Psychophysiology, 30 (1993), 98–119. Cambridge University Press. Printed in the USA.

METHODOLOGY

Distortion of ERP averages due to overlap from temporally adjacent ERPs: Analysis and correction

MARTY G. WOLDORFF
Department of Neurosciences, University of California–San Diego, La Jolla

Abstract

In studies of event-related potentials (ERPs), short interstimulus intervals (ISIs) are often employed to investigate certain neural or psychological phenomena. At short ISIs, however, the ERP responses to successive stimuli may overlap, thereby distorting the ERP averages. This paper describes a signal processing approach for analyzing the distortion of ERP averages due to such overlap. In general, the distortion is modeled in terms of mathematical convolutions of the ERP waveform elicited by each type of adjacent stimulus with the corresponding distribution in time of those stimuli relative to the averaging epoch. Using this framework, a number of implications of ERP overlap for experimental design and interpretation are examined, with special emphasis given to selective attention paradigms. It is shown that the possibility of confound due to ERP overlap is widespread in short-ISI experiments, and even the widely used procedure of stimulus randomization does not necessarily control for differential distortion of the ERPs to attended versus inattended stimuli. Problems due to ERP overlap can be particularly serious in short-ISI studies that examine how ERPs (and associated perceptual processes) are influenced by the nature of the preceding stimulus (i.e., stimulus sequence effects). A set of algorithms is presented for estimating and removing the residual distortion due to response overlap from recorded ERP averages. The use of these algorithms, collectively termed the Adjacent Response (Adjar) Technique, can alleviate many of the overlap-related problems that arise when short ISIs are used, thereby enhancing the power of the ERP technique.

Descriptors: ERP, Overlap, Selective attention, Adjar, Sequence effects, Short ISIs, Methodology, Experimental design

In studies of event-related potentials (ERPs), a series of stimuli is presented, and ERP averages are obtained by signal-averaging epochs of EEG time-locked to the stimuli. In many experiments, stimuli are presented at a relatively slow rate such that there is little or no overlap of the ERPs elicited by successive stimuli in the series. At higher stimulus presentation rates, however, the ERPs elicited by successive stimuli can overlap in time, and this can result in distortion of the ERP averages.

Despite the potential for waveform distortion, there are many experimental situations in which short ISIs are useful or even required. One such case is the use of ERPs to investigate mechanisms of human selective attention. Subjective experience suggests that when stimuli from several competing "channels" of input (e.g., tone pips to the left and right ears) are presented relatively slowly, it is rather difficult to attend very selectively to the stimuli in one channel and to "tune out" the others, whereas a high rate of stimulus delivery seems to enable a more selective focusing of attention. This view is strongly supported by empirical ERP evidence indicating that the early differentiation of processing of attended and inattended stimuli is enhanced by, or even requires, a faster rate of stimulus presentation (Hansen & Hillyard, 1984; Hillyard, Hink, Schwent, & Picton, 1973; Parasuraman, 1978; Schwent, Hillyard, & Galambos, 1976; Woldorff, Hansen, & Hillyard, 1987; Woldorff & Hillyard, 1991).

Rapid rates of stimulus presentation may also lead to interactions in the processing of successive stimuli in the sequence that might not occur at slower rates. Such potential interactions could include, for example, effects on stimulus processing of being preceded in the sequence by one stimulus type versus another or being preceded at one ISI versus another. The ability to apply ERPs to the study of such sequential effects at fast stimulus rates could yield considerable insight into the mechanisms of sensory, perceptual, and cognitive processes. Consid-

This work was supported by an NSF Graduate Fellowship and a McDonnell-Pew Cognitive Neuroscience Postdoctoral Fellowship to the author, by ONR Contract N00014-89-J-1743 and NIMH Grant MH 25594 to Steven A. Hillyard, and by NINCDS Program Project Grant NS 17778 (M. S. Gazzaniga, Director).

I thank Steven A. Hillyard, Jonathan C. Hansen, Marta Kutas, George R. Mangun, Hans-Jochen Heinze, and Ken Paller for various helpful discussions during the development of this work, as well as for useful comments on earlier drafts.

Address requests for reprints to: M. Woldorff, University of Texas Health Sciences Center, Research Imaging Center, 7703 Floyd Curl Drive, San Antonio, TX 78284-6240.

erable problems can arise in analyzing such effects, however, when the responses elicited by those various preceding events overlap and distort the ERPs of interest.

Another basic reason for studying perceptual processing at short ISIs is that there are many real-life situations in which environmental stimuli occur (and must be processed) very rapidly. Not being able to simulate the conditions of these phenomena in the laboratory and apply ERPs to their study places additional constraints on the ERP technique. Use of short ISIs also allows more responses to be collected per unit time, which can facilitate the acquisition of averaged ERPs with higher signal/noise ratios.

The main purpose of this paper is to analyze the distortion of ERP averages that occurs when short ISIs are used and the ERP responses elicited by adjacent stimuli overlap in time. Although a number of theoretical articles have examined and modeled the ERP averaging process (e.g., Aunon, McGillem, & Childers, 1981; Brillinger, 1981; Hansen, 1983) and several have presented methods to remove overlap under certain conditions (e.g., Eysholdt & Schreiner, 1982; Hansen, 1983), the present paper is specifically focused on developing a framework for analyzing the problem of adjacent-response overlap and for examining its implications for experimental design and interpretation. In general, the emphasis here will be on selective attention paradigms, although the principles have broader applicability to other experimental questions and protocols, as well as to other physiological response measures besides ERPs.[1]

In addition to examining potential problems due to ERP overlap, a set of algorithms is presented for estimating and removing the distortion due to such overlap from recorded ERP averages. These algorithms, collectively termed the Adjacent Response (Adjar) Filter or Technique, have a number of applications for short-ISI experiments, and their use can enhance the effectiveness of ERPs for investigating neural and psychological processes. An application of the technique to the ERP data from two short-ISI selective attention experiments can be found in Woldorff and Hillyard (1991).

Methods That Have Been Employed to Deal With Overlap

Various approaches have been taken in attempt to resolve the problem of overlapping ERP responses. One approach has been to increase the high-pass cutoff of the filter, thereby attenuating the longer latency, lower frequency portion of the ERPs. This technique artificially "forces" the response to be over, or at least to appear to be over, by the time the next stimulus arrives, thus eliminating or reducing the overlap. Such high-pass filtering may achieve a reasonable solution if just the early high-frequency waves of the ERP are of interest (e.g., the auditory brain-stem-evoked responses). However, when the longer latency waves are of interest or when these contain significant power in the higher frequencies, simple high-pass filtering for this purpose may be of limited value.

A very different approach capitalizes on the overlap of adjacent ERP responses by delivering stimuli at a particular constant ISI selected to stimulate certain components of successive responses in phase, thus resulting in a sinusoidal steady-state response (SSR). Examples include visual steady-state responses (Regan, 1972, 1982) and the auditory 40-Hz response (Galambos, Makeig, & Talmachoff, 1981). These responses are useful probes of brain activity in both clinical and experimental contexts.

If, however, the goal is to obtain an ERP reflection of the processing of the individual occurrences of the stimuli in a sequence (i.e., the transient responses), a different solution is required. One common solution has been to randomly vary or "jitter" the ISIs across a range wide enough so that during averaging the overlapping adjacent responses tend to cancel each other out. The degree to which jittering the ISIs can mitigate the distortion of the final ERP averages will be examined in detail in the next section.

Another approach to the problem of overlapping responses has been to focus on the *difference* between the responses to a stimulus under two different conditions (such as when the stimulus is attended versus ignored). It has generally been assumed that if the stimuli are randomly presented, any distortion of the final ERP averages due to overlapping adjacent responses will be equivalent for each stimulus type. However, there is a subtle flaw in this reasoning, and there are circumstances under which differential waveform distortion from overlap can be produced despite stimulus randomization. This problem will also be considered in detail.

ISI Jitter as a Low-Pass Filter With a Time Shift

At stimulus rates where successive ERP waveforms overlap, every response included in the average (except the first and last in the sequence) will have superimposed upon it portions of the ERP response to the preceding stimulus and portions of the response to the succeeding stimulus. Randomly varying or jittering the ISIs around a mean value can partially cancel or "smear out" these overlapping adjacent responses, thereby mitigating the distortion of the final average. An empirical rule of thumb is that the effective jitter range needs to be larger than the period of the slowest dominant waves in the overlapping responses. Figure 1 shows ERPs from subjects performing a selective attention task with different ISI jitter ranges. The ERPs in Figures 1a and b are clearly distorted by overlapping adjacent responses, with those in Figure 1a being dramatically worse because of the much narrower jitter range (60 ms versus 150 ms). In Figure 1c, however, where the range of ISIs was the widest (120–320 ms), the waveforms appear relatively undistorted. Note that the 200-ms jitter in this case was, indeed, larger than the period of the dominant waves in these responses.

Given a few assumptions, the effect of ISI jitter on the overlap from adjacent responses can be approximated as a low-pass filtering operation on the adjacent response with either a negative or positive time shift. To consider this process more quantitatively, a number of terms must be defined and certain assumptions detailed.

Event Distributions

Consider a hypothetical experiment in which there is only one type of stimulus, S, eliciting only one type of response, an ERP waveform $R(t)$. Let the ISI range be 200–400 ms, thereby yielding a jitter width, T_{jw}, = 200 ms. Assume that $R(t)$ is always

[1] In this paper, unless otherwise noted, the term *response* will be used in the physiological sense (in contrast to the behavioral sense, such as a *button press response*). Although the context will usually be ERP responses, the analyses and discussions will generally be applicable to any physiological response—that is, to any physiological activity that is associated in time with a stimulus event and is reflected by a time-extended waveform measured by the experimenter.

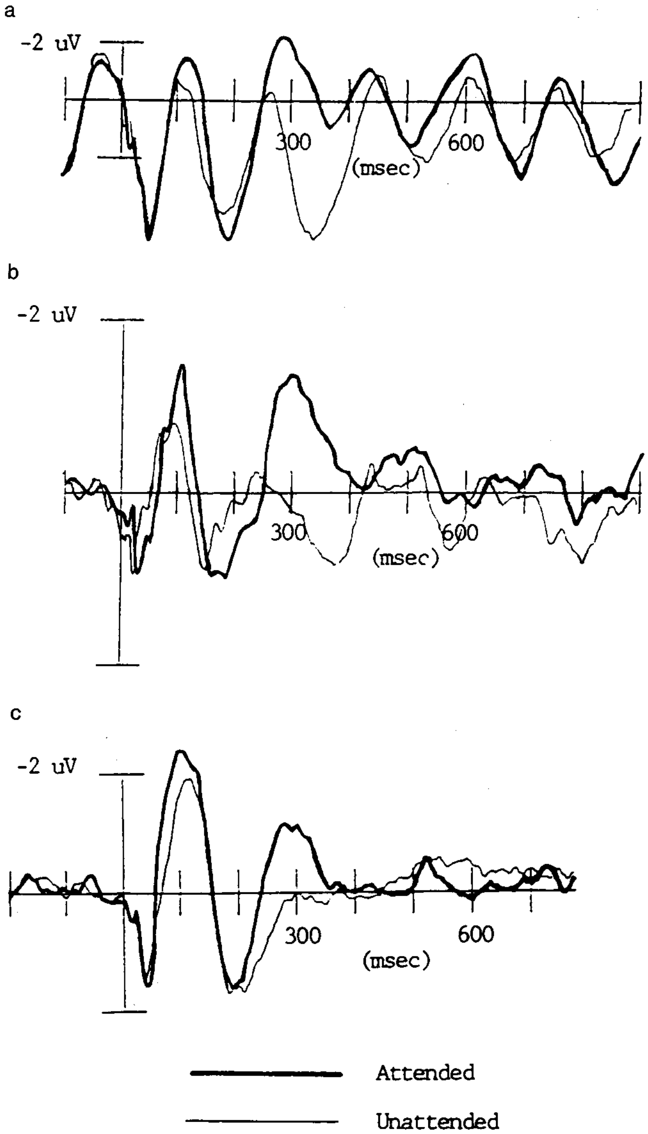

Figure 1. Effects of ISI jitter range on the averaged ERP response: (a) 130–190 ms, (b) 125–275 ms, (c) 120–320 ms. Responses are from single subjects performing a selective attention task.

