separate figures for each frequency. The
me within a frequency, but the color scal-
seline normalization and should also be

phical distributions of power with com-
differences more prominent in some fre-
be causing these differences?

19 Intertrial Phase Clustering

In chapter 13 you learned that the analytic signal resulting from convolution between an EEG time series and a complex wavelet, or resulting from the filter-Hilbert method, can be conceptualized as a vector in a complex polar plane with a magnitude (length of the vector) and a phase angle (the angle in radians relative to the positive real axis). These phase angles provide information about the timing of frequency-band-specific activity and are the main focus of this chapter.

If you want to compute the consistency of time-domain EEG traces over trials, you simply average the activity at each time point across trials to form an event-related potential (ERP). Computing the consistency of time-frequency power over trials is the same as for ERPs: average the frequency-band-specific power at each time point across trials. However, computing the consistency of time-frequency phase values over trials is not so simple because phase values cannot be averaged together in the same way that voltage values or power values can be averaged. Learning how to average phase values and compute the consistency of phase values over trials is useful not only for examining the timing of frequency-band-specific activity but also forms the basis of several phase-based connectivity methods.

19.1 Why Phase Values Cannot Be Averaged

It would seem easy simply to average together the phase angle time series across trials, the way you would average single-trial EEG traces to form an ERP. Unfortunately, however, this is inappropriate. Phase angles are circular, which means that, for example, 0.05 and 6.2332 might seem far apart from each other when considered as numbers on a number line, but are actually very close to each other when considered as radians. Indeed, averaging these two numbers together gives a result close to π, which is in the opposite side of polar space from angles 0.05 and 6.2332 radians. Thus, it is not appropriate to average radian values together as if they were normal numbers (figure 19.1).

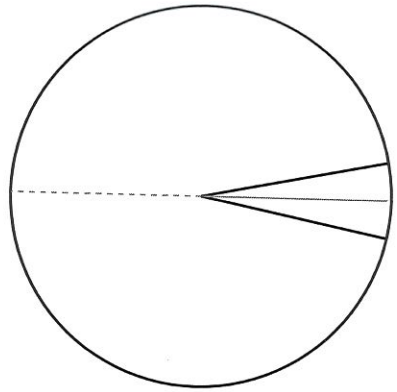

Figure 19.1
Two vectors (black lines) have similar angles; the average vector (solid gray line) reflects their proximity, but a vector formed by averaging their angles in radians (dashed gray line) does not reflect the two vectors. This is an illustration of why averaging phase angle radian values is inappropriate.

Remember that phase angles can be represented as vectors with unit length on a circle, and that Euler's formula (e^{ik}) provides a convenient mathematical description of those vectors in a complex plane. Thus, a population of phase angles can be represented as a population of vectors on a circle. For example, figure 19.2A,C shows phase angles from all trials at 200 ms and at 800 ms post–stimulus onset. Each vector (shown as a gray line) was formed by taking the phase angle from one trial and setting the vector length to 1 (rather than the vector length reflecting the similarity between the time series and the wavelet). The histograms on the bottom row show these same phase angle distributions in a different and perhaps more familiar way. You can see that at 200 ms, the vectors are clustered around one region of the circle, whereas at 800 ms they are more scattered. Another way to describe these distributions of phase angles is that the distribution of phase angles over trials is less uniform at 200 ms compared to the distribution at 800 ms. The extent to which these vectors are clustered (or nonuniformly distributed) is the measure of phase clustering across trials. This should make sense conceptually: if the timing of an oscillatory process is the same or similar at each repetition of a stimulus or other experiment event, their phase angles should take on similar values across trials.

19.2 Intertrial Phase Clustering

The measure of phase angle clustering illustrated conceptually in figure 19.2 is called intertrial phase clustering (ITPC). ITPC measures the extent to which a distribution of phase

ector (solid gray line) reflects their proxim-
(dashed gray line) does not reflect the two
radian values is inappropriate.

