Analyzing Neural Time Series Data: Theory and Practice

Mike X Cohen

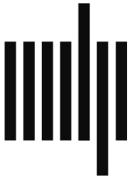

The MIT Press

From The MIT Press

MITCogNet

This book was set in ITC Stone Serif Std by Toppan Best-set Premedia Limited, Hong Kong. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Cohen, Mike X, 1979–
Analyzing neural time series data : theory and practice / Mike X Cohen.
pages cm.—(Issues in clinical and cognitive neuropsychology)
Includes bibliographical references and index.
ISBN 978-0-262-01987-3 (hardcover : alk. paper) 1. Neural networks (Neurobiology) 2. Neural networks (Computer science) 3. Computational neuroscience. 4. Artificial intelligence—Biological applications. I. Title.
QP363.3.C633 2014
612.8'2339—dc23
2013016461

10 9 8 7 6 5 4 3 2 1

26 Phase-Based Connectivity

Many of the concepts and procedures for measuring phase-based connectivity are similar to those for ITPC, so understanding the material presented in chapter 19 will help you to understand the material in this chapter.

26.1 Terminology

There are several terms that are used in the literature to describe phase-based connectivity, including phase-locking value/statistic/factor, phase synchronization, and phase coherence. Arguably, these are suboptimal terms, partly because they are interpretations of results rather than descriptions of methods and partly because the same terms are used to indicate different methods (for example, the term phase-locking value can be used either for ITPC or for connectivity, and yet these two analyses have very different meanings and interpretations). "Phase correlation" is sometimes used, but this is a poor term because a correlation indicates a linear bivariate relationship that can take positive or negative values; phase-based connectivity measures are not linear, nor can they have negative values.

Here, the term intersite phase clustering (ISPC) is preferred. ISPC is a concise description of the method (clustering in polar space of phase angle differences between electrodes, voxels, or neurons) without any assumptions or interpretations regarding putative neurophysiological mechanisms being made. Other analysis-specific terms (e.g., phase-lag index) are introduced later in this chapter.

26.2 ISPC over Time

ISPC is similar to ITPC presented in chapter 19. Recall from equation 19.1 that ITPC is defined as the length of the average vector of many unit vectors whose phase angles are obtained by a point in complex space resulting from the convolution between a complex wavelet and the

data (or from applying the Hilbert transform to bandpass-filtered data). ISPC works in a similar fashion, but rather than taking the average of phase angles, you take the average of phase angle *differences* between electrodes over time (Lachaux et al. 2000; Mormann et al. 2000).

$$\text{ISPC}_f = \left| n^{-1} \sum_{t=1}^{n} e^{i(\phi_{xt} - \phi_{yt})} \right| \tag{26.1}$$

in which n is the number of time points, and ϕ_x and ϕ_y are phase angles from electrodes x and y at frequency f. The only difference between equation 26.1 and equation 19.1 is that equation 26.1 has a subtraction of phase angles from two electrodes rather than phase angles from one electrode. When implementing this equation in Matlab, remember to multiply the difference of the phase angles by the imaginary operator and not only the first phase angle. That is, write `exp(1i*(angles1-angles2))` and not `exp(1i*angles1-angles2)`.

To illustrate this method, it is useful to compute ISPC between two neighboring electrodes without applying any spatial filters. Here, electrodes Pz and P1 are used. Much of the apparent connectivity between these electrodes is due to volume conduction and therefore should not be interpreted in terms of connectivity between distinct brain regions. Nonetheless, using similar time series provides a visually intuitive learning exercise. Figure 26.1A shows bandpass-filtered time series from these two electrodes and their phase angles, and figure 26.1B shows the amplitude and phase angle differences between these two electrodes.

Figure 26.1C shows the individual phase angles from each electrode. You can see that the phase angles from the two electrodes are uniformly distributed throughout polar space. This is not surprising because there are many cycles at this frequency band in the time segment shown. However, when you subtract the phase angles from the two electrodes at each time point (figure 26.1D), you can see that the phase angle differences are nonuniformly distributed in polar space. In fact, they are strongly clustered around zero, which, in this case, reflects that much of the apparent connectivity is due to volume conduction.

Figure 26.2 shows that the actual phase lag between electrodes does not matter. What matters is that the phase lag is consistent across time. Thus, activity measured by pairs of electrodes that are 0 ms, 10 ms, or 100 ms lagged from each other can be equally strongly synchronized when measuring synchrony through ISPC.

ISPC (along with many other phase-based connectivity measures) is symmetric (that is, nondirectional). This means that ISPC A $\rightarrow$ B is the same as ISPC B $\rightarrow$ A. You can confirm this yourself by swapping the order of the electrodes in the online Matlab code to see that it has no effect on ISPC.

So far, ISPC was computed over a window of time from one trial. It is likely that you will want to examine changes in ISPC over time. There are two options for examining task-related

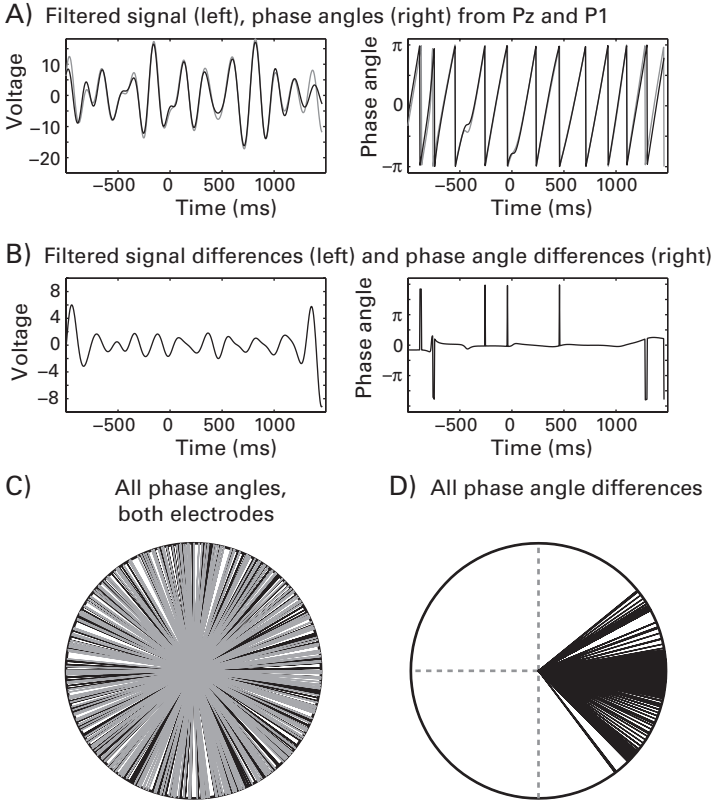

Figure 26.1

Illustration of phase angle differences between two electrodes over time. ISPC is defined as the length of the average vector of phase angle difference vectors (the vectors in panel D). Amplitude information, shown in the left plots of panels A and B, is not taken into consideration in this analysis. The apparent large jumps in the right-hand plot of panel B are due to small changes between -2π and $+2\pi$. This figure is actually the final frame of a movie that shows phase angle differences collecting over time. The online Matlab code will produce the movie.