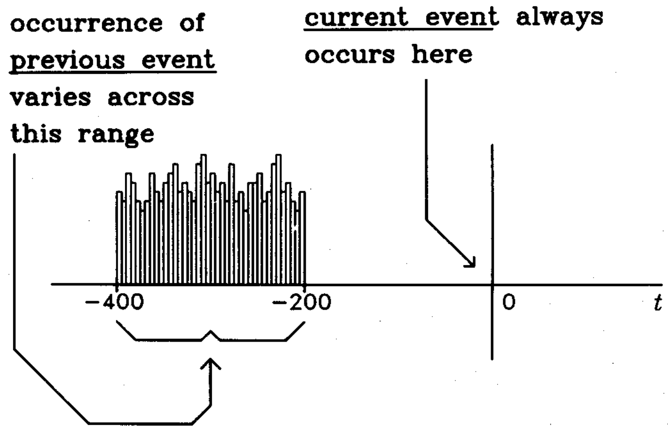

Figure 2. Event distributions. Prestimulus histogram indicating the temporal distribution of previous events relative to the current one (i.e., at $t = 0$) during time-locked averaging at ISIs of 200–400 ms.

approximately the same no matter when it occurs (i.e., time invariance) and that the electrical fields of overlapping responses add linearly (i.e., superposition).

The ERP averaging process involves time-locking of the reference epoch to the occurrence of each stimulus in turn. When a particular stimulus is time-locked to the averaging epoch, it is the *current stimulus*, occurring at $t = 0$. Relative to this, the immediately preceding event must have occurred between 200 and 400 ms before, because that is the specified ISI range. Figure 2 shows a possible *previous-event distribution*, $D_p(t)$, which is a time histogram of the occurrences of the previous stimuli relative to the current stimuli that are included in the average. [There is also an analogous *subsequent-event distribution*, $D_s(t)$, not presently being considered.] The event distribution in Figure 2 appears approximately rectangular but is not perfectly flat because it is meant to represent a typical histogram of actual occurrences of the previous stimuli in an experiment, rather than a theoretical probability distribution.

To keep the responses to current and previous stimuli distinct in this discussion, the response waveform for the current stimulus will be referred to as R_c and the response waveform for the previous stimulus as R_p. Because we are assuming that there is only one stimulus and response type in the present discussion, R_c and R_p will always have the same waveshape, R. In the more general case, however, where a number of different stimulus types and response types occur, R_c and R_p would not necessarily have the same waveshape. For example, in a typical auditory attention experiment in which the delivery of left- and right-ear stimuli is randomized, one might want to obtain an ERP average of all left-ear responses (i.e., the R_c's) that were preceded by right-ear responses (i.e., the R_p's). However, we will continue to assume for the present discussion that the R_p's are all basically identical to one another.

In almost all applications relevant here, response waveforms would be sampled (digitized) at discrete time points, and event distributions would be calculated as histograms. Thus, these functions are assumed throughout this paper to be discrete functions of time. In general, the most appropriate bin size of the event distribution histograms would be the same as the sample period of the digitization.

The Averaging Process

Consider what happens when the responses to 1,000 stimuli are averaged. The EEG averaging epochs, time-locked to each of the 1,000 stimuli, are summed (time point by time point), and the result (for each time point) is divided by 1,000, yielding what will be called here the *recorded ERP*. Let us assume that essentially all of the random background EEG averages out to zero (i.e., to a flat line). Because the response $R_c(t)$ was elicited consistently at $t = 0$, the resultant recorded ERP will include an average of 1,000 such responses.

However, the overlapping portions of the previous responses will also be included in the average. Unlike random ongoing EEG, however, they do not necessarily average out or approach zero with enough trials. An intuitive picture of the *averaged previous response* (AvR_p) can be gained as follows:

1. Shift waveform R_p to begin at a particular time point of the event distribution $D_p(t)$.
2. Weight it by the value of $D_p(t)$ at that point (i.e., multiply it by the number of times the previous response began at that point).
3. Do this at each time point of $D_p(t)$.
4. Add all the weighted waveforms together.
5. Divide by 1,000.

Because it has been assumed that the duration of these responses is longer than the ISI, the later portion of the AvR_p will overlap and distort the average of the current responses (the R_c's). This distortion may or may not be negligible relative to the remaining EEG noise, but it is not random activity that will average out to zero. In fact, with enough trials and subjects, it would almost invariably become statistically significant.

Normalized Event Distributions

The event distribution in Figure 2 of the actual number of occurrences of the previous events can easily be *normalized*, such that each value of the $D_p(t)$ function represents the proportion of trials in which the previous event occurred at that particular time point. To use such a normalized distribution to determine the AvR_p in the 1,000-stimulus example, a response waveform R_p would still be placed to begin at each previous time point and would be weighted by the value of $D_p(t)$ at that point, and all the weighted waveforms would be added together. The only difference would be that one would not then divide by 1,000, because this division already would have been taken into account in the normalization of $D_p(t)$.

Convolutions

The weighting and summing process described above is equivalent to a mathematical *convolution*. In particular, the averaged previous response (AvR_p) is also equal to the convolution of the functions $D_p(t)$ and $R_p(t)$, written

$$D_p(t) * R_p(t)$$

where $D_p(t)$ is the normalized previous-event distribution, $R_p(t)$ is the response waveform elicited by those previous events (in a reference frame before shifting backward in time), and the asterisk (*) is the traditional symbol for the convolution operation.

The formal mathematical definition of a (discrete) convolution commonly used in signal processing (see, e.g., Oppenheim and Schafer, 1975) is

$$C(t) = D_p(t) * R_p(t) = \sum_{k=-\infty}^{\infty} D_p(k) R_p(t-k).$$

This expression is equivalent to the intuitive weighting and summing process presented above, the only difference being the order in which terms are calculated and added.

As discussed below, the convolution $D_p * R_p$ is essentially equivalent to low-pass filtering of the $R_p(t)$ waveform and then shifting it backward in time. This filter operation can be examined conceptually in the time domain by thinking in terms of mutual cancellation of positive and negative values of the waveform as R_p shifts around in time. Consider a high-frequency component of R_p, one whose period is shorter than T_{jw} (the ISI jitter width), such as shown in Figure 3a. As R_p shifts across T_{jw} (Figure 3b), the positive and negative phases of such a component cancel each other fairly well. However, for a much lower frequency component (Figure 3c), the positive and negative phases would not cancel nearly as well for the same amount of shifting (Figure 3d). Thus, jittering the ISIs tends to cancel the higher frequencies of the overlapping adjacent responses but leaves the lower frequencies relatively unaffected.

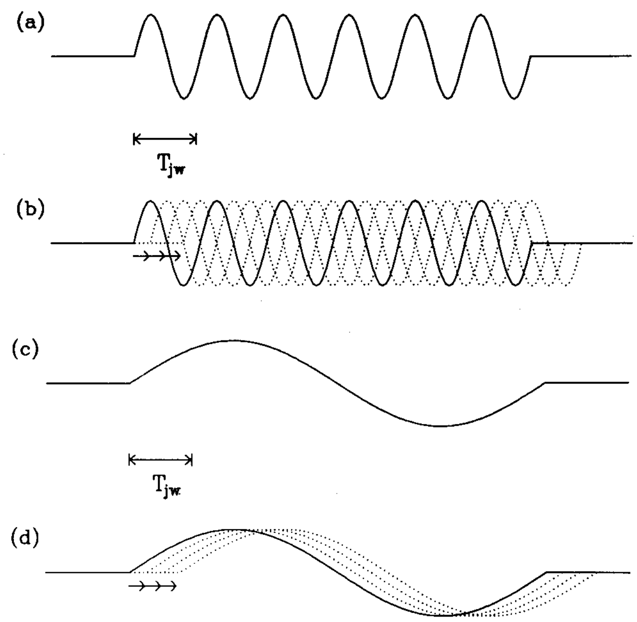

Figure 3. Schematic representation in the time domain of the low-pass filtering of adjacent responses that results from jittering the ISI. For a relatively high-frequency component (a), when the adjacent response is shifted across the jitter range, T_{jw}, the positive and negative phases cancel each other fairly well (b). However, for a much lower frequency component (c), the positive and negative phases do not cancel nearly as well (d) for the same amount of shifting.

A more precise examination of the low-pass filtering effects is possible in the frequency domain. The basic law of transforms between the time and frequency domains specifies that "a convolution in time is equivalent to a multiplication in frequency." That is,

	Time domain		Frequency domain
if the transform of	$D_p(t)$	is	$D'_p(f)$,
and the transform of	$R_p(t)$	is	$R'_p(f)$,
then the transform of	$D_p(t) * R_p(t)$	is	$[D'_p(f)][R'_p(f)]$.

If $D_p(t)$ is put into the reference frame shown in Figure 4a, one can calculate that the Fourier transform, $D'_p(f)$, is the *sinc* function shown in Figure 4b.[2] Because our interest is in examining the filtering effects on the previous response $R_p(t)$, consider just the magnitude of this transform at positive frequencies (Figure 4c). This is the *gain function* of the filtering operation,

[2]A *sinc* function is one that is governed by an equation of the form $y = \sin(x)/x$.

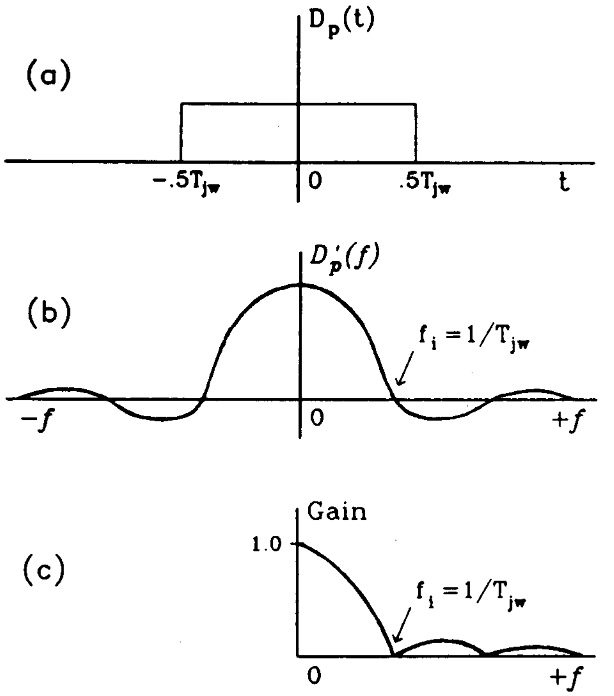

Figure 4. Frequency domain representation of the low-pass filtering of adjacent responses that results from jittering the ISI. (a) Adjacent-event distribution (of width T_{jw}) placed in a new reference frame so that it is centered around time 0. (b) Fourier transform (amplitude only) of the event distribution in (a). (c) Gain function (as a function of frequency) of the low-pass filtering effects on adjacent responses.