d as vectors with unit length on a
ent mathematical description of those
phase angles can be represented as a
19.2A,C shows phase angles from all
ach vector (shown as a gray line) was
setting the vector length to 1 (rather
the time series and the wavelet). The
angle distributions in a different and
the vectors are clustered around one
ttered. Another way to describe these
phase angles over trials is less uniform
he extent to which these vectors are
of phase clustering across trials. This
illatory process is the same or similar
nt, their phase angles should take on

ptually in figure 19.2 is called inter-
t to which a distribution of phase

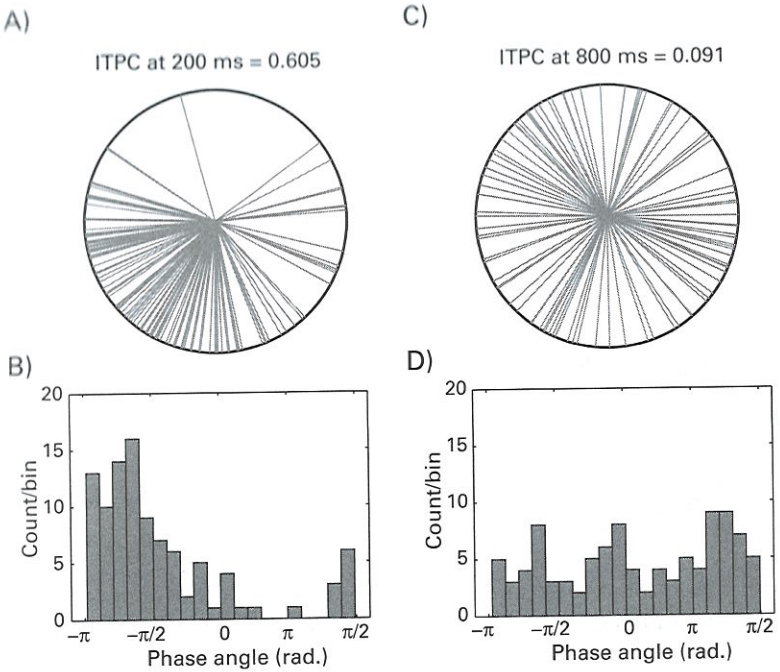

Figure 19.2
Phase-angle distributions at two time-frequency points. Each line in panels A and C corresponds to one trial, and the histograms in panels B and D show counts of trials per phase bin. It is clear that the phase angles are more clustered at 200 ms (panels A and B) compared to 800 ms (panels C and D).

angles at each time-frequency-electrode point across trials is nonuniformly distributed in polar space. In the literature you may see this measure referred to as "phase-locking value," "phase-locking factor," "phase resetting," "phase coherence," "intertrial phase coherence," or "cross-trial" instead of "intertrial." The term intertrial phase clustering is preferred here because it is a description of the analysis rather than an interpretation of the result (analysis terminology is further discussed in section 21.1).

How is the uniformity of the distribution of phase angles measured? This is done by computing the average vector (remember that the vectors are averaged, not their phase angles in radians) and then taking the length of that average vector. Figure 19.3 shows four examples of pairs of vectors, their average vector, and the length of the average vector. Note that the individual vectors always have unit length (length of 1.0), but the average vector has a length less than 1. The further apart the two vectors are from each other, the smaller the length of the average vector. Thus, the length of the average vector reflects the closeness of the two unit-length vectors. ITPC is this length.

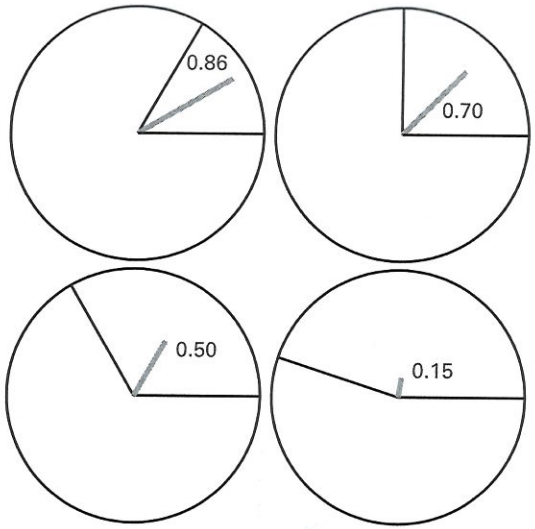

Figure 19.3
Example pairs of unit-length vectors (black lines) and their averages (gray lines). The numbers inside each circle indicate the length of the average vector. This number is the ITPC for those two vectors.

ITPC is bound between zero and one, with zero indicating completely uniformly distributed phase angles and one indicating completely identical phase angles. This should make sense from figure 19.3. Consider two extreme situations: at one extreme, all phase angles are perfectly uniformly distributed, and the average vector has a length of 0.0; at the other extreme, all phase angles are identical, and the average vector has a length of 1.0. It is not possible for ITPC to be negative or greater than one.