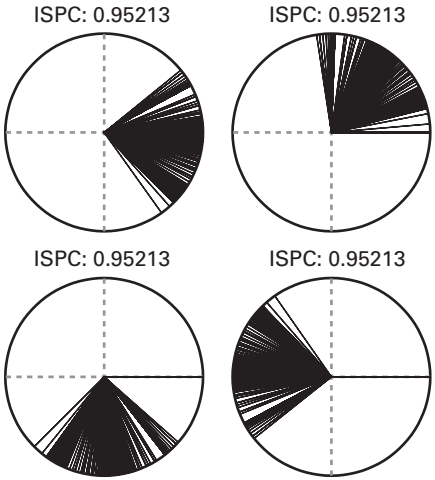

Figure 26.2
ISPC does not depend on phase values themselves, only the clustering of phase angle differences.

or time-varying ISPC. One option is to compute ISPC in sliding time segments, analogous to the way the FFT was computed in sliding time segments in the short-time FFT method (chapter 15). After ISPC has been computed in sliding time segments within a trial, the ISPC would be computed for each trial, and then the results can be averaged across trials (or correlated with a trial-varying behavior or experiment parameter; more on this later). Hereafter, this approach is called *ISPC-time*.

The selection of time segment length is partly related to the frequency and partly related to the task design. There is a trade-off between signal-to-noise ratio and temporal precision: longer time segments include more oscillation cycles and therefore will provide more robust and higher signal-to-noise estimates of ISPC-time. However, this comes at the expense of decreased temporal precision for time-varying and task-related modulations of ISPC-time. A segment length of at least three cycles (i.e., 1.5 before and after each time point) is recommended. If you have a task with long events such as a working memory delay or long stimulus presentation times, you could use more cycles. Because higher-frequency activity tends to have lower signal-to-noise ratio, and because higher-frequency cycles are, by definition, faster, you can have the number of cycles increase with increasing frequency, similar to how the number of cycles increases when creating Morlet wavelets. This option will modify the balance between temporal precision and signal-to-noise as a function of frequency. Thus, for example, you could use time segments corresponding to three to seven cycles for frequencies ranging from 4 to 60 Hz. It is also possible to keep the segment length fixed (e.g., at

300 ms), but this results in a variable number of cycles over frequencies. Figure 26.3 (plate 15) illustrates the effects of window segment length on ISPC-time results.

The second option for examining task-related or time-varying ISPC is to compute ISPC over trials instead of over time, which is discussed below.

26.3 ISPC-Trials

ISPC over trials (hereafter called *ISPC-trials*; the term "ISPC" is used when discussing both ISPC-time and ISPC-trials) is a similar but alternative method for assessing task-related

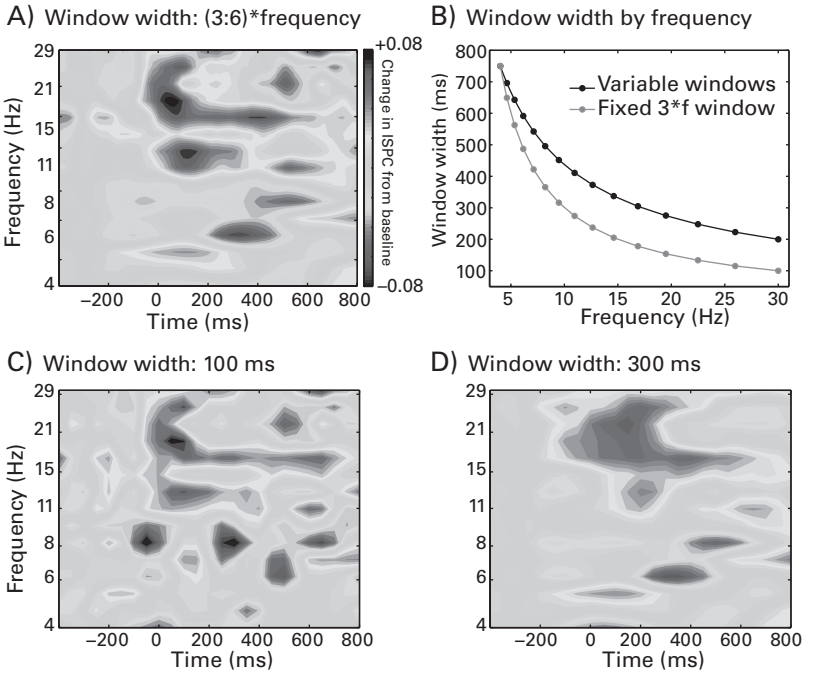

Figure 26.3 (plate 15)

ISPC-time between electrodes Fz and O1 using different time segment lengths. The results shown in panel A were obtained using a variable time segment length of three cycles at the lowest frequency (4 Hz) to six cycles at the highest frequency (30 Hz). The results in panels C and D were obtained using fixed time segment lengths for all frequencies. ISPC-time was computed over sliding time segments for each trial and then averaged across trials. All time-frequency plots have the same color scale (shown in panel A). ISPC strength from a pretrial baseline period of –400 to –200 ms was subtracted to highlight task-related effects. Panel B shows the length of the time segments used at each frequency when creating panel A (black line) and what the length of the time segments would be if a fixed ratio of three times the frequency (3*frequency or 3*f) had been used (gray line).

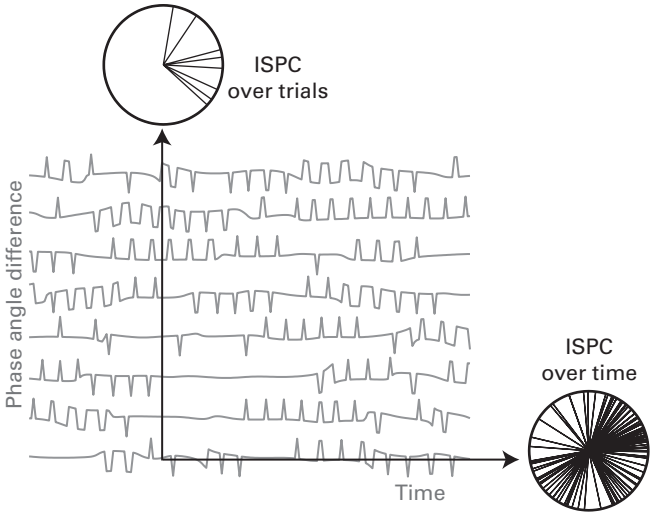

Figure 26.4
The conceptual difference between ISPC-time and ISPC-trials. Gray lines show time courses of phase angle differences over time for eight trials (each line is a trial). ISPC-time is computed from phase angle difference distributions over many time points for one trial, whereas ISPC-trials is computed from phase angle difference distributions at one time point over trials. Many connectivity analyses presented in this and subsequent chapters can be done over time or over trials.

phase-based connectivity. With ISPC-trials the assumption is slightly different than phase angle differences being clustered over time: the assumption is that connectivity produces a clustering of phase values at each time-frequency point relative to an experiment event over repeated trials (Lachaux et al. 1999). This is a subtle but important distinction. With ISPC-trials, the distribution of phase angle differences is generated at each time point over trials, whereas with ISPC-time, that distribution is computed at each trial over time points (figure 26.4). The distinction between computing connectivity over trials versus over time has implications both for analyses and for interpretations and will come up several times over the next few chapters; most connectivity analyses discussed in this book can be computed over time or over trials.