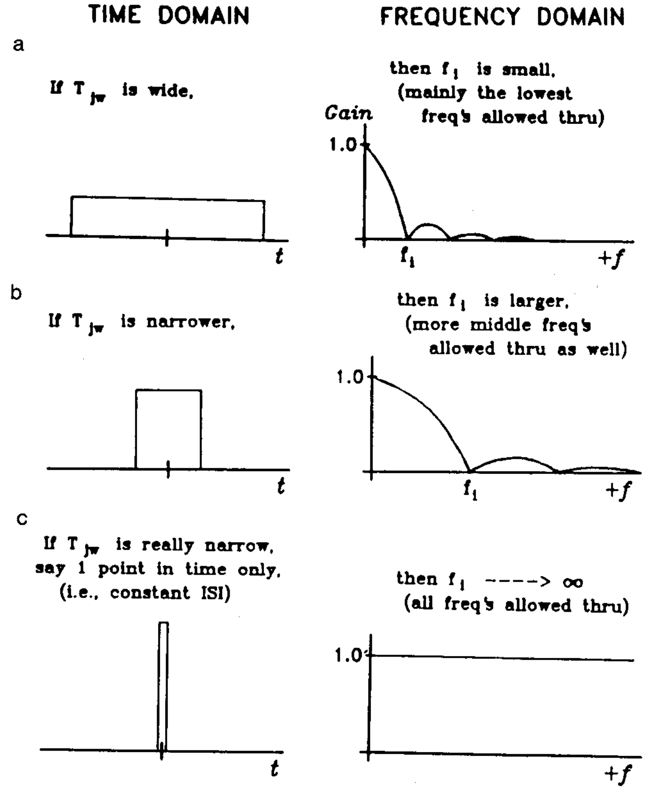

Figure 5. Relationship between the ISI jitter width, T_{jw}, and the low-pass filtering effects on adjacent responses.

representing the proportions of the various frequency components of the original signal (i.e., the previous response, R_p) that are still present after filtering; this gain function clearly resembles that of a low-pass filter. Most of the power that is allowed through is at frequencies less than $f_i = 1/T_{jw}$, a value closely related to the high-frequency cutoff of this low-pass filter. The final average of the overlapping previous responses, which is what ultimately constitutes the distortion of the current response average, is equivalent to passing the previous response through this low-pass filter and then shifting it to the left (i.e., backward) in time.

The relationship between the jitter width and the low-pass filtering function is shown schematically in Figure 5. If T_{jw} is wide (Figure 5a), then $f_i = 1/T_{jw}$ is small, and predominantly only the lower frequencies are allowed through. If T_{jw} is narrower (Figure 5b), f_i is larger, thereby allowing more of the higher frequencies through also. If T_{jw} is very narrow, such as when the ISI is constant, then f_i approaches infinity, and all frequencies are allowed through (Figure 5c). In this last case, the distortion of the current ERP average by the overlapping previous responses is not mitigated at all by any low-pass filtering; the average overlap is simply the latter part of the previous response shifted to the left.

The above analysis clarifies the basis of the empirical "rule of thumb" mentioned earlier regarding the choice of the jitter width, T_{jw}. If T_{jw} is as wide or wider than the period of the dominant waves of the overlapping responses, then $1/T_{jw} = f_i$ is lower than most of the frequency content of those responses, and the overlap will be substantially filtered or "smeared out."

Overlap From Subsequent Versus Previous ERP Responses

Given these assumptions, a similar analysis would hold for overlap resulting from subsequent ERP responses, with the shifts in time to the right rather than the left. However, some important practical differences make the effects of overlap from previous ERPs more problematic, and these will be the primary focus of this paper.

For subsequent ERPs, it is mainly the early part that overlaps the current ERP. The early components of most ERPs tend to be of fairly low amplitude and high frequency. Thus, the distortion from these components will generally be considerably attenuated after the low-pass filtering effects of jittering the ISIs. Furthermore, any residual distortion from such jittered subsequent ERPs overlaps the latter part of the current ERPs, which tends to be of higher amplitude and thus less susceptible to distortion. By contrast, the latter part of previous ERP responses overlaps the current ones. Because late waves tend to be of higher amplitude and lower frequency, they will not be filtered out as well by the ISI jitter. Moreover, they will overlap the early waves of the current ERP, where effects of interest would generally be small in magnitude.

Second-Order Event Distributions

Thus far, only adjacent events that were first order—that is, those that were immediately preceding or following—have been considered. In most situations, this is sufficient because the responses elicited by more remote events would not overlap R_c. However, overlap from these second-order (or higher) adjacent

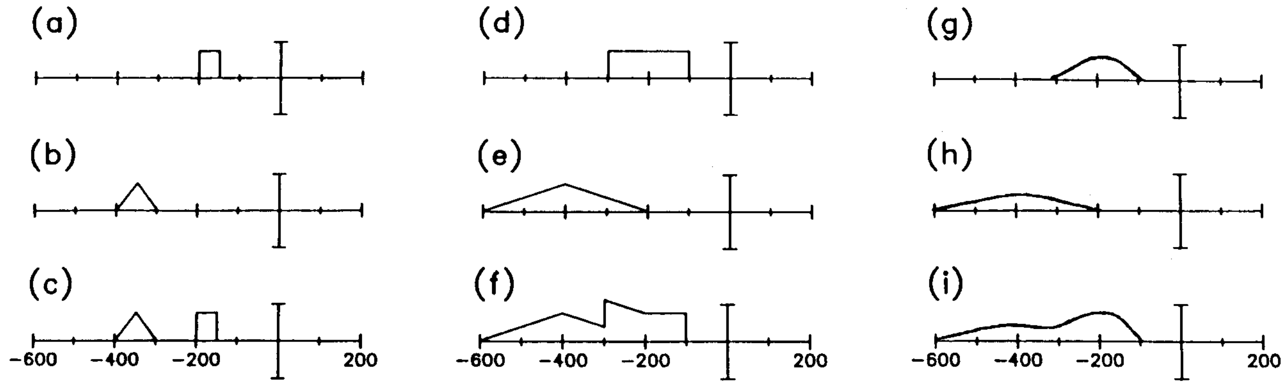

Figure 6. Temporal relationship of the first- and second-order previous-event distributions for three different ISI ranges. Left column plots (a, b, c) are for ISIs of 150–200 ms with a rectangular distribution, middle column (d, e, f) for ISIs of 100–300 ms with a rectangular distribution, and right column (g, h, i) for ISIs of 100–300 ms with a Gaussian distribution. Top row plots (a, d, g) indicate the first-order previous-event distributions (i.e., that of the immediately previous events), middle row (b, e, h) the corresponding second-order distributions (i.e., that of the events that occurred two previously), and bottom row (c, f, i) the total distributions.

events can become relevant. For example, for a range of ISIs from 150 to 200 ms with a rectangular distribution, the immediately preceding stimulus will have occurred 150–200 ms prior to any current event being averaged (Figure 6a). However, the stimulus *before that* could have occurred anywhere between 300 and 400 ms prior to the current event. The event distribution for this second-order event, relative to the current event, turns out to be the convolution of the first-order rectangular distribution with itself and shifted over, producing the triangular distribution shown in Figure 6b. Depending upon the particular responses elicited by the second-order event, as well as on the amplifier settings used, this response could overlap R_c substantially. The relevant previous-event distribution would then be the sum of the first- and second-order distributions (Figure 6c), and the total previous-response overlap would be the convolution of the previous response with this total distribution.

In the foregoing example, the first- and second-order distributions did not overlap in time. However, consider a rectangular ISI range of 100–300 ms (Figure 6d). In such a case, the second-order previous event could occur anytime between 200 and 600 ms before the current event (Figure 6e), and the first- and second-order distributions would overlap. The total previous-event distribution, however, would still be the sum of the two (Figure 6f).

Such overlap of the first- and second-order event distributions can widen the effective ISI jitter (e.g., compare Figures 6d and f), which can enhance to some extent its effectiveness as a low-pass filter of the adjacent responses. The effectiveness of this low-pass filtering can be enhanced even further, however, if this total event distribution were smooth throughout its extent (rather than having, for example, the sharp edge at −300 ms). One way to obtain a smoother total distribution would be to use first-order distributions that are Gaussian (Figures 6g–i), rather than rectangular, an approach that could be especially useful at very short ISIs.

Stimulus Randomization Does Not Control for Differential Overlap

Another common approach to the problem of overlapping responses has been to concentrate on the difference between the ERP waveforms elicited by a stimulus under two different conditions, rather than on the original waveforms themselves. For example, in a selective attention paradigm, the averaged ERP to a stimulus when it is attended is compared with the averaged ERP to that same stimulus when it is ignored. It is generally assumed that if a stimulus is equally likely to be preceded (and succeeded) by an attended stimulus as by an inattended stimulus, then any distortion from overlap will be about the same for the attended and inattended ERPs. It is further assumed that this justifies attributing any differences between these ERPs to the differential processing of the stimulus, rather than to differential overlap. There is a flaw in this argument, however.

Although stimulus randomization has implications for adjacent-ERP overlap, the reasons for its widespread use arose from other considerations. In particular, random presentation of relevant and irrelevant stimuli in selective attention experiments enables ruling out differential *preparation* as a confound (Hillyard et al., 1973; Näätänen, 1967, 1975); if subjects have no information to help them predict the next stimulus, they cannot differentially prepare for it. However, the assumption that stimulus randomization also controls for differential overlap from adjacent responses is not always valid. Although stimulus randomization does imply that any single stimulus is equally likely to be both preceded and succeeded by all other physical stimulus types, the responses to those various adjacent stimuli may well change under the different experimental conditions (e.g., direction of attention).

Consider a two-channel selective attention paradigm within the visual modality, in which the two stimulus channels (i.e., classes of stimuli to be attended or ignored) are left-field flashes (L's) and right-field flashes (R's). Assume that (a) stimulus randomization was perfect, (b) the ERP responses added linearly, (c) the ERP responses were time invariant, and (d) there were no other types of stimuli in the sequence (such as targets). We will focus on overlap from preceding ERPs.

The critical comparison in these experiments is generally between the ERP elicited by a stimulus when it is attended and the ERP elicited by the same physical stimulus when it is ignored (i.e., when the other channel is being attended). Because the physical stimulus is identical and stimulus randomization controls for differential preparation, any differences in the ERPs

	Previous Stimuli	Current Stimulus	Previous Stimuli	Current Stimulus
a	(L) (R)	L	(L) (R)	R

	Previous Responses	Current Attended Response	Previous Responses	Current Inattended Response
b	(A) (I)	A	(A) (I)	I
c	(L+) (R−)	L+	(L−) (R+)	L−
d	(L−) (R+)	R+	(L+) (R−)	R−
e	(C+) (F−)	C+	(C−) (F+)	C−

Figure 7. Stimulus/response types for current and previous events during various two-channel selective attention paradigms. L = left stimulus, R = right stimulus, A = attended response, I = inattended response, C = click, F = flash. A plus sign (+) indicates the attended response to that particular physical stimulus type, whereas a minus sign (−) indicates the inattended response.

are presumed to be consequences of the selective internal shifting of attention either toward or away from the stimulus. The question of concern here is whether any part of the difference between the averaged ERPs to attended and inattended stimuli possibly could have arisen from differential previous-response overlap.

Figure 7 enumerates the various combinations of current and prior events that can occur in the randomized sequence. Figure 7a shows the alternative physical stimulus types and represents successful stimulus randomization: L's were equally likely to be preceded by L's as by R's, and, likewise, R's were equally likely to be preceded by L's as by R's. The rest of Figure 7 (b–e) presents various response types. The left side indicates the responses to current stimuli when they were attended along with the possible previous response types that could have occurred in that condition, and the right side indicates the corresponding current and previous response types when those same current stimuli were ignored.