At this point, you should have an intuitive understanding of ITPC. The next step is to quantify ITPC mathematically.

$$\mathrm{ITPC}_{tf} = \left| n^{-1} \sum_{r=1}^{n} e^{ik_{tfr}} \right| \tag{19.1}$$

The double vertical bars indicate the absolute value or, in this case, the length of the average vector. This is necessary because the result of the averaging is a complex number (because the vectors are described by complex numbers) that contains both the length and the angle of the average vector. n is the number of trials; n^{-1} is a convenient shorthand for $1/n$, and, combined with the summation operator, indicates an average. e^{ik} is from Euler's formula and provides the complex polar representation of a phase angle k on trial r at time-frequency

verages (gray lines). The numbers inside
per is the ITPC for those two vectors.

cating completely uniformly distrib-
ical phase angles. This should make
is: at one extreme, all phase angles
tor has a length of 0.0; at the other
vector has a length of 1.0. It is not

nding of ITPC. The next step is to

(19.1)

this case, the length of the average
ging is a complex number (because
ains both the length and the angle
onvenient shorthand for *1/n*, and,
rage. e^{ik} is from Euler's formula and
ngle *k* on trial *r* at time-frequency

point *tf*. You can see that the *M* from equation 13.8 is not present in equation 19.1. This means that magnitude information is not taken into account when computing ITPC, and thus, the lengths of all vectors are implicitly set to 1.0. This formula can be expressed in Matlab code as:

```
abs(mean(exp(1i*k)))
```

where `k` is a vector of phase angles at one time-frequency point over trials. Notice that this equation does not average phase angles in radians; it averages complex vectors whose angles are defined by the phase angles in radians. Make sure you write `abs(mean(. . .))` and not `mean(abs(. . .))`. The former is the length of the mean vector (this is what you want), whereas the latter is the mean of the individual vector lengths. Because each vector has a length of 1, the mean of all vector lengths is exactly 1.

The angle of the average vector is the average phase, or the "preferred" position of the angles at that time-frequency point over trials (in figure 19.2A, the preferred phase angle would point down and to the left). In practice, the average phase angle is not often used. However, it would be possible for two experimental conditions to elicit comparable ITPC strength but significant differences in the preferred phase angle. Statistical methods for testing differences in preferred phase angle over conditions or time points are discussed in chapter 34.

Equation 19.1 will generate an ITPC value for a single time-frequency-electrode point. This equation is then repeated over many time points and many frequencies to generate a time-frequency map of ITPC at each electrode. An example ITPC time course and a time-frequency plot are shown in figure 19.4.

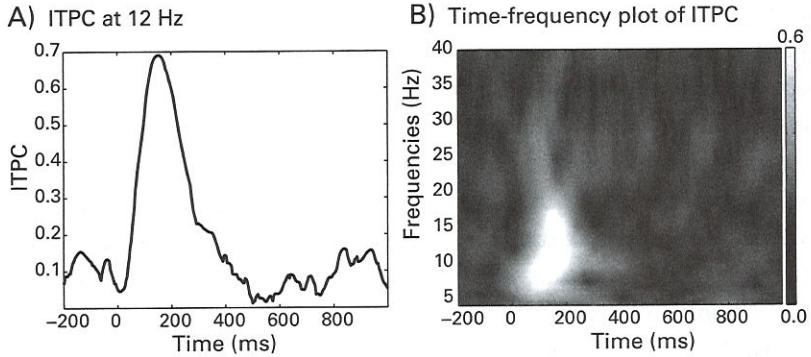

Figure 19.4
ITPC from electrode Pz at 12 Hz (panel A) and over time-frequency space (panel B).

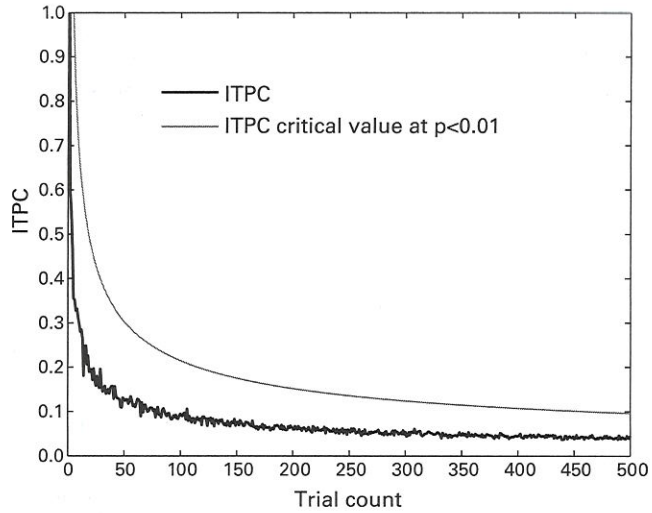

Figure 19.5
ITPC as a function of trial count for simulated random phase angles averaged over 50 simulations. Also shown is the critical ITPC value for each trial count, corresponding to a *p*-value threshold of 0.01.