ISPC-trials add a stronger constraint, which is that phase angle differences must be consistent over trials. With ISPC-time, the preferred phase angle difference can be different on each trial as long as the amount of clustering is similar on each trial. With ISPC-trials, the preferred phase angle difference must be similar at each time-frequency point over trials. For example, if the four panels in figure 26.2 corresponded to phase angle differences over time

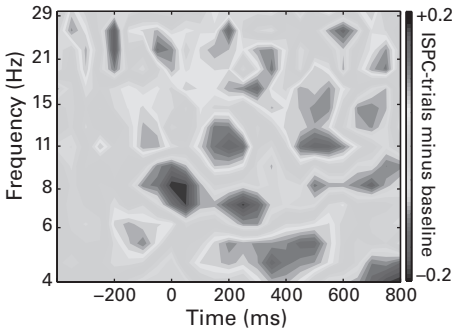

Figure 26.5 (plate 16)
Baseline-subtracted ISPC-trials using the same data as shown in figure 26.3 (plate 15).

for each of four different trials, ISPC-time would be very high, whereas ISPC-trials would be very low.

The equation for ISPC-trials is the same as equation 26.1, except that the *t* now refers to trial instead of time point, and *n* refers to the number of trials instead of the number of time points. An example result of ISPC-trials is shown in figure 26.5 (plate 16).

Should you compute ISPC-time or ISPC-trials? There are two advantages of ISPC-time. First, it is less sensitive to trial-to-trial jitter of experiment event timing (see figure 19.9) because even if the event timing is jittered, the timing of activity at different EEG electrodes will never be jittered relative to the timing of activity at other electrodes. Second, ISPC-time is insensitive to phase angle differences being different on different trials and can therefore be conceptualized as measuring total (non-phase-locked and phase-locked) rather than only phase-locked connectivity. For both of these reasons, ISPC-time is better suited for detecting high-frequency connectivity. The disadvantage of ISPC-time is that it requires longer within-trial time segments. This reduces temporal precision even more than was already done by time-frequency decomposition (as shown in figure 26.3B, plate 15). It may also be problematic if time segments extend into other trial events, the intertrial interval, or the trial buffer zones, which may contain edge artifacts. This problem becomes worse at lower frequencies, for which a window of three cycles could mean 1 s or longer. Thus, for lower frequencies, ISPC-time is better suited for experiments with long task events. In theory, it is possible to compute ISPC over both time and trials by concatenating all phase angle pairs over the time segment from all trials and computing one ISPC value. However, that may be a suboptimal strategy: it assumes that the preferred phase angle difference is the same over trials (which is a requirement of ISPC-trials but not ISPC-time) but still has the poor temporal precision of ISPC-time.

There are two advantages of ISPC-trials. First, it provides stronger evidence compared to ISPC-time for task-related modulations in connectivity because the connectivity must be in the same phase configuration on each trial. Second, ISPC-trials is computed at each time point individually, which means there is no loss of temporal precision beyond what was introduced by wavelet convolution or bandpass filtering. If you have hypotheses concerning the time course of connectivity over tens to hundreds of milliseconds, ISPC-trials should be preferred over ISPC-time. The two disadvantages of ISPC-trials are that it is sensitive to slight jitters in the uncertainty of the timing of experiment events, as shown in figure 19.9, and that it may fail to detect connectivity that occurs on each trial but with different phase values over trials. ISPC-trials cannot be performed on resting-state data.

ISPC-trials can be performed on results from any time-frequency decomposition method that provides phase values, including complex wavelet convolution, filter-Hilbert, short-time FFT, and multitaper. ISPC-time can be performed only on results from time-frequency decomposition methods that provide phase angle time series with the same temporal resolution as the original data. The two methods presented in this book that can be used for ISPC-time are complex wavelet convolution and filter-Hilbert.

26.4 ISPC and the Number of Trials

ISPC-trials, like ITPC, is sensitive to the number of trials used in the analyses. Figure 26.6 illustrates this point by showing ISPC-trials as a function of the number of trials included in the analysis. This figure is similar to figure 19.6A. More discussion of the relationship between the number of trials and ITPC was presented in section 19.3. Figure 26.6 also illustrates a strategy for determining whether you have enough trials for stable estimates of ISPC-trials within a condition. You can generate a figure like this by combining all trials from all conditions and noting the number of trials at which the lines appear to flatten for the frequencies you want to analyze. If the number of trials per condition exceeds this number, you probably have enough trials in that condition for a robust estimate of ISPC-trials. Based on the data used in figure 26.6, around 40 trials per condition should lead to stable estimates of ISPC-trials for most frequency bands. As with figure 19.6A, the results presented in figure 26.6 are based on one pair of electrodes and one subject and therefore should not be interpreted to indicate that 40 trials are always sufficient for ISPC-trials.

Three other ways of dealing with potentially low trial count are to compute $ISPC_Z$ (Rayleigh's Z, see equation 19.2), to use other phase-based connectivity measures that are less sensitive to the number of data points, or to apply statistical corrections to account

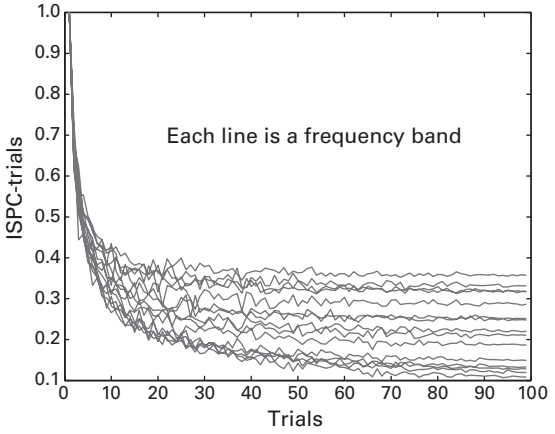

Figure 26.6
Relationship between ISPC-trials (at 300 ms poststimulus) and the number of trials used in the analysis. Results are averaged over 50 iterations of randomly selected trials.

for differences in trial count (Vinck et al. 2010). However, as discussed in chapters 7 and 19, the best strategy is to design the experiment such that you will have a sufficient number of trials in the experiment and a similar number of trials across the conditions you want to compare.

Note that ISPC-time is less sensitive to trial count than is ISPC-trials. This is because the inflation of ISPC-trials is due to having a small number of data points per se, not specifically a small number of trials. Thus, if you compute ISPC-time, you may have several hundred data points per ISPC-time value, even if there are only 25 trials. For this reason, if you are concerned that you have too few trials for ISPC-trials, it is safer to compute and interpret ISPC-time, even at the expense of reduced temporal precision.