An important but misleading ramification of successful stimulus randomization is shown in Figure 7b; regardless of whether a stimulus was attended (left side) or inattended (right side), it was equally likely to be preceded by attended responses (*A*'s) as by inattended responses (*I*'s). Although correct, this statement is misleading because it is incomplete. It leaves out that the preceding attended responses on the left side of Figure 7b do not represent the same waveforms as the preceding attended responses on the right side, and, similarly, the preceding inattended responses on the left and right sides are not identical. In Figures 7c and 7d, the attended responses to left and right stimuli are represented more specifically by the symbols L+ and R+, respectively, and the inattended responses by L− and R−. When the left stimuli are being attended (left side, Figure 7c), they are preceded by, and partially overlapped by, L+'s (attended left responses) and R−'s (inattended right responses). However, when the left stimuli are being ignored (right side, Figure 7c), they are preceded by L−'s (inattended left responses) and R+'s (attended right responses). Similar considerations apply for the current right-field flashes (Figure 7d). The problem is now clear: L+'s may not be the same as R+'s, even though they are both attended responses, and, likewise, L−'s may not be the same as R−'s even though they are both inattended responses.

In a visual experiment at a midline recording site, the ERPs labeled L+ and R+ may be quite similar to one another, as might the two inattended responses. However, for lateral sites, several flash-evoked ERP components tend to be larger contralateral to the side of the flash (e.g., Mangun & Hillyard, 1987, 1988). Thus, at a right occipital site (e.g., O2), L+'s may be substantially larger than R+'s, whatever other differences may exist between them. The overlap, therefore, on the current L+ response from previous attended responses (i.e., L+'s) will not be the same as the overlap on a current L− from previous attended responses (i.e., R+'s). By the same reasoning, the overlap from the corresponding previous inattended waveforms will not be equivalent. Thus, despite complete stimulus randomization, the current attended and inattended ERPs would be overlapped by different patterns of preceding waveforms.

Similar problems will arise at lateral sites in intramodal experiments in other modalities where the ERPs are highly lateralized (e.g., somatosensory). Auditory selective attention experiments, on the other hand, involve ERP responses that are not nearly so lateralized, and therefore would be less likely to produce this artifact.

Cross-modal attention experiments are particularly vulnerable to these problems, however, because of the substantial differences between the ERPs in the two modalities. Take, for example, an auditory/visual design in which the two stimulus channels are clicks (C's) and flashes (F's), and consider the prior overlap on click ERPs (Figure 7e). When the clicks are attended and therefore elicit current C+'s, the previous *A*'s and *I*'s (Figure 7b) are C+'s and F−'s, respectively, whereas when the clicks are ignored and therefore elicit current C−'s, the previous *A*'s and *I*'s are F+'s and C−'s. These two combinations of preceding waveforms could easily be different at electrode sites all over the head, including those on the midline. For example, at the Oz site, ERPs to clicks are small, regardless of whether they are attended or not. C+'s and C−'s would therefore be small at this site (relative to the flash ERPs), and most of the overlap on the current C+'s would come from F−'s, and most of the overlap on the C−'s would come from F+'s. Because there is certainly a strong attention effect on flashes at occipital sites, the F−'s and F+'s would be quite different. Therefore, the ERP averages for attended versus inattended clicks could be quite different at Oz solely because of differential previous-ERP overlap, thereby giving the false impression of an attention effect at that site on the auditory ERPs. Analogous errors could be induced at other scalp locations.

The foregoing analysis illustrates how, despite complete stimulus randomization, an actual attention effect upon the ERP to a certain stimulus at a particular latency can produce artifactual attention effects in the ERPs to other stimuli or in the ERP to the same stimulus at the wrong latency. Perhaps one of the greatest dangers of such overlap is that it could masquerade as an attention effect at a very short latency. Even if such artifactual effects are small, the early ERP components are also small, and any attention effects on them would probably also be small. Given that attention effects on early components have significant implications for human information processing models, the potential problem of differential previous-response overlap may be important to consider whenever ERPs are used to probe the mechanisms of selective attention to streams of short-ISI input.

Quantification of Overlap From Previous Responses

In this section a scheme will be presented for mathematically describing the distortion of ERP averages due to overlapping previous responses in experiments where the ISI has been jittered and there is more than one type of previous response. Such a quantification enables an analysis of the differential overlap distortion of ERPs that are to be compared for experimental differences. As an example, a simple two-channel selective attention experiment will be examined in detail.

In a two-channel selective attention experiment a subject receives two classes of stimuli, which will be referred to here as 1's and 2's. These 1's and 2's could be, for example, left-field visual flashes and right-field visual flashes or left-ear clicks and right-ear clicks, or flashes and clicks if the experiment is cross-modal. The two attention conditions are (a) attend to the 1's and (b) attend to the 2's.

Rather than assume that stimulus randomization was successfully accomplished, cases when it was and when it was not will be examined. It will be assumed, however, that identical stimulus sequences were presented for both attention conditions, a procedure often implemented as part of counterbalancing.

In a stimulus sequence of 1's and 2's, there are four event distributions relevant to the present analysis:

$D_{11}(t)$ = the event distribution of 1's preceding 1's

$D_{21}(t)$ = the event distribution of 2's preceding 1's

$D_{12}(t)$ = the event distribution of 1's preceding 2's

$D_{22}(t)$ = the event distribution of 2's preceding 2's.

Because the same stimulus sequence is assumed to be presented for the two attention conditions, each of these four event distributions is the same for both attention conditions. However, because successful stimulus randomization is not being assumed here, these event distributions are not necessarily equal to each other. Appropriately normalized versions of these distributions are used in the analyses.

Employing two attention conditions results in four types of ERP responses:

When the 1's are being attended, there are A_1's (responses to attended 1's) and I_2's (responses to inattended 2's). When the 2's are being attended, there are I_1's (responses to inattended 1's) and A_2's (responses to attended 2's).

Again assume that the waveforms of each of these responses are time invariant.[3]

These definitions allow mathematical modeling of the averaged previous-response overlap distorting each of the four types of current ERPs. As described earlier, if the set of previous responses (i.e., the R_p's) consists of essentially one type, the averaged previous response (AvR_p) can be modeled as the convolution of the previous-event distribution with the previous-response waveform. In the present case, because each current response can be preceded by (and partially overlapped by) two kinds of responses, the averaged previous response would be equal to the sum (or weighted average) of two convolutions. For example, A_1's (responses to attended 1's) can be preceded by A_1's and I_2's. Thus, the averaged previous response for A_1's is[4]

$$AvR_p \text{ for } A_1\text{'s} = (D_{11} * A_1) + (D_{21} * I_2). \tag{1}$$

Similarly, the averaged previous response for I_1's is

$$AvR_p \text{ for } I_1\text{'s} = (D_{11} * I_1) + (D_{21} * A_2). \tag{2}$$

The critical comparison in this type of experiment is between the ERP elicited by stimuli when they were attended and the ERP elicited by those same stimuli when they were inattended; thus, for example, the average of the A_1's and that of the I_1's would be compared. The main interest here, therefore, is to evaluate the differential distortion of these two ERP averages due to the overlapping previous responses. To do this, the difference between the averaged previous response for the A_1's and that for the I_1's is examined; this difference we will call the *differential averaged previous response* (*Diff.* AvR_p):

$$\begin{aligned} &\textit{Diff. } AvR_p \text{ for 1's} \\ &\quad = (AvR_p \text{ for } A_1\text{'s}) - (AvR_p \text{ for } I_1\text{'s}) \\ &\quad = (D_{11} * A_1) + (D_{21} * I_2) - (D_{11} * I_1) - (D_{21} * A_2) \\ &\quad = [(D_{11} * A_1) - (D_{11} * I_1)] - [(D_{21} * A_2) - (D_{21} * I_2)]. \end{aligned}$$

Because convolution obeys the distributive property, this can be simplified to

$$\textit{Diff. } AvR_p \text{ for 1's} = [D_{11} * (A_1 - I_1)] - [D_{21} * (A_2 - I_2)]. \tag{3}$$

[3]As described for the Adjar technique, the results for previous-response overlap generally change little when this assumption is violated. (Certain analogous analyses for subsequent-response overlap could be somewhat affected, however, depending on the violation—see Woldorff, 1989, Appendix 2C).

[4]Note that all variables in this expression, including AvR_p for A_1's, are functions of time, although the time function symbolism has been dropped for simplicity of expression. The more precise expression for Equation 1 would be: $AvR_p(t)$ for A_1's = $[D_{11}(t) * A_1(t)] + [D_{21}(t) * I_2(t)]$. Also note that D_{11} and D_{21} are assumed here to have been normalized with the same constant (i.e., the total number of current 1's), so that the two convolutions in this sum are appropriately weighted relative to each other.

The differential previous-response overlap for the 1's (i.e., the differential distortion of the ERP average for attended 1's and the ERP average for inattended 1's) is the longer latency portion of this differential averaged previous response. Thus, the differential overlap distortion for the 1's can be expressed as the longer latency portion of the difference between two convolutions, namely the convolution of D_{11} with the attention effect for 1's and the convolution of D_{21} with the attention effect for 2's.

Similarly, the differential overlap for 2's can be expressed as the later portion of

$$\textit{Diff. } AvR_p \text{ for 2's} = [D_{22} * (A_2 - I_2)] - [D_{12} * (A_1 - I_1)]. \tag{4}$$

Equations 3 and 4 encapsulate the differential previous-response overlap that can occur in conventionally collected ERP averages in a two-channel selective attention experiment. By setting conditions on some of the variables involved, these equations can be simplified to examine the differential overlap in various experimentally relevant cases.

Case 1

There are no attention effects for either channel. That is,

$$A_1 - I_1 = 0 \quad \text{(or more precisely: } A_1(t) - I_1(t) = 0\text{)},$$

and

$$A_2 - I_2 = 0 \quad \text{(or more precisely: } A_2(t) - I_2(t) = 0\text{)}.$$

Clearly, if there are no attention effects for either channel, Equations 3 and 4 both reduce to zero, and there will be no differential AvR_p.

Case 2

All event distributions are equal:

$$D_{11} = D_{21} = D_{12} = D_{22} = D;$$

and the attention effects for the two stimulus channels are equal:

$$(A_1 - I_1) = (A_2 - I_2) = (A - I).$$

Substituting in Equations 3 and 4,

$$\textit{Diff. } AvR_p \text{ for 1's} = [D * (A - I)] - [D * (A - I)] = 0,$$

$$\textit{Diff. } AvR_p \text{ for 2's} = [D * (A - I)] - [D * (A - I)] = 0.$$

Under these conditions there is also no differential AvR_p and, hence, no differential overlap. The condition that all the event distributions are equal will be met when stimulus randomization and ISI randomization have been successful. However, to assure no differential overlap, the additional condition that the attention effect is symmetric (i.e., equal) for the two stimulus channels must also be true.

Case 3

All event distributions are equal:

$$D_{11} = D_{21} = D_{12} = D_{22} = D;$$

but the attention effects for the two stimulus channels are unequal:

$$(A_1 - I_1) \neq (A_2 - I_2).$$

Then, substituting into Equations 3 and 4:

$$\begin{aligned}\textit{Diff. } AvR_p \text{ for 1's} &= [D * (A_1 - I_1)] - [D * (A_2 - I_2)] \\ &= D * [(A_1 - I_1) - (A_2 - I_2)];\end{aligned}$$

$$\begin{aligned}\textit{Diff. } AvR_p \text{ for 2's} &= [D * (A_2 - I_2)] - [D * (A_1 - I_1)] \\ &= D * [(A_2 - I_2) - (A_1 - I_1)] \\ &= -(\text{Diff. } AvR_p \text{ for 1's}).\end{aligned}$$

This case is especially important because there are many experimental situations in which the attention effects for the two stimulus channels will not be equal to each other at every electrode site. The above analysis confirms that in such cases, despite perfect randomization of stimuli and ISIs, there can be differential overlap distorting the ERP averages; moreover, the analysis specifies that the differential overlap for one of the two channels will be approximately equal in magnitude and shape but of opposite polarity to the differential overlap for the other channel. In addition, the magnitude of the differential overlap is proportional to the asymmetry of (i.e., difference between) the attention effects for the two stimulus channels.