19.3 Strength in Numbers

Trial count influences ITPC. Because ITPC cannot be below zero, noise and sampling errors are more likely to increase rather than decrease ITPC, particularly when there are few trials. This is illustrated in figure 19.5, which shows ITPC as a function of the number of data points (trials) using randomly generated phase angles. You can see that even with random numbers, ITPC values are fairly high with small sample sizes. The gray line shows the critical ITPC value corresponding to $p < 0.01$; any ITPC value above this line would be considered statistically significant (section 34.5 discusses how to compute and interpret this critical value). You can also see that even with randomly generated phase angles, ITPC values do not reach 0 after 500 trials.

Figure 19.6 shows ITPC as a function of trial count using real data (electrode FCz at 6 Hz). From a total of 99 trials, random subsets of trials were selected, ITPC was computed, and then the average ITPC between 100 ms and 450 ms was taken. Figure 19.6A shows ITPC as a function of the number of randomly selected trials, similar to the plot in figure 19.5. The value of ITPC stabilizes with around 20 trials, although it becomes statistically significant only after around 45 trials (using a *p*-value threshold of 0.01). Figure 19.6B shows the time course of ITPC at electrode FCz for all 99 trials compared to 20 randomly selected trials.

les averaged over 50 simulations. Also
g to a *p*-value threshold of 0.01.

w zero, noise and sampling errors
icularly when there are few trials.
ction of the number of data points
that even with random numbers,
gray line shows the critical ITPC
s line would be considered statis-
and interpret this critical value).
angles, ITPC values do not reach

real data (electrode FCz at 6 Hz).
ed, ITPC was computed, and then
gure 19.6A shows ITPC as a func-
e plot in figure 19.5. The value of
tatistically significant only after
19.6B shows the time course of
mly selected trials.

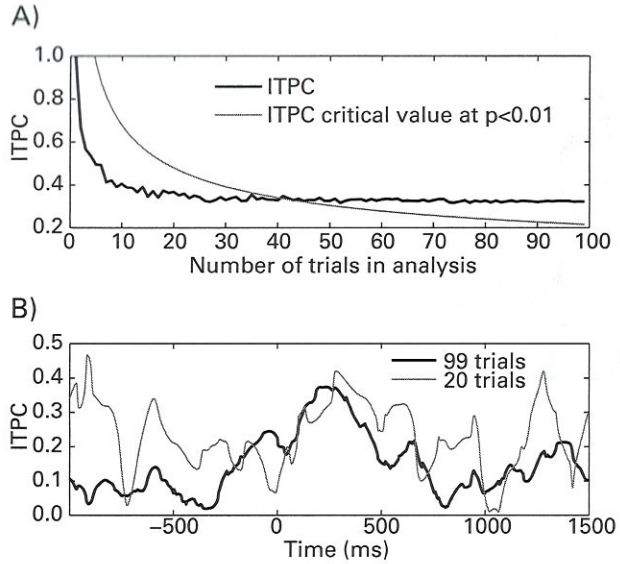

Figure 19.6
Relationship between ITPC and trial count in real data. Panel A shows the same analysis as was shown in figure 19.5 (ITPC as a function of trial count) but for real data instead of randomly generated phase angles. Each point is the average of 50 iterations of random trial selection. Panel B shows the time courses of ITPC for 20 randomly selected trials and all 99 trials. Phase angles were extracted via complex wavelet convolution with a 6-Hz wavelet.

Figure 19.6A shows the relationship between ITPC and trial count for one frequency band; figure 19.7 shows results of this analysis computed over a wider range of frequencies. You can see that the effect of trial count on ITPC depends on the frequency. For example, ITPC values for frequencies below 15 Hz seem to stabilize after around 20 trials, whereas ITPC values from higher frequencies seem to require more trials to stabilize. Keep in mind that figures 19.6 and 19.7 are based on data from one electrode, one subject, and one time window; the results of this test are likely to vary depending on the characteristics of your data. As with figure 18.13, you should interpret this figure not as an absolute reference for how many trials you need but rather as a way to examine in your own data whether you likely have enough trials for a stable estimate of ITPC within each condition. You can perform this analysis on your data by pooling all trials from all conditions together and examining whether the ITPC-by-trial-count function stabilizes with the number of trials corresponding to the number of trials in each condition.