26.5 Relation between ISPC and Power

As mentioned several times in this book, power and phase are mostly independent, except that with zero power, phase cannot be estimated. Section 19.6 showed that signal amplitude has minimal effect on ITPC except when power was zero (and this was not a big problem if power was only phasically zero and the data were bandpass filtered). A related example is shown here. Of two electrodes between which ISPC-trials was computed, activity from one electrode was multiplied by 0.0001 to dampen its power. This had no appreciable effect on

the resulting estimate of ISPC, as seen in figure 26.7. Dampening activity from both electrodes or multiplying one or both time series by 10,000 also had no effect. Furthermore, the scatter plots in figure 26.7D show that in this dataset there is no clear relationship between ISPC-trials and power over time. This example suggests that the estimation of phase is not directly linked to power. As with ITPC in section 19.6 and figure 19.10, this illustration does not mean that you should be unconcerned about the possible effects of power on phase-based measures, but it does suggest that within reasonable ranges, and as long as power is not zero, phase-based measures of connectivity can be interpreted independent of power.

26.6 Weighted ISPC-Trials

Weighted ISPC-trials will allow you to test for a statistical relationship between ISPC-trials and trial-varying behavior or experiment variables such as reaction time or stimulus features (Cohen and Cavanagh 2011). The procedures for wISPC-trials are identical to those for wITPC (section 19.7), except that phase angle differences between two electrodes are used instead of the phase angles from one electrode. The wISPC-trials can also help dissociate connectivity from volume conduction: if the wISPC-trials is significant, but there is no analogous correlation between power and that same trial-varying variable, the trial-varying phase angle differences are unlikely to be due to volume-conducted activity.

Correlating ISPC-time with a trial-varying variable does not require the same procedures as wITPC-trials because the connectivity value can already be computed on each trial. Thus, you can simply correlate (preferably using Spearman's correlation, because ISPC values are nonnormally distributed) the ISPC value at each time-frequency point over trials with the trial-varying variable.

26.7 Spectral Coherence (Magnitude-Squared Coherence)

Spectral coherence is similar to ISPC, but the phase values are weighted by power values. In many cases spectral coherence and ISPC will provide similar results. However, because spectral coherence also incorporates power information, results from spectral coherence are likely to be influenced by strong increases or decreases in power. For example, if connectivity increases but power simultaneously decreases, spectral coherence may provide biased results (Lachaux et al. 1999). Typically, you will see the equation for spectral coherence as follows:

$$Coher_{xy} = \left| \frac{S_{xy}}{S_{xx}S_{yy}} \right| \tag{26.2}$$

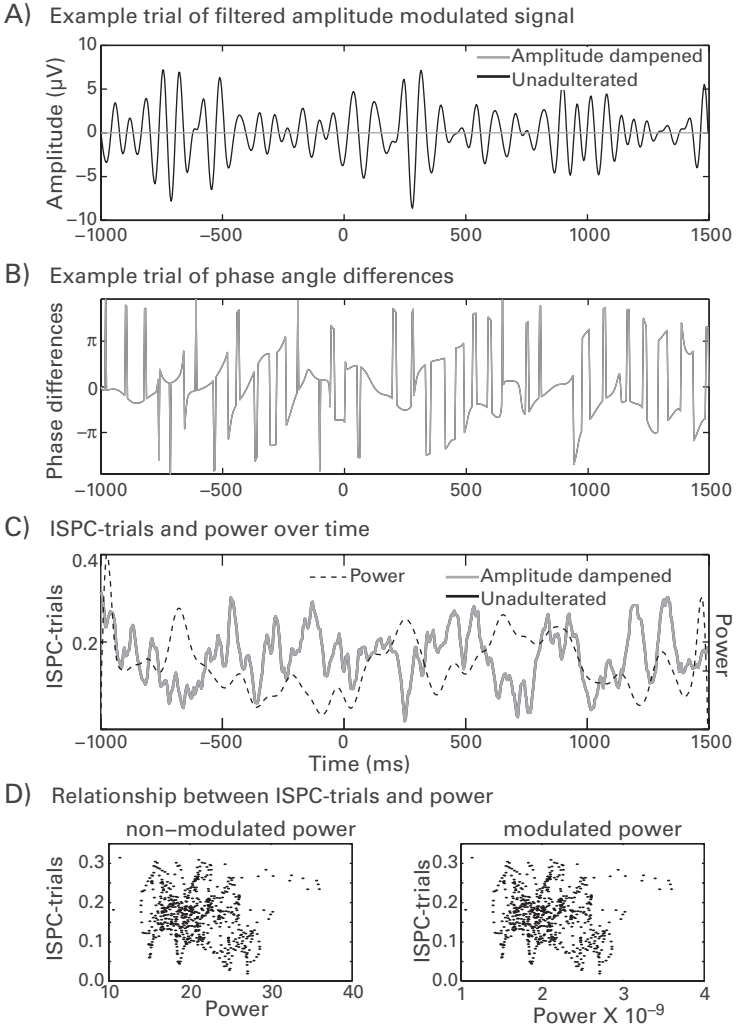

Figure 26.7

ISPC-trials was computed between electrodes Fz and O1, filtered from 10–20 Hz, with and without multiplying one of the time series by 0.0001 to dampen its amplitude (see panel A; the time course is not zero but too small to be visible on this scale). In the example phase angle differences from one trial in panel B and in the ISPC-trials in panel C, the time courses for the modulated and unmodulated data are overlapping and thus not distinguishable in the figure. The power time course from electrode Fz is overlaid, illustrating no relationship between power in one electrode and connectivity between electrodes (this lack of clear relationship holds for O1 power as well). The scatter plots in panel D show ISPC by power of one of the electrodes (each dot is a time point).

S_{xy} is the cross-spectral density between activities at electrodes X and Y, and S_{xx} and S_{yy} are the autospectral densities for electrodes X and Y (this is explained further below). Sometimes, the numerator is squared, which is called magnitude-squared coherence.

For consistency within this book, it is useful to rewrite and explain spectral coherence using a Euler-like format, which is equivalent to equation 26.2 but might be more intuitive in the context of the information you learned in chapters 11–13.

$$C_{xy} = \left| n^{-1} \sum_{t=1}^{n} |m_{tx}||m_{ty}| e^{i\phi_{txy}} \right|^2 \tag{26.3}$$

m_x and m_y are the analytic signals X and Y (the vertical bars surrounding them indicate to take the magnitude of the analytic signals, e.g., via the Matlab function `abs`), φ_{xy} is the phase angle difference between electrodes X and Y, and t refers to trials or time points, depending on whether coherence is computed over time or over trials. This formula may look complicated, but if you break it up into chunks, you will see that it comprises concepts you have already learned. First consider the $e^{i\phi}$ part, which is Euler's formula. These are the phase angle differences between electrodes X and Y. Those phase angle differences are modulated (or weighted) by the signal magnitudes (similar to how wISPC was weighted by a trial-varying variable). The sum and n^{-1} indicates the average over magnitude-modulated phase values. Finally, the length of the complex result (that is, the length of the average vector) is taken. You can see that with the magnitude modulators dropped, equation 26.3 is similar to equation 26.1. You might notice that a problem with equation 26.3 is that the result scales with power, and power changes over frequency, time, task events, and so on. This is why spectral coherence is then normalized by power, by treating equation 26.3 as the numerator and the combined power as the denominator, as in equation 26.4.

$$Coher_{xy} = \frac{C_{xy}}{\left(n^{-1} \sum_{t=1}^{n} |m_{tx}|^2 \right) \left(n^{-1} \sum_{t=1}^{n} |m_{ty}|^2 \right)} \tag{26.4}$$