As can be seen in the above equations the amount of differential overlap is determined not by the degree of asymmetry between the ERP responses themselves for the two stimulus channels but rather by the degree of asymmetry of the attention effect. However, those experiments with significantly different responses for the two stimulus channels also tend to have unequal attention effects. (For cross-modal experiments this is obvious; for lateral asymmetries during intramodal visual attention see, for example, Mangun & Hillyard, 1988.) Thus, these situations are quite vulnerable to such artifacts.

Also, the differential overlap is not directly equatable with the asymmetry of the attention effects but arises from the convolution of this asymmetry with the event distribution. This is therefore equivalent to low-pass filtering the asymmetry of the attention effects and shifting the result backward in time. Thus, not only does jittering the ISI mitigate adjacent-response overlap in general, it also directly mitigates the risk of differential overlap upon attended and inattended waveforms. In contrast, using a constant ISI in experiments with unequal attention effects for the two stimulus channels clearly carries with it a greater risk of this kind of artifact.

The foregoing analysis implies that the observation of presumed attention effects in the two channels that are similar in waveshape but of opposite polarity could be an important clue that differential overlap has occurred. However, because differential overlap from previous responses would summate with any real attention effects on the current response, the opposite polarity clue could be easily camouflaged. Thus, this clue may be most evident during the prestimulus baseline period, where no real attention effects could exist.

Even prestimulus effects, however, may be obscured under certain conditions. Consider the case where the prestimulus baseline period is fairly short and the differential overlap is of rather low frequency. In such a case the baselining of the current ERP, which sets the voltage to zero during the prestimulus period (or some portion thereof), will tend to negate differences in this period and displace the differential overlap, and therefore the artifactual "attention effect," toward the longer latencies of the current waveforms. This possibility underscores the more general point that including a substantial length of prestimulus period in the averaging epoch is important in checking for previous-response overlap, whether differential or not.

Case 4

Event distributions (D's) are not all equal, but the attention effects for the two stimulus channels are equal:

$$(A_1 - I_1) = (A_2 - I_2) = (A - I).$$

There are many different ways in which the event distributions might be unequal. For example, event distributions could be skewed toward different ends of the ISI range (Figure 8a), or they could have different overall frequencies (Figure 8b).

One reason event distributions might show such inequalities is simply an inadequate number of trials. If the number of trials is low, just specifying a particular ISI probability distribution (e.g., rectangular) and the relative probabilities of stimulus types may not necessarily result in actual ISI distributions and stimulus ratios very close to the desired values. If one particular stimulus presentation sequence that is biased in some way is used repeatedly (such as across subjects), the distributions would then be unequal in a systematic way. One way to mitigate this possibility, therefore, is to employ a number of different stimulus presentation scenarios.

Event distributions may be unequal even when the number of trials is large, however. For example, say production of a random sequence of 1's and 2's in equal proportion is desired. To also control for extremes in the *local probability* of the sequence, one might add the restriction that no more than, say, four of one type of stimulus could occur in a row. However, such a restriction will slightly undermine the randomness of the random sequence generator, in that it will result in a slightly greater probability for a "switch" of stimulus type in the sequence (i.e., 1,2 or 2,1) rather than for a stimulus repetition or "same" (1,1 or 2,2). Thus, the D_{12} and D_{21} distribution values will be slightly larger on the average than those of D_{11} or D_{22} (Figure 8b). This case of unequal numbers of sames and switches can produce particularly insidious artifacts, as shown below.

Subcase (of Case 4). Number of sames is not equal to the number of switches:

$$D_{11} = D_{22} = D_{sa}$$

and

$$D_{12} = D_{21} = D_{sw} \neq D_{sa},$$

but the attention effects for the two stimulus channels are equal:

$$(A_1 - I_1) = (A_2 - I_2) = (A - I).$$

Then, substituting in Equations 3 and 4:

$$\begin{aligned} \textit{Diff. } AvR_p \text{ for 1's} &= [D_{sa} * (A - I)] - [D_{sw} * (A - I)] \\ &= (D_{sa} - D_{sw}) * (A - I); \end{aligned}$$

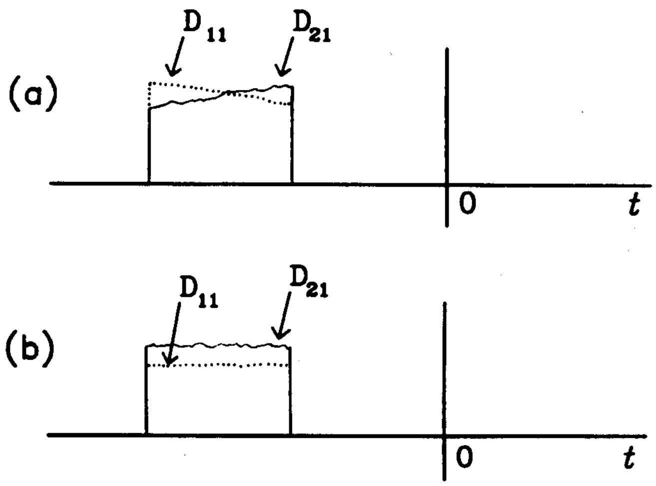

Figure 8. Examples of unequal previous-event distributions. D_{11} is the distribution of 1's preceded by 1's, D_{21} the distribution of 1's preceded by 2's. (a) Skewing across the ISI range; (b) unequal total numbers.

$$\begin{aligned} \textit{Diff. } AvR_p \text{ for 2's} &= [D_{sa} * (A - I)] - [D_{sw} * (A - I)] \\ &= (D_{sa} - D_{sw}) * (A - I) \\ &= (\textit{Diff. } AvR_p \text{ for 1's}). \end{aligned}$$

Thus, even though the attention effects are identical for the two stimulus channels, differential averaged previous response can still result. Furthermore, the resulting differential overlap would be the *same* for the two stimulus channels. There is no clue of opposite polarity to possibly alert the experimenter that differential overlap may be occurring and contributing to whatever actual attention effects may be present. Indeed, obtaining a similar attention effect for the two channels in an intramodal experiment is a result that normally would tend to give an experimenter more confidence in the reliability of the effect.

A relatively simple way to check for this problem is to make sure that the stimulus sequences contain approximately equal numbers of sames and switches. If so, this particular artifact may be ruled out. If they are not equal, then such an artifact might have occurred, and further investigation may be warranted along the lines described later in this paper.

Although this section has focused on differential overlap between attended and inattended ERPs in selective attention experiments, differential overlap is a potential problem whenever one wants to compare ERPs (or any physiological response measures) that were collected at stimulus rates fast enough to result in overlapping responses. Unequal averaged overlap on the ERP responses being compared can masquerade as significant differences between the responses and/or seriously distort true differences that might exist.

Further Problems Due to Response Overlap

Comparisons Between Difference Waves and Individual ERPs

Even if a combination of stimulus randomization and ISI jitter has effectively eliminated *differential* overlap between two ERP averages that are to be compared, there still could be substantial, albeit approximately equal, overlap in each of these individual averages. In such a case, the difference wave will have this overlap subtracted out, whereas the individual averages will not. Accordingly, any conclusions concerning the relationship between the waveform structure, distribution, or componentry of the difference wave and that of the individual ERPs could be erroneous.

Sequential Analysis

When ERPs are averaged such that adjacent stimuli are *not* randomized, the problem of differential overlap at short ISIs is even greater, placing constraints upon what phenomena may be investigated with ERPs. One of the most serious constraints is the difficulty in analyzing potentially important sequence effects—that is, the physiological and psychological effects of the previous stimulus type on the ERP to the current stimulus. The problem arises because a sequential analysis requires dividing the ERP "Full Averages" into subaverages based on different subsets of the previous stimulus types and/or ISI subranges. This procedure, which clearly violates the assumption of randomized adjacent stimuli, results in the physiological effects of the previous stimulus type (or ISI) on the current ERP being confounded with differential overlap from the ERPs to those differing previous stimuli. Further, because each of the current ERP subaverages is preceded by only one particular class of previous response, perhaps occurring within a narrowed ISI range, there is generally more residual distortion from previous responses (less "smearing out") in each of the subaverages than in the Full Average.

There are two main variables of interest in a sequential analysis: previous ISI subrange and previous stimulus (or response) type. To illustrate this analysis, a selective attention example will again be employed, but sequential effects also play an important role in other classes of psychophysiological phenomena (e.g., memory formation and decay, language processing).

ISI subrange. Consider some relatively undistorted averages from an auditory selective attention experiment. Figure 9a shows the ERP averages for attended and inattended left-ear tones. Each of these two averages are Full Averages because the ERP responses that were included in them were preceded by all possible previous stimulus types across the entire ISI range of 120–320 ms. The only criterion for including a response in one of the averages or the other was whether it was attended or not.

Suppose, however, one wanted to examine the effects on these ERPs of having the previous stimulus occur 120–220 ms versus 220–320 ms before. The data can be reaveraged accordingly (Figures 9b and c). From the observation that the prestimulus baselines are not flat and are quite different between the two pairs of subaverages, it is apparent that there is substantial previous-ERP overlap. For example, the negative wave peaking at about −70 ms in Figure 9b is the partially "smeared-out" N1 of the preceding ERPs, whereas the positive wave peaking at about −70 ms in Figure 9c results from the P2s of the preceding ERPs. This distortion certainly continues past the prestimulus time into the current waveforms. Therefore, any true effect of these ISI differences on the current ERP is confounded with the differential previous-ERP overlap.[5]

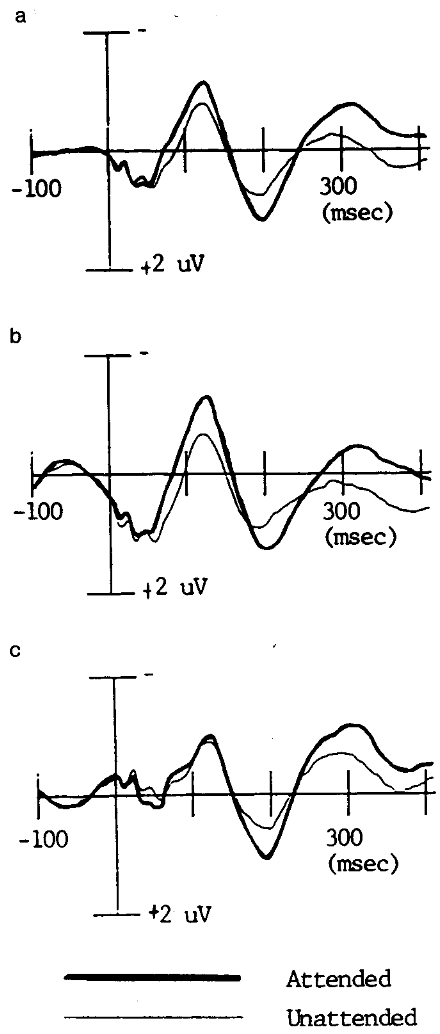

Figure 9. Data example of previous-response overlap when attempting sequential analysis based on previous ISI subrange. ERPs are to left standard (nontarget) tones at the C3 site from a short-ISI selective attention experiment, grand averaged across subjects. The attended and inattended Full Averages (a) consist of ERP responses preceded by all possible previous response types elicited across the entire possible ISI range of 120–320 ms. These trials were sorted into subaverages (b, c) based on whether the previous stimulus came within the ISI subrange of 120–220 ms (b) or 220–320 ms (c). Note the substantial overlap from previous responses distorting the subaverages in (b) and (c) and that this overlap is quite different between (b) and (c).