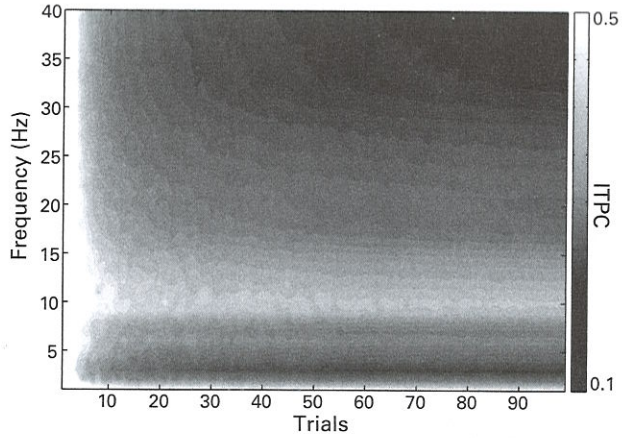

Figure 19.7
ITPC as a function of trial count and frequency. Note that where there is strong ITPC (here, around 10 Hz), trial count seems to have less influence on the strength of ITPC. Results are averaged over 50 iterations of randomly selected trials.

Low trial count can be more deleterious when ITPC is compared between conditions that differ in trial count. Consider figure 19.6A, and imagine comparing ITPC between a condition with 15 trials and a condition with 60 trials. The condition with 15 trials is likely to have larger ITPC simply because of lower trial count, regardless of what neurocognitive processes may have occurred during those two conditions. On the other hand, if your two conditions have trial counts of 60 and 70, there is less cause for concern about possible spurious results due to trial count differences.

19.4 Using ITPC When There Are Few Trials or Condition Differences in Trial Count

If possible, try to avoid this issue all together by designing the experiment such that each condition has at least 30 trials and all conditions have roughly the same number of trials. However, this is not always possible. If you have differences in trial count across conditions and would like to compare condition differences in ITPC, there are three strategies you can apply.

The first strategy is to match conditions for trial count within subject. The procedures, advantages, and limitations of selecting trials across conditions are discussed in section 7.4.

The second strategy is to apply a condition-specific baseline subtraction or percentage change transformation. The positive bias in ITPC due to trial count affects all time points

ere there is strong ITPC (here, around 10
f ITPC. Results are averaged over 50 itera-

s compared between conditions that
e comparing ITPC between a condi-
ndition with 15 trials is likely to have
ess of what neurocognitive processes
e other hand, if your two conditions
ncern about possible spurious results

ion Differences in Trial Count

ning the experiment such that each
roughly the same number of trials.
nces in trial count across conditions
C, there are three strategies you can

unt within subject. The procedures,
ditions are discussed in section 7.4.
baseline subtraction or percentage
to trial count affects all time points

(indeed, this bias is even present with random data, as shown in figure 19.5), and thus, subtracting a baseline ITPC value will help minimize the bias. Because ITPC is not affected by $1/f$ power-law scaling, linear baseline subtraction is appropriate. Decibels, percentage change, and Z-transform are also acceptable baseline normalizations. However, this is not necessarily an ideal strategy. Low trial count not only increases ITPC, it also increases the susceptibility of ITPC to noise or nonrepresentative data. Thus, the baseline ITPC may be noisy, as can be seen in figure 19.6B.

The third strategy is to transform the ITPC to $ITPC_Z$, also known as Rayleigh's Z. The formula for $ITPC_Z$ follows:

$$ITPC_Z = n*ITPC^2 \tag{19.2}$$

in which n is the number of trials (see figure 19.8). Note that $ITPC_Z$ cannot be interpreted as a standard statistical Z value (that is, a value drawn from a distribution with mean of zero and variance of one) because it cannot be below zero, unless you somehow have a negative number of trials. However, a p-value can be computed from $ITPC_Z$ to determine statistical significance. This is discussed in section 34.5.