The numerator is equation 26.3. The denominator is simply the product of the average power values from electrodes X and Y (the averaging is done over trials or time points, depending on whether you are computing coherence over time or over trials). This normalization factor accounts for signal magnitudes and thus puts coherence on a scale from 0 to 1, with 1 being complete coherence and 0 being complete independence. These coherence values are thus interpreted in the same way you would interpret ISPC values. Note that although the equation is normalized by the total power, individual phase angle vectors are still weighted

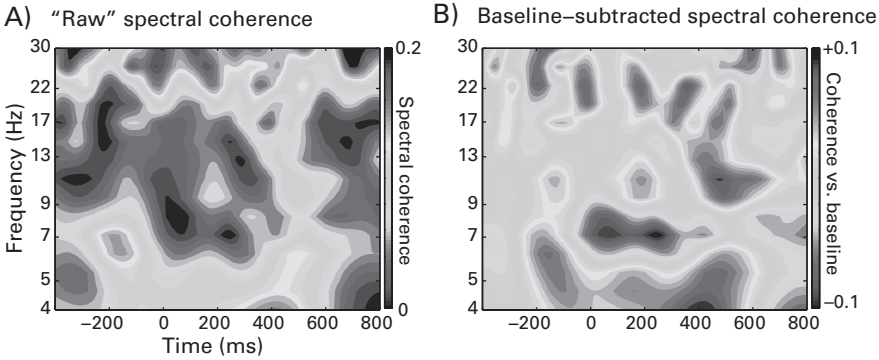

Figure 26.8 (plate 17)
Time-frequency spectral coherence over trials between Fz and O1, shown in "raw" units (no baseline subtraction; panel A) and after linear frequency band-specific baseline subtraction (panel B). These results can be compared with those presented in figure 26.5 (plate 16).

by power, and the mapping between power and phase angle can be nonrandom if particular regions of phase space are associated with relatively increased or decreased power. This may occur, for example, after stimulus onset, when there are both robust changes in power and simultaneous nonuniform phase angle distributions. This is why spectral coherence results can be influenced by power despite being normalized by the total power in the denominator. Figure 26.8 (plate 17) shows an example of spectral coherence computed over trials.

A brief aside on Matlab code: remember from section 13.5 that multiplying an analytic signal by its conjugate is the same as squaring its absolute value but is about twice as fast. Translated into Matlab code, this means that `sig1.*conj(sig1)` is the same thing as `abs(sig1).^2`, where `sig1` is an analytic signal (the result of complex wavelet convolution or filter-Hilbert). To compute the cross-spectral density, you can use `sig1.*conj(sig2)` (this is S_{xy} in equation 26.2). This is equivalent to and faster than, although perhaps less intuitive than, the Euler notation: `abs(sig1).*abs(sig2).*exp(1i*(angle(sig1)-angle(sig2)))`.

Equations 26.4 and 26.2 are for the magnitude of coherence, which quantifies the strength of the coherence. If you do not square and take the absolute value (magnitude) of the numerator, the resulting coherence is a complex number, and you can therefore also extract the phase value, which may be useful for testing the phase lag between the two electrodes or phase lag differences between two conditions. The online Matlab code shows you how to do this. More information on what to do with these phase values is presented in section 26.10.

26.8 Phase Lag-Based Measures

Because effects of volume conduction are instantaneous within measurement capabilities of M/EEG acquisition and within frequencies typically investigated in M/EEG research (Plonsey and Heppner 1967; Stinstra and Peters 1998), spurious connectivity results that are caused by two electrodes measuring activity from the same source will have phase lags of zero or π (π if the electrodes are on opposite sides of the dipole). (For the remainder of this chapter, "zero-phase" will be used instead of "zero- or π-phase" for convenience.) Thus, it seems sensible to avoid spurious connectivity results due to volume conduction by avoiding zero-phase-lag connectivity. There are several phase-based connectivity measures that ignore zero-phase-lag connectivity, including imaginary coherence (Nolte et al. 2004), phase-slope index (Nolte et al. 2008), phase-lag index (Stam, Nolte, and Daffertshofer 2007), and weighted phase-lag index (Vinck et al. 2011). These measures are insensitive to volume conduction, although in some cases they may still be susceptible to mixing sources (Peraza et al. 2012).

Imaginary coherence (Nolte et al. 2004) was developed as a way to apply spectral coherence without concern for spurious connectivity due to volume conduction. Computing imaginary coherence uses almost the same equation as that for spectral coherence, except the imaginary part of the spectral coherence is taken before the magnitude. The online Matlab code will show you how to compute imaginary coherence.

The phase-lag index (Stam, Nolte, and Daffertshofer 2007) measures the extent to which a distribution of phase angle differences is distributed toward positive or negative sides of the imaginary axis on the complex plane (that is, whether the vectors are consistently pointing "up" or "down" in polar space when the imaginary axis corresponds to a vertical line). The idea is that if spurious connectivity is due to volume conduction, the phase angle differences will be distributed around zero radians (as in figure 26.1D). In contrast, non-volume-conducted connectivity will produce a distribution of phase angles that is predominantly on the positive or on the negative side of the imaginary axis. Thus, with the phase-lag index, the vectors are not averaged, but instead, the sign of the imaginary part of the cross-spectral density is averaged. If all phase angle differences are on one side of the imaginary axis, the phase-lag index will be high. In contrast, if half of the phase angle differences are positive and half are negative with respect to the imaginary axis, the phase-lag index (PLI) will be zero.

$$PLI_{xy} = \left| n^{-1} \sum_{t=1}^{n} \operatorname{sgn}\left(imag\left(S_{xyt}\right)\right) \right| \tag{26.6}$$

in which *imag(S)* indicates the imaginary part of the cross-spectral density at time point (or trial) *t* and is extracted from a complex number in Matlab using the function `imag`; *sgn*

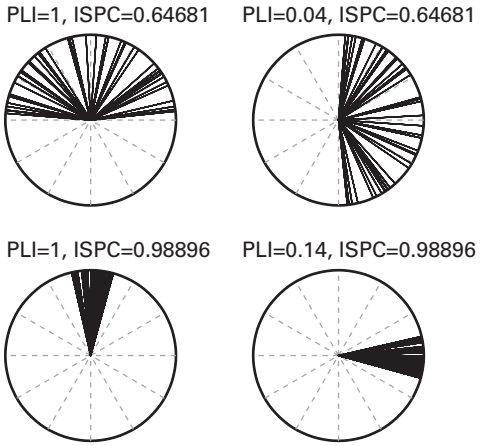

Figure 26.9
Comparison of the phase-lag index (PLI) and ISPC under different phase angle distributions. ISPC depends only on the clustering of the distribution, regardless of its mean phase direction, whereas PLI is sensitive to the phase angle directions rather than clustering per se. Note that the cases on the right side of the figure may be difficult to interpret—zero phase lag can indicate either true zero-phase-lag connectivity or spurious connectivity due to volume conduction.

indicates the sign (–1 for negative values, +1 for positive values, and 0 for zero values). The phase-lag index is less sensitive to outliers, but it is also less sensitive to the amount of clustering in the distribution. That is, if the phase values are spread out in polar space but all on one side of the imaginary axis, phase-lag index will still be high. This is shown in figure 26.9.