One comparison that could validly be made would be to derive the attentional difference wave formed by subtracting the inattended from the attended ERP in Figure 9b and compare it with the analogous difference wave from Figure 9c, thereby comparing the ERP attention effects for the two previous ISI subranges. However, this comparison clearly illustrates the caveat mentioned above that even though the attentional difference waves for these two ISI subranges could be compared with each other, they could not be compared with the individual ERP subaverages from which they were derived. This is because the individual subaverages have substantial previous-response overlap distorting their waveform structure, whereas the difference waves would (presumably) have this overlap subtracted out.

[5]Note that the presence of substantial overlap in the individual subaverages, but not in the Full Averages, underscores the relationship between the width of the ISI jitter and its effectiveness as a low-pass filter.

Previous stimulus (or response) type. Use of ERPs to investigate the physiological and psychological effects of the type of the previous event is also subject to severe confounds at short ISIs. For example, consider partitioning the Full Average for attended left-ear responses (L+'s) and the Full Average for inattended left-ear responses (L−'s) into two subaverages each, based on the two possible previous response types in each condition:

Full Average		Subaverages
L+	→	[1] (L+).L+
	→	[2] (R−).L+

Full Average		Subaverages
L−	→	[3] (L−).L−
	→	[4] (R+).L−

There are a number of comparisons one might want to make among such subaverages. For example, one might want to compare Subaverage 1 with Subaverage 2 to evaluate the effect of previous stimulus type on an attended left-ear ERP response. Or, one might want to compare the attention effect for left-ear stimuli when the previous stimulus was in the same ear (i.e., Subaverage 1 minus Subaverage 3) with the attention effect for left-ear stimuli when the previous stimulus was in the opposite ear (i.e., Subaverage 2 minus Subaverage 4). However, because these various subaverages all have different previous ERP response waveforms overlapping them, analyses such as these would be confounded (see Woldorff, 1989, for further discussion of these comparisons).

Procedures for Removing the Overlap: The Adjar Technique

Having examined some of the ways in which adjacent-response overlap can distort ERP waveforms in psychophysiological experiments, we may now consider an approach for removing the overlap distortion in recorded data. A set of algorithms for estimating and removing adjacent-response overlap under certain conditions is proposed. This method, termed the *Adjar* (adjacent response) technique (or filter), is based on the conceptual framework just developed. Although this technique will generally be most effective when its application is planned as part of the design phase of an experiment, there are many cases where all aspects of the procedure, including evaluating whether and how to implement it, could be performed after data collection. This method will be illustrated primarily in the context of selective attention experiments, but its principles can be applied more broadly.

There are two main levels to the Adjar technique, a noniterative level for removing previous-response overlap only (Level 1) and an iterative level for removing distortion from both previous and subsequent overlap (Level 2).

Level 1 enables the study of sequential effects (i.e., the physiological and psychological effects of previous event type and ISI) in short-ISI ERP experiments by subtracting out the differential distortion from previous-response overlap. However, the ISI jitter must have been sufficient to have yielded Full Averages (i.e., averages based on all the previous stimulus types and ISIs) that are relatively undistorted, as these are used to estimate the overlapping previous responses that need to be subtracted out. This level may also be employed to remove previous-response overlap from Full Averages themselves if they are only slightly distorted.

Level 2 is more complicated to implement than Level 1, but it is also more powerful. Basically, it allows removal of overlap from Full Average waveforms that were initially more distorted (more than just slightly) because they were obtained using shorter ISIs with insufficient jitter (e.g., see Figure 1b). This distortion, which may have resulted from overlap by both previous and subsequent responses, is removed by a series of iterations that converge toward the best estimates of the "true" Full Average waveforms. In addition, if the effects of previous event type and ISI are of interest, these Adjar-corrected Full Average waveforms can then be used to estimate the previous-response overlap distorting sequence-based subaverages.

Adjar Level 1

Filtering subaverages. Using the example of a selective attention experiment, this level would start with the standard set of ERP data consisting of the attended Full Average and the inattended Full Average (an example of these for left-ear tones in our example experiment was shown in Figure 9a). These are called "Full" Averages because the responses that went into them were preceded by all of the possible previous stimulus types at all of the possible previous ISIs over the entire 120–320-ms range. To proceed with Level 1 analyses, these Full Averages must be relatively undistorted (obtained with a sufficiently wide ISI jitter range). A useful criterion for determining whether this has been accomplished is that the prestimulus baseline should be relatively flat; including a substantial length of baseline in the recording epoch will clearly facilitate such a check. There may be some residual random EEG fluctuations in the baseline, but there should not be what looks like a time-locked signal (e.g., see Figures 1a and b); this will usually be clearest in the grand average (see Figures 9b and c). A better indication of the amount of overlap distortion of the Full Averages can be gained through application of the technique.

To illustrate Level 1, consider reaveraging these data such that the attended and the inattended Full Averages for left-ear stimuli are each subdivided into four subaverages, based on factorial combinations of previous ISI subrange (short or long) and previous stimulus type (left or right tone) (Figure 10). The grand average waveforms of these eight left-tone subaverages for the Cz site in our attention experiment example are shown in Figures 11a–d. If averaged together, the four attended left-tone subaverages of this figure would be equal to the attended left-tone Full Average for this site, and the four inattended subaverages would be equal to the inattended Full Average.[6]

One can easily infer from Figure 11 that there is substantial differential previous-response overlap superimposed upon these various subaverages from the observation that the prestimulus baselines are not flat and look quite different in each of the sub-

[6] In this experiment there were also infrequent, slightly deviant "target" tones (intensity decrements) in each ear; the task was to detect those in the attended ear from amongst the more frequent "standard" tones. For this discussion, however, only ERP averages to standard tones (preceded by only standard tones) are considered. A detailed discussion of adjacent target ERPs can be found in Woldorff (1989).

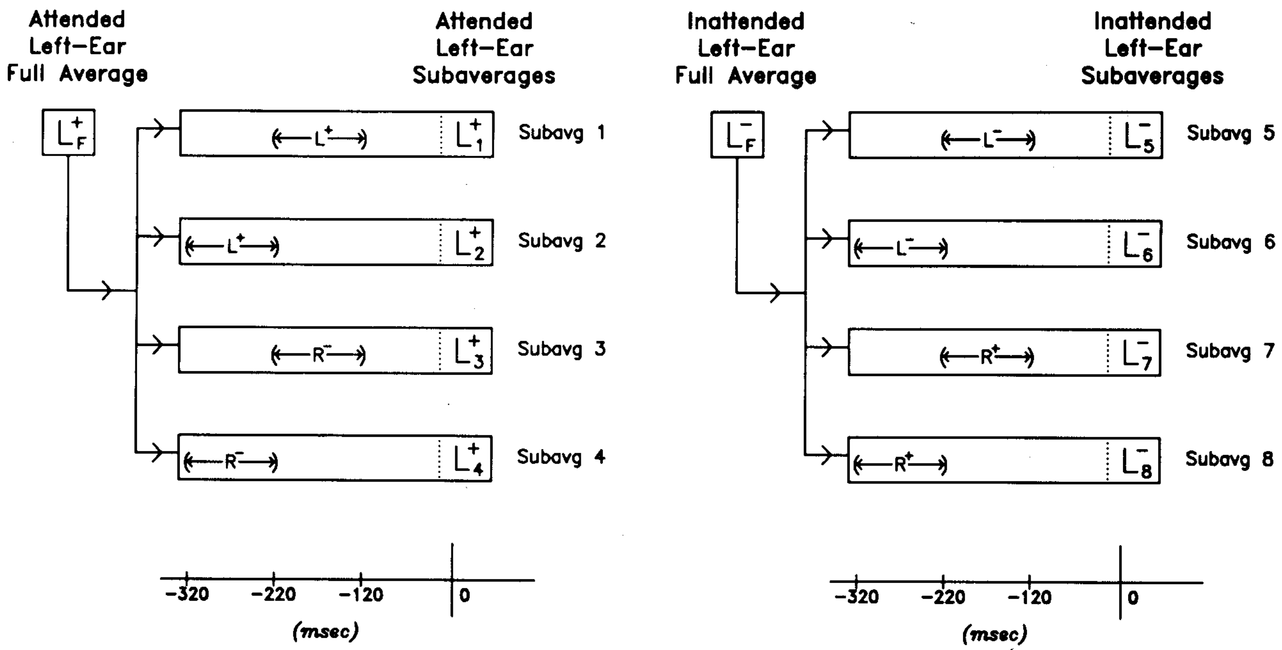

Figure 10. Schematic representation of the subdivision of each of the attended and inattended Full Averages of left-ear ERP responses into four subaverages based on previous stimulus type and ISI subrange. L's and R's represent ERPs elicited by left- and right-tone responses, respectively, which are either attended (+) or inattended (−). The subscript F stands for Full Average. The large letter on the right side of the box for each subaverage represents the current response, and the smaller preceding letter represents the previous response elicited in the indicated ISI range. For example, Subaverage 1 consists of all attended left-ear responses that were (immediately) preceded by an attended left-ear response elicited 120–220 ms previously, and so on.

averages. In all of these cases, the distortion undoubtedly continues past the prestimulus period into the current waveform itself. The goal of the Level 1 Adjar procedure is to subtract out these overlapping responses so that we can compare the true attended and inattended responses elicited by the current stimuli when preceded by these various subsets of previous events.

Consider Subaverage 1 for a particular subject. The intent is to derive an estimate of the average of the overlapping previous ERP responses, which in this case are the ERPs elicited by attended left-ear tones that occurred 120–220 ms previously. As outlined earlier and further described below, the average of these overlapping previous responses can be modeled as the convolution of the attended left-tone ERP waveform [i.e., the $R_p(t)$], with the appropriate previous event distribution [$D_p(t)$]. The first step toward removing the distortion is therefore determining this normalized event distribution. Because the digitization rate for this experiment was 2 ms, the number of attended left stimuli that were preceded by attended left stimuli in each previous 2-ms interval in the short (120–220-ms) ISI range must be ascertained. These 50 numbers (i.e., one for each 2-ms previous ISI interval) are converted to percentages to yield the normalized previous-event distribution corresponding to this subaverage. These normalized numbers are $D_p(-120)$, $D_p(-122)$, $D_p(-124)$, and so forth.

Assuming all these previous attended left-tone ERP responses were essentially identical and that a good estimate of the attended left-ear ERP waveform was available, the next step would be to convolve such a waveform with the above-described previous-event distribution. That is, one would take such a waveform and shift it over to the left 120 ms and weight it by $D_p(-120)$, shift it over 122 ms and weight it by $D_p(-122)$, and so on. For the present case, this shifting and weighting must be done 50 times, once for each of the 50 2-ms intervals. These 50 time-shifted and weighted waveforms can then be added up to yield, for this particular subject, the convolution of the previous attended left-tone response with the appropriate previous-event distribution. The later portion of this convolution waveform should then be a good estimate of the distortion present in Subaverage 1 due to the attended left-tone responses elicited 120–220 ms previously. This convolution waveform can then be subtracted from Subaverage 1 to obtain an undistorted, or filtered, Subaverage 1.

Two important issues must be addressed in this scenario, however. First, which attended left-tone ERP waveform should be convolved in the above way? Second, considering that all those previous attended left-tone ERPs were probably not all identical, why should the convolution of one attended left-tone ERP waveform provide a good estimate for their average? The answer to the first question is that, in general, the best waveform to use for each subject for this convolution is the Full Average of attended left-tone ERPs for that subject because the Full Average is the best estimate available for those preceding attended left-tone ERPs. This proposition and the answer to the second question are discussed below.

