cybern* 2010;102:327–40.
- Thierry AM, Gioanni Y, Degenetais E, Glowinski J. Hippocampo-prefrontal cortex pathway: anatomical and electrophysiological characteristics. *Hippocampus* 2000;10:411–9.
- Verwer RW, Meijer RJ, Van Uum HF, Witter MP. Collateral projections from the rat hippocampal formation to the lateral and medial prefrontal cortex. *Hippocampus* 1997;7:397–402.
- Winterhalder M, Schelter B, Hesse W, Schwab K, Leistriz L, Klan D, et al. Comparison directed of linear signal processing techniques to infer interactions in multivariate neural systems. *Signal Process* 2005;85:2137–60.