The weighted phase-lag index (not to be confused with wISPC; weighted phase-lag index is not computed in relation to any trial-varying variable) is an extension of the phase-lag index in which angle differences are weighted according to their distance from the real axis (Vinck et al. 2011).

$$wPLI_{xy} = \frac{n^{-1}\sum_{t=1}^{n}\left|imag\left(S_{xyt}\right)\right|\operatorname{sgn}\left(imag\left(S_{xyt}\right)\right)}{n^{-1}\sum_{t=1}^{n}\left|imag\left(S_{xyt}\right)\right|} \tag{26.7}$$

You can see equation 26.6 embedded in the numerator of equation 26.7, but equation 26.7 also scales the sign of the angles by the magnitude of the imaginary component (see also figure 1 in Vinck et al. 2011); thus, vectors further away from zero or π radians have a larger influence on the estimate of connectivity. As with spectral coherence, the weighting

term would result in a scaling of wPLI values. Thus, the denominator unscales the final result. The weighted phase-lag index was further developed by introducing a debiasing term to correct for some inflation due to sample size (Vinck et al. 2011). This is shown in the online Matlab code.

The phase-slope index was developed to measure directed phase-based connectivity (Nolte et al. 2008). Consider that if there is a directed functional connection from area A to area B with a 10-ms phase lag, the spectral representation of the 10-ms phase lag (that is, the phase delay) will increase with increasing frequency. This is because a 10-ms delay corresponds to a larger phase lag at 10 Hz compared to 5 Hz (respectively, 0.628 and 0.314 radians, using equation 26.8 below). Thus, the phase-slope index measures whether the slope of the phase lag is consistently positive or negative over several adjacent frequency bins (computed from the Fourier transform). Furthermore, the sign of the slope indicates whether the net connectivity flows from region A to B or the reverse. You can specify which frequency bands to use, thus making the phase-slope index a frequency-band-specific phase-based measure of directed connectivity. The main limitation of the phase-slope index is that it measures only the overall asymmetry of directed connectivity; if there is equally strong bidirectional connectivity the phase-slope index will be close to zero. For this reason, it may be useful to compare results from the phase-slope index to results from ISPC or another non-directional phase-based connectivity measure. Because the phase-slope index is based on FFTs of time series data, it can be computed only over time, not over trials. The phase-slope index is not further discussed here, but the online Matlab code contains a function to compute the phase-slope index (`data2psiX.m`), which is a modified version of the Matlab code that accompanies the Nolte et al. (2008) paper.

Like other measures of connectivity discussed in this chapter, phase-lag-related measures can be computed over trials at each time point or over time at each trial. Figure 26.10 (plate 18) shows a comparison of results from ISPC, phase-lag, and weighted phase-lag indices over trials and over time.

In figure 26.10 (plate 18), you can see that the phase-lag index over trials and weighted phase-lag index over trials (panels A and B) look very similar to each other. The weighted phase-lag index over time (panel D) looks different from the weighted phase-lag index over trials but looks similar to ISPC-time (panel C; this is the same result as shown in figure 26.3A [plate 15] with different color scaling). This figure further highlights that connectivity measures can reveal different dynamics when assessed over time versus over trials.

There are two limitations of phase-lag-based measures. First, if the two electrodes have slightly different frequency concentrations, phase-lag indices can fluctuate rapidly as the

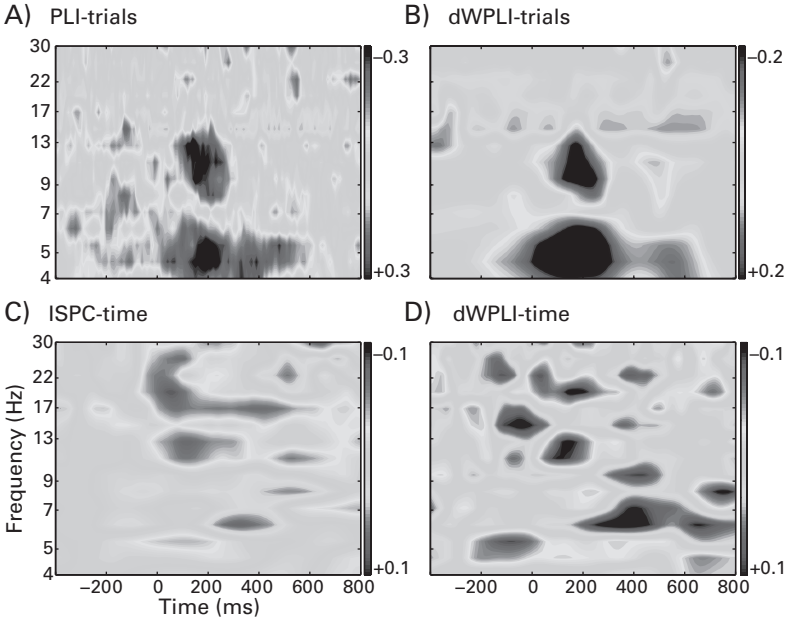

Figure 26.10 (plate 18)
Visual comparison of different measures of baseline-subtracted phase-based connectivity computed between Fz and O1. The color scaling differs across plots to facilitate qualitative comparisons.

phase angle differences spin around polar space. That is, although the analysis parameters are identical, because of some frequency smoothing from wavelet convolution or filtering, one electrode may have more energy at a slightly higher frequency than the other electrode. In this case the preferred phase difference angle will spin around polar space, creating large transient fluctuations in phase-lag indices as the preferred phase difference angle crosses zero or π radians. This can be seen in figure 26.11. Figure 26.11 intentionally shows a somewhat extreme example in which a short time segment length was chosen to illustrate this point. Nonetheless, this phenomenon can happen in task-related analyses if the time segment lengths are short or if there is strong phase clustering combined with rotating phase angle differences. Such an effect can also occur if the frequency of activity at one electrode changes slightly over time, which has been shown to occur in real data (Burgess 2012). This effect can be minimized by increasing the number of time points or trials in the analysis and by selecting time-frequency decomposition methods that increase frequency specificity.

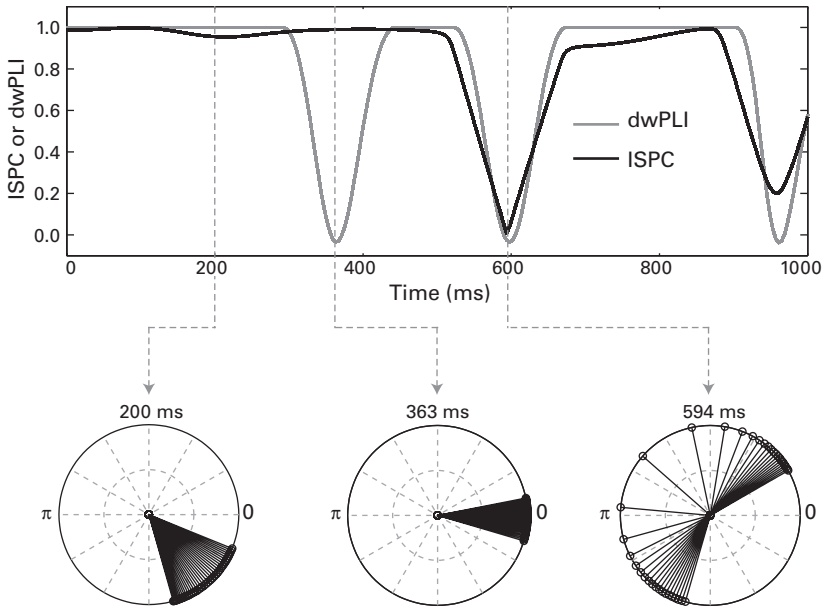

Figure 26.11
Illustration of one potential limitation of phase-lag-based measures. If the preferred phase difference angle is not stationary over time, phase-lag-based connectivity measures will transiently go toward zero as the distribution passes through 0 or π radians (this happens here at 363 ms). Polar plots show the phase angle differences over time corresponding to selected center time points. This figure is a screen shot of a movie, which you can view using the online Matlab code.