The proposition that the convolution of the Full Average provides a good estimate of the average of the corresponding previous responses is central to the Adjar technique. A distinction should be made, however, between an ideal "true" Full Average ERP, which is defined here as one completely free of both residual EEG noise and adjacent-response overlap, and the recorded Full Average ERP, which is the one actually available to the experimenter. The present Level 1 analysis will assume

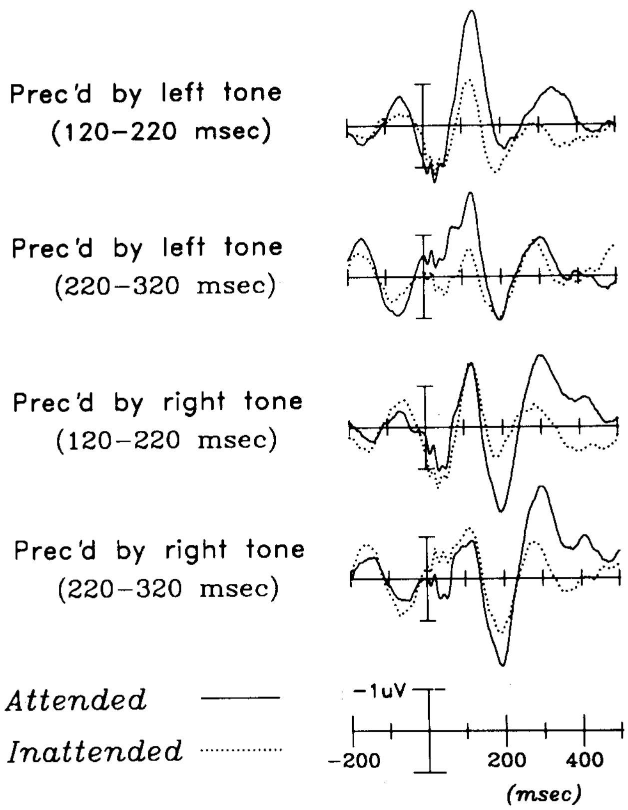

Figure 11. Grand average ERPs at the Cz site for the eight left-tone subaverages depicted in Figure 10. Note the substantial differential previous-response overlap distorting these various subaverages. Data from Woldorff and Hillyard (1991).

that the recorded Full Average ERP is a very good estimate of the true ERP; therefore this distinction can be ignored for now, allowing us to first focus on why the (true) Full Average response itself can serve as a good estimate of the previous responses. In Appendix 2, this issue is examined in more detail, and the magnitude and nature of the error involved in using the recorded Full Average response waveform are discussed.

In earlier analyses in this paper, waveform invariance generally was assumed—namely, that the ERP response to a particular type of stimulus, say an attended left-ear tone, is always approximately the same no matter when it occurs. However, in investigating sequence effects, we are now acknowledging that responses may differ as a function of characteristics of the preceding event, and, in fact, that these differences may be of considerable interest. Thus, the proposition that any one single response waveform, such as the attended left-tone Full Average, can be used to estimate the overlap contributed by all of the different previous attended left-tone responses that actually occurred might not seem obvious. The key to this proposition, however, is that the combination of responses that make up the Full Average is approximately the same as the combination that would be expected to make up the *average* of the previous responses elicited at any particular prior time point (or time bin), and it is the average of the previous responses at each of the various prior time points that actually contributes to the total averaged previous response (i.e., to that which we have been calling AvR_p).

This point may be clarified through the following example. Consider all the previous attended left-tone responses that were elicited at the particular previous time point $-t_i$ (or, more appropriately, in the $-i$th time bin), where $-t_i$ is between −120 and −220 ms. Assume there were M_i of such responses (i.e., $D_p(-t_i) = M_i/N$, where N = the total number of current trials) and that these M_i previous responses were not all identical, but could vary according to which event came *before them* and how long ago. However, these second-order previous events (i.e., the events previous to the M_i immediately previous responses) would have been randomly distributed, and hence the M_i immediately previous responses were randomly likely to be any of the various possibilities for attended left-tone responses. Therefore, the best estimate of the average of these M_i responses is the average of all these possibilities, which is just the Full Average.[7] In other words, the contribution from the M_i responses that actually occurred at $-t_i$ ms to the total averaged previous response (i.e., that which resulted from the responses from all previous time points) is approximately equal to what the contribution would have been if all of the M_i responses had actually been identical to the Full Average. Because this argument holds for the arbitrary previous time point of $-t_i$ ms, it is true in general at all the previous time points. Therefore, the convolution (i.e., the shifting-weighting-summing) of the Full Average with the entire event distribution will give a good approximation of the (total) averaged previous response, AvR_p.

Thus, it does not matter that the attended left-tone responses that actually occurred at all those various previous time points were not all of one type but could have been any of a number of different possibilities. The sum of the contributions from all of these possibilities to the (total) averaged previous response is (approximately) equal to what the contribution would have been if all of the previous responses had been equal to the Full Average response. (For further discussion of this issue, see Appendix 2.)

Once the estimate of the (total) averaged previous response for Subaverage 1 has been obtained for a particular subject, it is then simply subtracted from Subaverage 1 for that subject. The other seven left-ear subaverages would be Adjar filtered in an analogous way, using the appropriate event distribution and appropriate Full Average to estimate the previous-ERP overlap in each case.

The Level 1 procedure described above was applied to the eight left-tone subaverages for all the subjects from our example experiment. The grand averages of the estimated previous-response overlap for each of these eight subaverages are shown in the middle column of Figure 12, and the grand averages of the resultant filtered subaverages are shown in the right column. A comparison of these filtered subaverages with the unfiltered subaverages (left column) shows how much less distorted the prestimulus baselines have become and therefore how close the previous-response overlap estimates were to the actual overlap in the prestimulus period. Although such a result does not guarantee that all of the distortion from the overlapping previous responses has been successfully subtracted out, it indicates that the bulk of it has been removed. Moreover, it provides support for the argument that convolving a single response waveform

[7]Note that this conclusion would still hold when there was additional random variation of these M_i responses that was *unrelated* to what event came before them—that is, the Full Average would still be the best estimate of the average of the M_i responses.

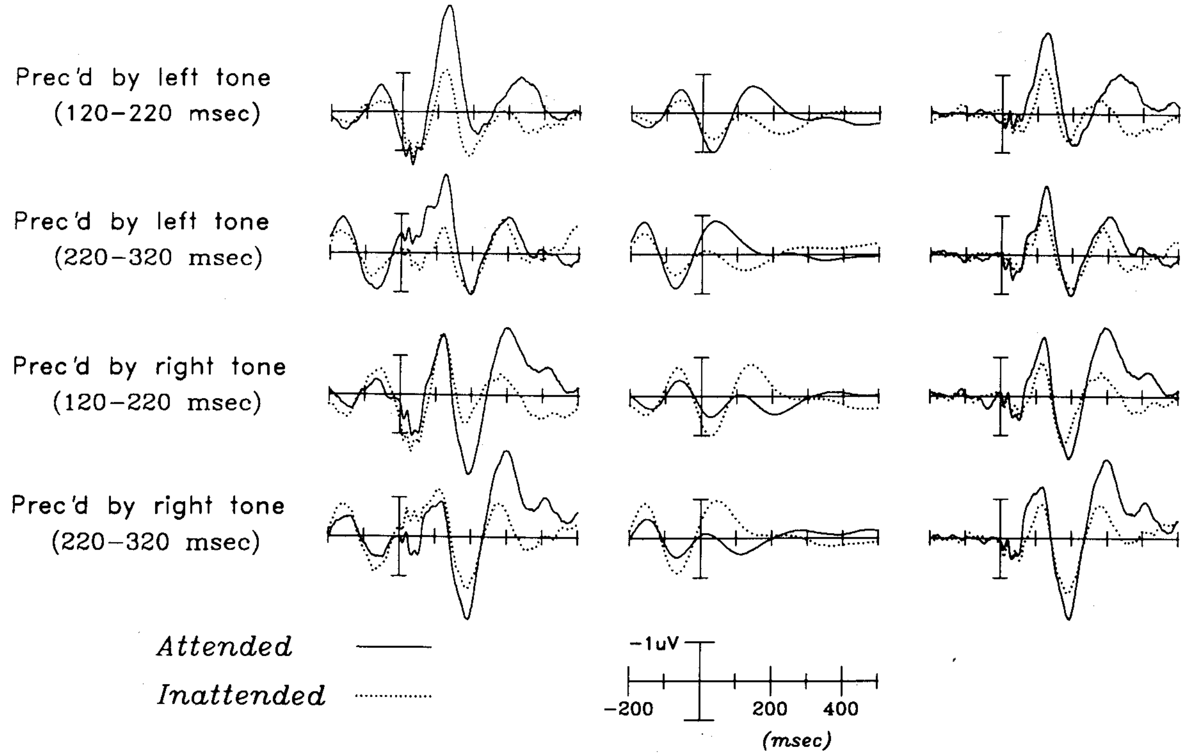

Figure 12. Adjar-filtering process applied to the eight left-tone subaverages depicted in Figure 10 and shown in unfiltered form in Figure 11. The original unfiltered subaverages are shown again (left column), juxtaposed with the estimates of the previous-response overlap (middle column) and the corresponding Adjar-filtered subaverages (right column) obtained by the subtraction of these estimates. Data from Woldorff and Hillyard (1991).

does indeed provide a good estimate of the average of the preceding responses that probably were varying. Now that this distortion has been removed, the influence on the left-tone ERPs and ERP attention effects of what stimulus came before and how long ago it occurred can now be examined. (The filtered subaverages also can be collapsed across previous ISIs to concentrate on the effects of previous stimulus type, or across previous stimulus type to concentrate on the effects of previous ISIs.)

It is important to emphasize that this filtering operation subtracts out overlapping scalp potential changes that were time-locked to the previous stimulus, leaving those that were time-locked to the current stimulus. Comparisons between the filtered subaverages will therefore reflect processing interactions that are time-locked to the current stimulus. However, such interactions must be interpreted with care. The most straightforward effects these interactions would be likely to reflect are changes in the processing of the current stimulus as a function of previous events, such as an N1 component or subcomponent of the current ERP being bigger or smaller depending on the preceding event type or ISI. These changes will certainly be time-locked to the current stimulus. However, if there are changes in the processing of the previous stimulus because the current stimulus arrived while that processing was still occurring, then these changes may well be time-locked to the current stimulus as well (i.e., to the arrival of the information necessary for the interaction), rather than to the previous stimulus. If so, these changes would remain as part of the response to the current stimulus rather then having been subtracted out during the Adjar filtering.

One might argue that this second type of change should be attributed to the processing of the previous stimulus even though it is actually time-locked to the current stimulus. However, if it is time-locked to the current stimulus, then in some sense it should also be considered a reflection of the processing of the current stimulus. It is probably most accurate to attribute such an effect to the processing interaction of the two stimuli, where the processing of the current stimulus interacts with the ongoing processing of the previous stimulus. Regardless, the removal of distorting, overlapping ERPs that are time-locked to the previous stimulus should prove to be a useful technique for understanding the true sequential interaction effects.

Filtering Full Averages. The Level 1 technique may also be employed to remove small amounts of previous-response overlap from the Full Averages themselves given that they are only slightly distorted. This technique might be particularly useful when a combination of randomization and ISI jitter has eliminated most of the overlap, but some overlap may still remain that could be distorting the small early waves of the current Full Average ERPs.

