


Because ERPs are embedded in a larger EEG signal, almost all ERP studies rely on some sort of averaging procedure to minimize the EEG noise, and the averaging procedure is typically accompanied by a process that eliminates trials containing artifacts or followed by some procedure to correct for artifacts. These procedures appear to be relatively simple, but there are many important and complex issues lurking below the surface that one must understand before applying them. This chapter will discuss the underlying issues and provide several practical suggestions for averaging and for dealing with artifacts.

The Averaging Process

Basics of Signal Averaging

Figure 4.1 illustrates the traditional approach to signal averaging. First, EEG-epochs following a given type of event (usually a stimulus) are extracted from the ongoing EEG. These epochs are aligned with respect to the time-locking event and then simply averaged together in a point-by-point manner. The logic behind this procedure is as follows. The EEG data collected on a single trial is assumed to consist of an ERP waveform plus random noise. The ERP waveform is assumed to be identical on each trial, whereas the noise is assumed to be completely unrelated to the time-locking event. If you could somehow extract just the ERP waveform from the single-trial EEG data, it would look exactly the same on every trial, and averaging together several trials would yield the



same waveform that was present on the individual trials. In contrast, if you could somehow extract just the noise from the EEG data, it would be random from trial to trial, and the average of a large number of trials would be a flat line at zero microvolts. Thus, when you average together many trials containing both a consistent ERP waveform and random noise, the noise is reduced but the ERP waveform remains.

As you average together more and more trials, the noise remaining in the averaged waveform gets smaller and smaller. Mathematically speaking, if R is the amount of noise on a single trial and N is the number of trials, the size of the noise in an average of the N trials is equal to $(1/\sqrt{N}) \times R$. In other words, the remaining noise in an average decreases as a function of the square root of the number of trials. Moreover, because the signal is assumed to be unaffected by the averaging process, the signal-to-noise (S/N) ratio increases as a function of the square root of the number of trials.

As an example, imagine an experiment in which you are measuring the amplitude of the P3 wave, and the actual amplitude of the P3 wave is 20 μV (if you could measure it without any EEG noise). If the actual noise in the EEG averages 50 μV on a single trial, then the S/N ratio on a single trial will be 20:50, or 0.4 (which is not very good). If you average two trials together, then the S/N ratio will increase by a factor of 1.4 (because $\sqrt{2} = 1.4$). To double the S/N ratio from .4 to .8, it is necessary to average together four trials (because $\sqrt{4} = 2$). To quadruple the S/N ratio from .4 to 1.6, it is necessary to average together sixteen trials (because $\sqrt{16} = 4$). Thus, doubling the S/N ratio requires four times as many trials and quadrupling the S/N ratio requires sixteen times as many trials. This relationship between the number of trials and the S/N

ratio is shown in Figure 4.1. The top waveform shows the raw EEG over a period of about 2 seconds, during which time two stimuli were presented. The left column shows segments of EEG for each of several trials, time-locked to stimulus onset. The right column shows the effects of averaging one, two, three, four, five, six, seven, or eight of these EEG segments. Negative is plotted upward.

◀ **Figure 4.1**