A second, related, limitation is that condition differences in phase-lag-based measures can reflect either differences in connectivity or differences in the preferred phase of the connectivity. For example, imagine two conditions with equal phase angle difference clustering, but condition A is associated with a phase angle difference of $\pi/2$ and condition B is associated with a phase angle difference of $\pi/5$. Because of measurement noise, there may be more individual phase angle difference values below zero (negative values on the imaginary plane) in condition B compared to condition A. This will cause phase-lag-based indices to be lower for condition B compared to condition A, although the strength of phase clustering is the same (see also figure 26.9). For this reason, condition differences in phase-lag-based indices should be inspected, for example in combination with the gv-test (explained in section 26.10). If the preferred phase angles of the two conditions are not statistically significantly different, condition differences in phase-lag-based measures can be safely interpreted; if the preferred phase angles are statistically significantly different, the condition

difference in phase-lag-index could be due to differences in the preferred phase difference angles.

In other words, phase-lag-based measures assume that both the phase lag and the frequencies of activities of the two electrodes are stationary for the duration of time used in the analysis. If these two assumptions are not met, then it is possible that the same amount of phase coupling can cause spurious time-varying changes in the apparent strength of phase-lag-based measures.

26.9 Which Measure of Phase Connectivity Should You Use?

The phase-based measure of connectivity that is most appropriate depends in part on how hypothesis-driven versus exploratory and data-driven your analyses are. If you have a priori hypotheses about a small number of specific connectivity patterns, ISPC combined with tests against volume conduction might be a better option because ISPC is maximally sensitive to detecting connectivity, regardless of the phase angle differences.

On the other hand, if you have no or few hypotheses and will thus do a lot of data exploration by testing for connectivity over many electrode pairs, time points, and frequency bands, it might be better to use a connectivity measure that is insensitive to volume conduction, even at the risk of ignoring some potentially true connectivity patterns, because inspecting each of hundreds or thousands of results for possible contamination by volume conduction is impractical. With measures such as phase-lag index or imaginary coherence, you can be more confident about minimizing volume conduction contamination. On the other hand, because they are also sensitive to the mean phase angle of the phase angle distribution, phase-lag-based measures may be best suited for resting state or tasks in which connectivity strength is not compared across conditions.

Another consideration is whether to compute connectivity over trials or over time. This decision is less influenced by your orientation toward hypothesis testing versus data exploration but rather depends on your task design and expectations about the results. Differences between connectivity over trials versus over time are discussed in section 26.3, but briefly: connectivity over time is more sensitive to detecting high-frequency connectivity and for resting-state data or tasks that have long event durations (i.e., at least several hundred milliseconds) due to the poor temporal precision of connectivity over time; connectivity over trials has higher temporal precision and is therefore better able to identify the time course of changes in connectivity and is also better able to identify transient changes in connectivity.

26.10 Testing the Mean Phase Angle

Typically when ISPC is computed, the angle of the mean vector is ignored (as it is with ITPC from chapter 19). The angle of the mean vector (also called the "preferred angle") can be extracted in Matlab by using `angle(mean(. . .))` instead of `abs(mean(. . .))`. The angle of connectivity may be useful for two reasons. First, it provides the phase lag or the time lag between the two electrodes. The phase lag has units of radians but can be converted to time in milliseconds using equation 26.8.

$$\text{lag(ms)} = \frac{1000\phi_d}{2\pi f} \tag{26.8}$$

in which ϕ_d is the phase angle difference between the two electrodes in radians (change the 1000 to 1 to convert to seconds). Time lag can be used as supportive evidence for directionality of connectivity but does not provide unambiguous evidence for directionality, nor does it imply causality (see section 25.2).

The second reason the phase angles from ISPC might be useful is to test whether a phase angle or a phase angle difference is statistically significantly different from some specified phase angle. This can be used either to test for possible contamination by volume conduction or to test for condition differences in the preferred phase angle of connectivity (e.g., if the phase lag is longer in one condition compared to another condition). This can be done with the "*v*-test" (Durand and Greenwood 1958; Zar 1999).

$$u = nISPC\cos(\phi - \Phi)\sqrt{\frac{2}{n}} \tag{26.9}$$

in which n is the number of trials or time points, ϕ is the observed phase angle difference, and Φ is the hypothesized phase angle against which to test (an intermediary step folded into equation 26.9 is called v, hence the name). Remarkably, the quantity u is normally distributed under the null hypothesis (figure 26.12). This means that you can directly convert the u from equation 26.9 into a p-value, as you would for a normal Z-score. This can be done with the Matlab function `normcdf` (which is in the Matlab statistics toolbox), or you can look up key values in a p-value table (e.g., $u = 1.96$ corresponds to $p < 0.05$ two-tailed).

To run this test, specify Φ as a phase value against which to test (make sure Φ is radians and not milliseconds or angles in degrees). In the test for volume conduction you would hope to obtain a nonsignificant p-value (that is, the null hypothesis that the phase angle difference is not zero cannot be rejected).

Unfortunately, this test has four disadvantages when it comes to testing phase angles with typical EEG datasets. First, the *v*-test works well for a small number of data points but is less

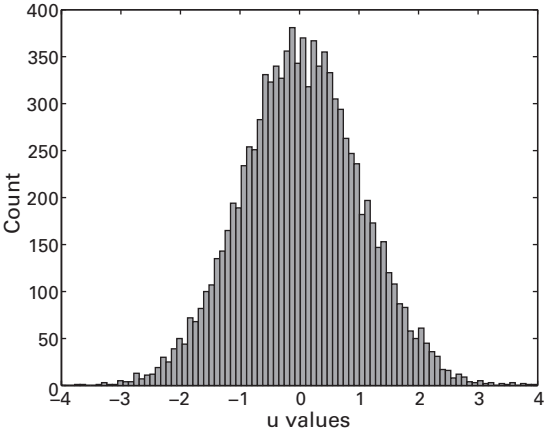

Figure 26.12
A distribution of *u* values from equation 26.9 under the null hypothesis, generated by computing ISPC between 640 pairs of random vectors (10,000 tests were done to create this distribution). In this particular simulation, the mean was –0.015, and the variance was 1.0003.

useful when there are more than 20 data points. This can be seen in figure 26.13B, where $N = 300$ (a good number of trials or time points for an EEG study with a very high signal-to-noise ratio) leads to half of polar space being "significantly" equal to zero radians. Second, the *v*-test is symmetric, meaning that the *p*-value for phase angle φ is exactly 1 minus the *p*-value for $2\pi - \varphi$ (this results from the cosine component in equation 26.9, which gives a cyclic response). Third, it has narrow slopes, that is, *p*-values tend to be very close to 0 or very close to 1, particularly for large *N*. Fourth, it produces many false positives, particularly with many data points. In 10,000 simulations of 640 pairs of random phase angles, around 5–6% of phase couplings are statistically significant using an α of $p < 0.05$. Although it may seem unsurprising that random data are considered significant 5% of the time when using an α of 5%, a modification of the *v*-test presented below is more sensitive and therefore results in fewer Type-I errors (false positives).