19.5 Effects of Temporal Jitter on ITPC and Power

Because ITPC reflects the precise timing of band-specific activity, it is sensitive to temporal jitter or temporal uncertainty of experiment events. This temporal jitter can arise, for

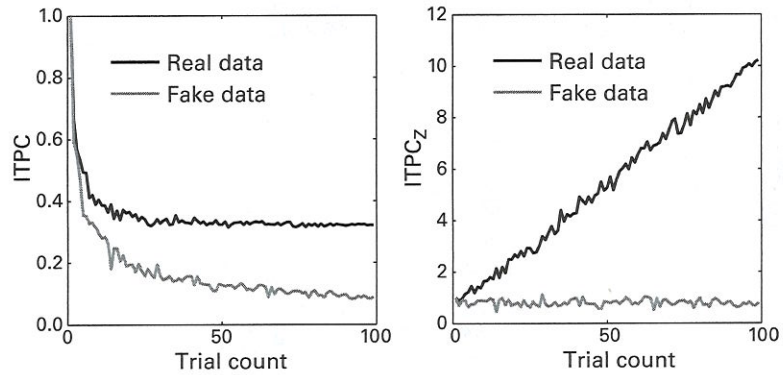

Figure 19.8
Comparison of ITPC and $ITPC_Z$ for randomly generated phase angles (gray lines) and real data (black lines). You can see that $ITPC_Z$ remains flat with increasing trial count for random data, whereas $ITPC_Z$ increases with trial count for real data, even though the "raw" ITPC values remain stable. This reflects increased reliability of $ITPC_Z$ with additional data.

example, from uncertainties in monitor display rates or time locking to events with uncertain onset times, such as EMG activations or pupil responses.

The negative impact of temporal jitter becomes worse at higher frequencies because the cycles become shorter. Imagine a temporal jitter of 10 ms; this jitter corresponds to only 2% of a cycle at 2 Hz but 40% of a cycle at 40 Hz. The negative effects of temporal jitter on ITPC can be demonstrated by introducing temporal jitter into real data. In the plots in figure 19.9 (plate 10), a random time lag between 4 and 40 ms was added to each trial, and then ITPC and power were computed. Temporal jitters had a strong negative impact on ITPC above 6 Hz and obliterated the increase in ITPC at 100 ms from 22 to 35 Hz. In contrast, power was largely unaffected by the temporal jitters. This result shows that phase values are more temporally precise measurements of frequency-band-specific activity compared to power values.

If you use a 60-Hz monitor and there is uncertainty as to when the stimuli are drawn on the monitor with respect to when the experiment marker is recorded in the EEG acquisition

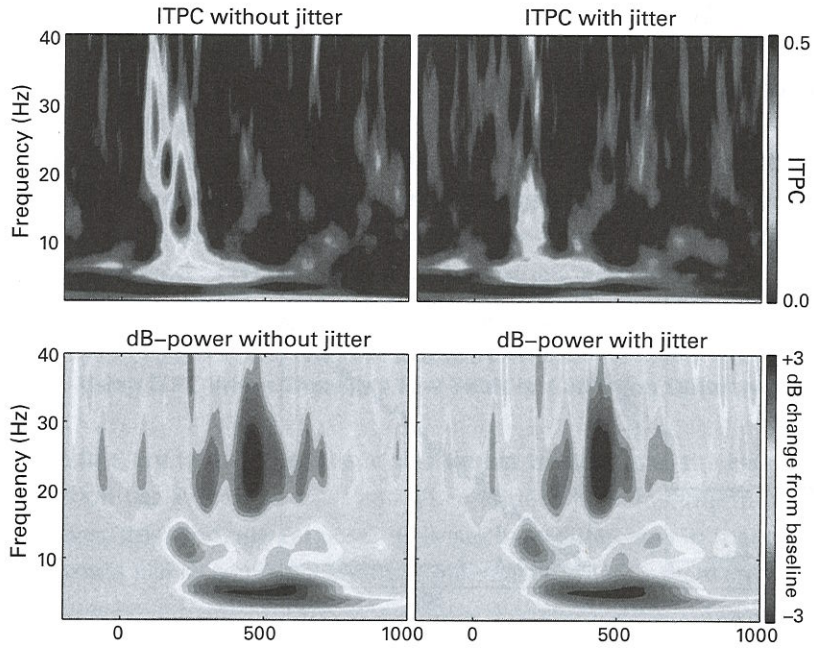

Figure 19.9 (plate 10)
Temporal jitters of less than 40 ms can have deleterious effects on ITPC (top row), particularly at frequencies above 10 Hz. In contrast, temporal jitters have little noticeable effects on power.

e locking to events with uncer-

higher frequencies because the
this jitter corresponds to only
ive effects of temporal jitter on
real data. In the plots in figure
s added to each trial, and then
rong negative impact on ITPC
from 22 to 35 Hz. In contrast,
esult shows that phase values
d-specific activity compared to

when the stimuli are drawn on
ecorded in the EEG acquisition

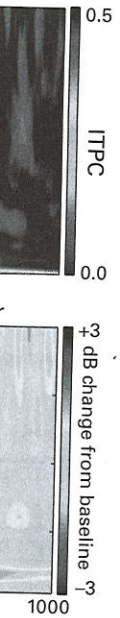

top row), particularly at frequen-
cts on power.

system, this could produce jitters up to 34 ms, which is the same range used in figure 19.9 (plate 10). Therefore, if you want to examine ITPC in the alpha band or higher, it would be a good idea to compare the timing of the experiment markers with the timing of the actual stimulus onset, with an oscilloscope if possible. If ITPC in relatively high frequencies is important for your analyses, you might also consider using a monitor with at least a 120-Hz refresh rate or using LEDs or auditory tones that are more temporally precise. The timing uncertainty related to monitor refresh rates can be further compounded if you have stimuli appear in different locations on the monitor. Monitor pixels are drawn in a loop that lasts the duration of the refresh rate; if stimulus locations differ randomly across trials and this is not taken into consideration in the analyses, temporal jitters of up to one refresh (e.g., 17 ms for a 60-Hz monitor) could be introduced into the data. If there are unavoidable and large temporal jitters in the experiment events, it might be best not to analyze ITPC except at low frequencies.

19.6 ITPC and Power

Several times in this book it is noted that phase angles are independent of power except in situations of very low power. The issue is that with decreasing power, phase becomes more difficult to estimate. Consider the extreme case, in which there is zero power at a particular frequency band: the band-specific activity is zero, there are no oscillations, and therefore phase is undefined. This extreme situation of zero power is unlikely to occur in real EEG data, partly because the brain generates broadband activity in addition to frequency-band-specific activity and partly because noise often increases broadband power. However, decreased power may nonetheless affect the signal-to-noise ratio of phase estimates, and so ITPC may in fact be affected by power. This point was illustrated using simulated data (Muthukumaraswamy and Singh 2011), suggesting that phase-based measures can be influenced by power when the signal-to-noise ratio is very low (less than –10 dB; in that simulation, power had little effect on phase-based measures for simulated data with higher signal-to-noise ratios).

Figure 19.10 shows the relationship between ITPC and power using real data. ITPC was computed over time for the original data and the data tapered with a 1-Hz sine wave with an amplitude fluctuation between 0 and 2. This caused the amplitude of the data to fluctuate between zero and double its original amplitude. This was done on the broadband signal and on the signal after bandpass filtering from 10 to 20 Hz. The broadband ITPC shows some differences between the unmodulated and modulated signals around times when the amplitude modulation was close to zero but seems intact otherwise, including when the amplitude was

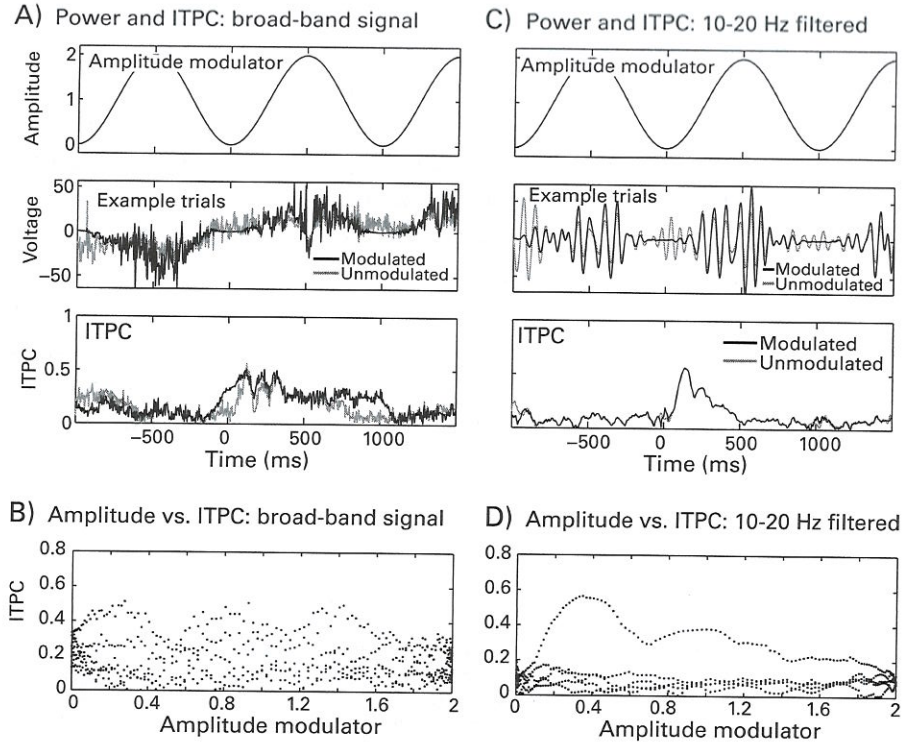

Figure 19.10

The relationship between power and ITPC. Real data were modulated by a 1-Hz sine wave (panels A and C, top row), resulting in significant power decreases and increases over time (panels A and C, middle row). ITPC based on broadband activity was somewhat affected by power being attenuated to zero, whereas ITPC based on band-limited activity was less affected by power (third row; the black and gray lines in panel C are mostly overlapping). Panels B and D show the relationship between ITPC and the modulating power signal.