Therefore, I propose a modification of the *v*-test, called the Gaussian *v*-test, or *gv*-test. The *gv*-test simply replaces the cosine component with a Gaussian component, as seen in equation 26.10. This minor change addresses several disadvantages of the *v*-test (see figure 26.13 for comparisons between the *v*-test and the *gv*-test): it is robust to a large number of data points because the width of the Gaussian scales with the number of points; it is asymmetric and thus tests precisely for one region of phase space rather than testing polar opposites of phase space; and it produces fewer false positives. In the same 10,000 simulations described above, there were around 0.3–0.5% false positives. Therefore, the *gv*-test is a

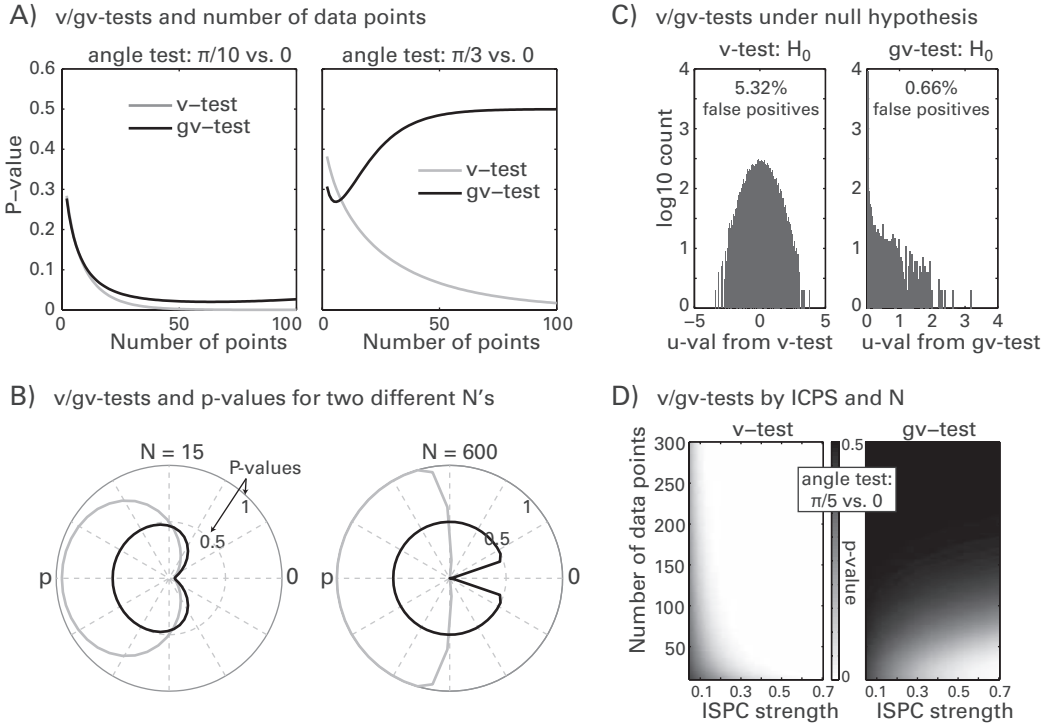

Figure 26.13

Comparison of *v*-test and *gv*-test. Panel A shows *p*-values (*y*-axis) as a function of the number of points (time points or trials) in the analysis (*x*-axis), for $\pi/10$ versus 0 radians (left) and $\pi/3$ versus 0 radians (right). Panel B shows *p*-values (radial axis) as a function of phase angles in radians in a test against zero radians. Panel C shows *u*-values from *v*-test and *gv*-test for random data. Panel D shows *p*-values (gray-scale) as a function of the number of data points (*y*-axis) and ISPC strength (*x*-axis).

suitable alternative to the *v*-test in situations of large number of data points, as is typically utilized in EEG studies. For a small number of data points, the *gv*-test and the *v*-test will provide similar results.

$$u = nISPCe^{-\phi_d^2/4\pi/n}\sqrt{\frac{2}{n}} \tag{26.10}$$

in which ϕ_d is the difference between the observed phase value and the phase value against which to test [the same as $\varphi - \Phi$ in equation 26.9], and $4\pi/n$ is $2s^2$, where s is the square root of $2\pi/n$.

Thus, after ISPC has been computed, the *gv*-test is one possible analysis to test whether the result may have been due to volume conduction. If the *gv*-test provides evidence that the phase angle differences are 0 or π, the result could indicate zero-phase-lag connectivity or volume conduction (in this case, further inspection would be required, as described in section 25.10); if the *gv*-test provides evidence that the phase angle differences are not 0 or π, the result cannot be due to volume conduction.

One point to keep in mind about testing phase angles across subjects is that the precise phase value will depend on a variety of factors including cortical folding and dipole orientation, which may differ across subjects even if the neurocognitive process is the same.

26.11 Describing These Analyses in Your Methods Section

There are many possibilities for phase-based measures of connectivity. Thus, the two most important things to include in the Methods section are a justification for why you chose one measure over the other measures and a clear description of the analysis performed. The clear description is important because of inconsistent terminology in the literature and because of the number of possibilities for computing phase-based connectivity, including whether connectivity was computed over time or trials (this is often unclear in publications). Include the formulas, statistical procedures or transformations (e.g., *z*-normalization for wISPCz), the number of trials per condition, and, if relevant, the number of cycles from the wavelet convolution or the length of time if the phase-based connectivity measure was computed over time; these details provide information regarding the temporal precision of the results. If you make modifications to standard analysis approaches, you can also consider including your Matlab code as an appendix or in an online supplemental materials section.

26.12 Exercises

1. Select one seed electrode and one frequency band and compute phase-based connectivity between that seed electrode and every other electrode. Use two methods for phase-based connectivity that were presented in this chapter, one that is volume conduction independent (e.g., PLI) and one that could produce spurious connectivity due to volume conduction (e.g., ISPC). Do not apply a baseline subtraction. Make topographical plots of seeded connectivity in a time window of your choice (e.g., 300–350 ms). What are the similarities and differences between results from the two methods, and what might be the reasons for the similarities and differences?

2. Now apply a baseline subtraction to the results (you can choose the baseline time period). Are there any changes in the plots after baseline subtraction (note that the color scaling will be different after baseline subtraction), and how do results from the two analyses compare with each other after baseline subtraction?

3. From the results in exercise 1 above, pick one "target" electrode (any electrode other than the seed) and provide evidence, using additional data analyses if necessary, for or against that measure of phase-based connectivity between that electrode and the seed being driven by volume conduction.