doubled. Thus, extreme signal attenuation can have an adverse effect on ITPC, but signal enhancement had little effect in this example. Figure 19.10B shows that there was no clear relationship between the modulating signal amplitude and ITPC over time.

The bandpass-filtered data showed very similar ITPC time courses for the unmodulated and modulated activity, consistent with the idea that the phase information remains intact when the amplitude information is strongly attenuated. Although it may appear from figure 19.10D that there is a relationship between the modulating signal and ITPC (e.g., the increase in ITPC when the amplitude modulator is around 0.4), this is because a trough of the

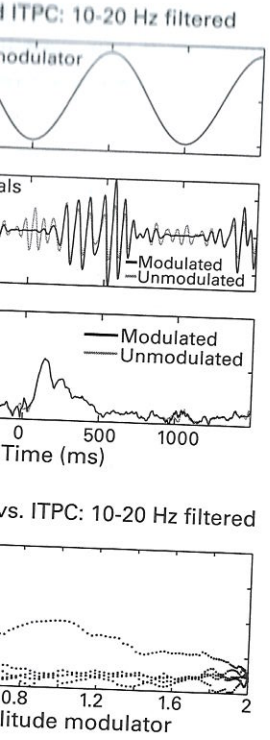

ated by a 1-Hz sine wave (panels A and
es over time (panels A and C, middle
l by power being attenuated to zero,
power (third row; the black and gray
ne relationship between ITPC and the

dverse effect on ITPC, but signal
0B shows that there was no clear
ITPC over time.
me courses for the unmodulated
hase information remains intact
Although it may appear from
ulating signal and ITPC (e.g., the
4), this is because a trough of the

1-Hz amplitude modulator coincided with a stimulus-related increase in ITPC. For example, in the online Matlab code, if you change the amplitude-modulating signal from a sine to a cosine, the relationship between modulating amplitude and ITPC will reverse because the phase of the modulating amplitude will shift.

The negative effect of very low power on ITPC was observed only in the broadband signal; the amplitude of the data had relatively little impact on the ITPC of bandpass-filtered data. This is because the time periods of zero power were relatively brief and because the phase values at each time point of a wavelet convolution could still be estimated from surrounding points (because stationarity is assumed during the wavelet period). Thus, longer periods of sustained zero (or near-zero) power might hinder estimation of frequency-band-specific phase values. The more important point of figure 19.10, however, is that for all nonzero points of the amplitude modulator period, power modulation had no effect on ITPC. You can try this yourself in the online Matlab code by increasing the magnitude of the power modulator envelope. It will have no effect except at power values of zero. Thus, changes in power within realistic ranges of EEG data do not necessarily lead to spurious inflations or deflations of ITPC except in the unusual case when power is exactly zero for an extended period of time (chapter 26 further illustrates this point).

In general, these findings suggest that power and ITPC are not necessarily coupled and that ITPC can be safely interpreted independently of power. Nonetheless, the potential relationship between power and ITPC should not be dismissed because of this one illustration. It is a good idea to examine the relationship between ITPC and power in your data before confidently interpreting ITPC and power in different ways. This could be done, for example, by showing qualitative differences between power and ITPC over time, frequency, electrodes, or conditions.

19.7 Weighted ITPC

Weighted ITPC (wITPC) is an extension of ITPC that provides a more direct link between phase angles and trial-varying behavior or experiment parameters (Cohen and Cavanagh 2011; Cohen and Voytek 2013). The wITPC addresses two limitations of interpreting ITPC with respect to cognitive processes. The first limitation is that ITPC can arise due to several task-related but not condition-specific factors, including stimulus-evoked responses and general orienting or attention responses. A second limitation of ITPC is that it precludes discovery of phase dynamics that are related to the task but are not clustered within the same range across trials. That is, if phase values are related to trial-varying behavior or experiment properties but are not related to stimulus onset, there may not be a significant ITPC.