In such a case, because the recorded Full Average ERPs are only slightly distorted, they are still fairly good estimates of the true Full Average responses and therefore may be used to estimate average previous responses. Assuming subaverages are being investigated, the procedure is a simple extension of the one previously outlined. Because a Full Average is just an average of subaverages, simply averaging the filtered subaverages together will yield filtered Full Averages. In addition, the overlap estimates calculated for each of the subaverages can be

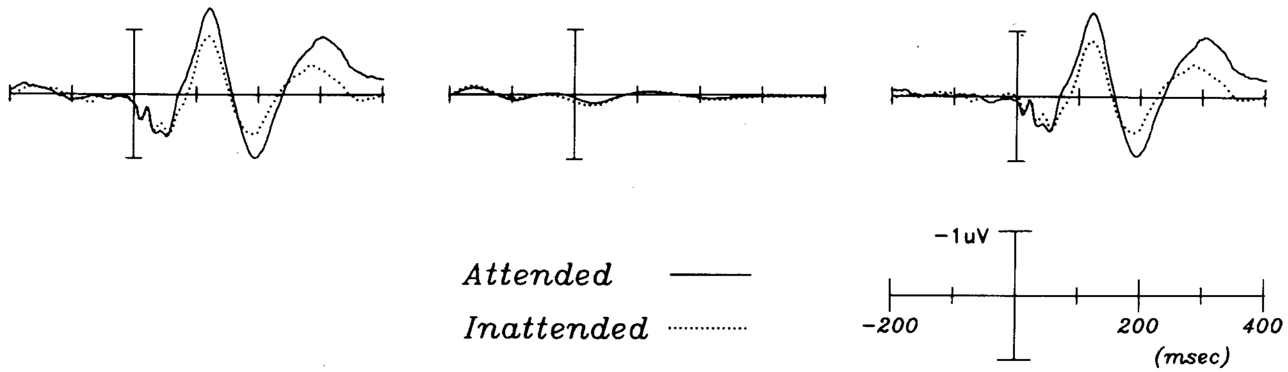

Figure 13. Adjar filtering of the ERP Full Averages for the left-ear tones at the C3 site. Left column: original unfiltered Full Averages; middle column: estimates of previous overlap; right column: Adjar-filtered Full Averages. These resulted from collapsing (averaging) together the appropriately corresponding subaverages for this site. Although there was some distortion from previous responses even in these Full Averages, it was fairly small in magnitude and was not different for the attended versus the inattended average. Data from Woldorff and Hillyard (1991).

averaged together (in an appropriately weighted way) to examine the estimated total previous-response overlap that was distorting each of the original Full Averages. This, in fact, is the most direct way to examine whether the overlap on the Full Averages actually was negligible and whether it was differential or not (e.g., different for the attended Full Average than for the inattended Full Average).

For the selective attention example, Figure 13 shows some unfiltered and filtered left-tone Full Averages and the corresponding total overlap that was estimated and removed. Note that there was indeed some overlap in each of the Full Averages, but it was not very large, and it was not differential for the attended and inattended stimuli.

In the procedure just outlined to obtain filtered Full Averages, the original unfiltered Full Average waveforms are used to estimate the overlap on *themselves* so that it can be removed, although this was accomplished by filtering subaverages and then collapsing them together. If one were interested only in Full Averages, the steps involving subaverages could be omitted. This would then be equivalent to calculating each of the convolution terms on the right side of an equation such as Equation 1 and summing them (or, depending on normalization considerations, averaging them together in an appropriately weighted way).

Adjar Level 2

It has been argued above (also see Appendix 1) that in cases where the Full Averages are at most only slightly distorted, using them as the waveforms to be convolved in the application of Level 1 will still be effective in removing previous-response overlap. If Level 1 has already been applied, a straightforward way to gauge the relative degree of distortion of the Full Averages is to examine the estimated total previous-response overlap distorting them, as outlined earlier. If this distortion looks small relative to the Full Averages themselves, then the assumption that they were only slightly distorted is probably reasonable, and the application of Level 1 alone adequate.

In cases where the Full Averages are more than just slightly distorted, such a single-step process might not provide adequate removal of overlap. It is argued in Appendix 2, however, that even in such cases, assuming the ISI was jittered, applying the algorithms of Level 1 will generally provide more gain than loss in removing systematic distortion due to previous-response overlap, while introducing only negligible additional random error. This suggests that application of these algorithms in an appropriate iterative way should remove more and more of any remaining distortion with every iteration. This is the approach of Level 2.

The Level 2 procedures are more complex in implementation than those of Level 1 but have correspondingly greater analytic capability. Level 2 allows removal of overlap distortion, due to both previous and subsequent responses, from Full Average waveforms that are distorted because they were obtained under conditions of insufficient ISI jitter. As a final step, Level 2 can also include the procedures of Level 1 for analyzing the effects of previous event type and ISI, in that once relatively undistorted Full Averages have been obtained, they can be used to estimate the previous-response overlap distorting sequence-based subaverages.

Level 2 involves a series of iterations that result in a convergence toward the true Full Average waveform. There are various ways that the estimation and convergence process could be implemented, depending on the specifics of the data set, but they would generally all employ a similar approach. One sequence of steps is outlined below:

1. Start by assuming that the original Full Average waveforms are the best estimates available of the overlapping subsequent responses.
2. Convolve these subsequent-response estimates with the appropriate subsequent-event distributions, obtaining a first estimate of the overlapping averaged subsequent response, AvR_s, for each current Full Average.
3. Subtract these estimates from each original Full Average, thereby obtaining better estimates of these current Full Average responses, especially their later portions.
4. Now take these better estimates of the current responses to be the best estimates of the previous responses.
5. Convolve these previous-response estimates with the appropriate previous-event distributions, obtaining an estimate of the overlapping averaged previous response, AvR_p, for each current Full Average.

6. Subtract out these averaged previous-response estimates from the *original* Full Averages (not those obtained in Step 3), thereby obtaining even better estimates of the current responses, especially their early portions.
7. Now iterate. That is, use the better estimates (those obtained in Step 6) of the current responses as better estimates of the subsequent responses. Return to Step 1 and repeat process. Continue iterations until changes between the waveform estimates obtained in successive iterations are negligible.
8. The final outcome of these iterations will be estimates of both the averaged previous-response overlap and the averaged subsequent-response overlap for each original Full Average waveform, along with two sets of Full Average waveforms: (A) one set with the final previous-response estimates subtracted out (but not the subsequent) and (B) one set with the final subsequent-response estimates subtracted out (but not the previous). As a final step, subtract the final subsequent-response estimates from each corresponding Full Average in set A, thereby deriving Full Averages with both previous- and subsequent-response estimates subtracted out. (Or, equivalently, subtract both final previous- and final subsequent-response estimates from the original Full Averages.)[8]

An application of these procedures to simulated data is shown in Figure 14. A series of double calibration pulses (square waves) were generated at short ISIs (150–350 ms). The result of conventional ERP-like averaging of these (shown in the thin trace) clearly has distortion from overlap from both the previous and subsequent double-cal "responses." By applying the methodology summarized in the steps listed above, the more veridical waveform (shown in the thick trace) was obtained.

Applications of the Level 2 procedures to waveforms that are more ERP-like are presented in Figures 15 and 16. For Figure 15, an auditory ERP waveform obtained in our laboratory was taken as the original "true" undistorted Full Average waveform for a simulated short-ISI experiment. A computer then generated the overlap-distorted average that would be obtained if such an ERP waveform were repeatedly elicited at ISIs of 180–300 ms. It was then assumed that the overlap-distorted waveform was what the experimenter would have recorded and was therefore what would be available for application of the Level 2 Adjar procedures. Figure 15d shows the final waveform estimate obtained with the procedures after three full iterations (i.e., including carrying out final Step 8), superposed on the true original waveform for comparison.

Figure 16 shows a more complicated simulation with two types of ERP waveforms. An auditory ERP and a visual ERP were taken as the original "true" undistorted responses, and the

[8]It may seem that a more straightforward way to carry out the Level 2 iterations is to calculate and subtract the subsequent- and previous-overlap estimations in parallel at the same iteration, rather than calculating and subtracting them in separate sequential iterative steps. Although the sequential approach is more complicated, it both converges more quickly and removes second- and higher-order overlap at the same time the first-order is removed (see Woldorff, 1989, Appendix 2B). More importantly, however, the parallel approach can run into a nettlesome but subtle convergence problem—namely, that the solutions can begin to converge in the first several iterations, but then, under certain conditions, start to diverge as the iterations continue. This divergence is due to a selection and accumulation of power at the frequency whose period is twice the average ISI. This approach is therefore not recommended.

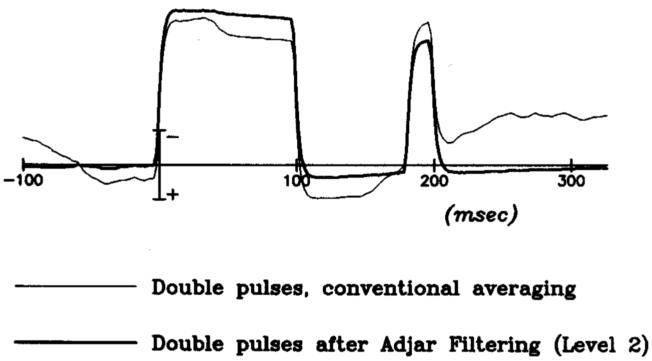

Figure 14. Example of the application of Level 2 Adjar technique to simulated data. Light trace is the original average of 200 double-pulse "responses," where the first pulse of the pair began at time 0 and lasted 100 ms and the second began at 180 ms (relative to the trigger) and lasted 20 ms. ISIs of the triggers ranged from 150 to 350 ms. Note the distortion from overlap from both the previous and subsequent double-pulse responses. Bold line is after applying Level 2 Adjar filtering.

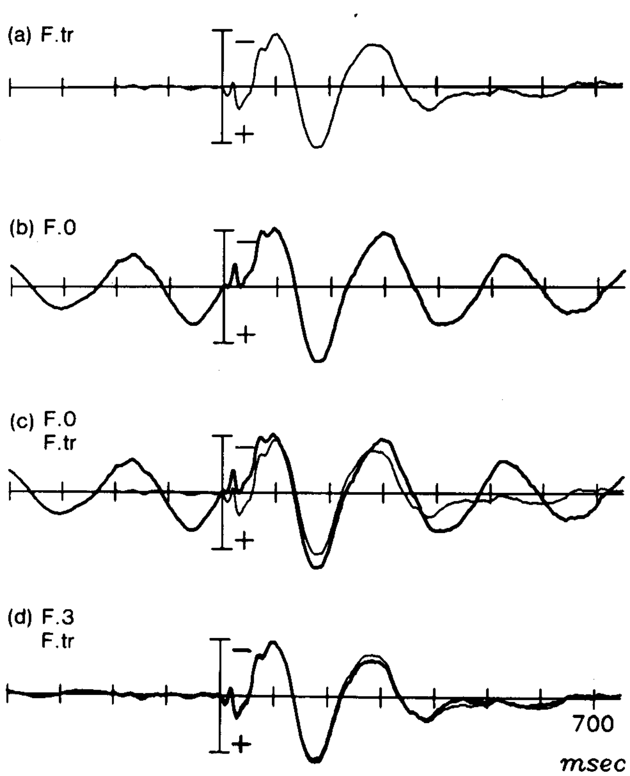

Figure 15. Simulation of Level 2 procedures with an auditory ERP. (a) Original auditory ERP waveform, used here as the original "true" undistorted Full Average waveform in a simulated short-ISI experiment. (b) Computer-generated overlap-distorted average that would be obtained if such an ERP waveform were repeatedly elicited in an experiment at ISIs of 180–300 ms (includes first-, second-, and third-order overlap, both previous and subsequent). (c) Overlap-distorted waveform (bold trace) and original undistorted waveform (light trace), superposed for comparison. The distorted waveform was then taken as the ERP average that would have been recorded and therefore as what would be available to the experimenter for application of the Level 2 Adjar procedures. (d) The final waveform estimate (bold trace) obtained after three full iterations (i.e., including carrying out Step 8), superposed on the original undistorted waveform (light trace) for comparison.

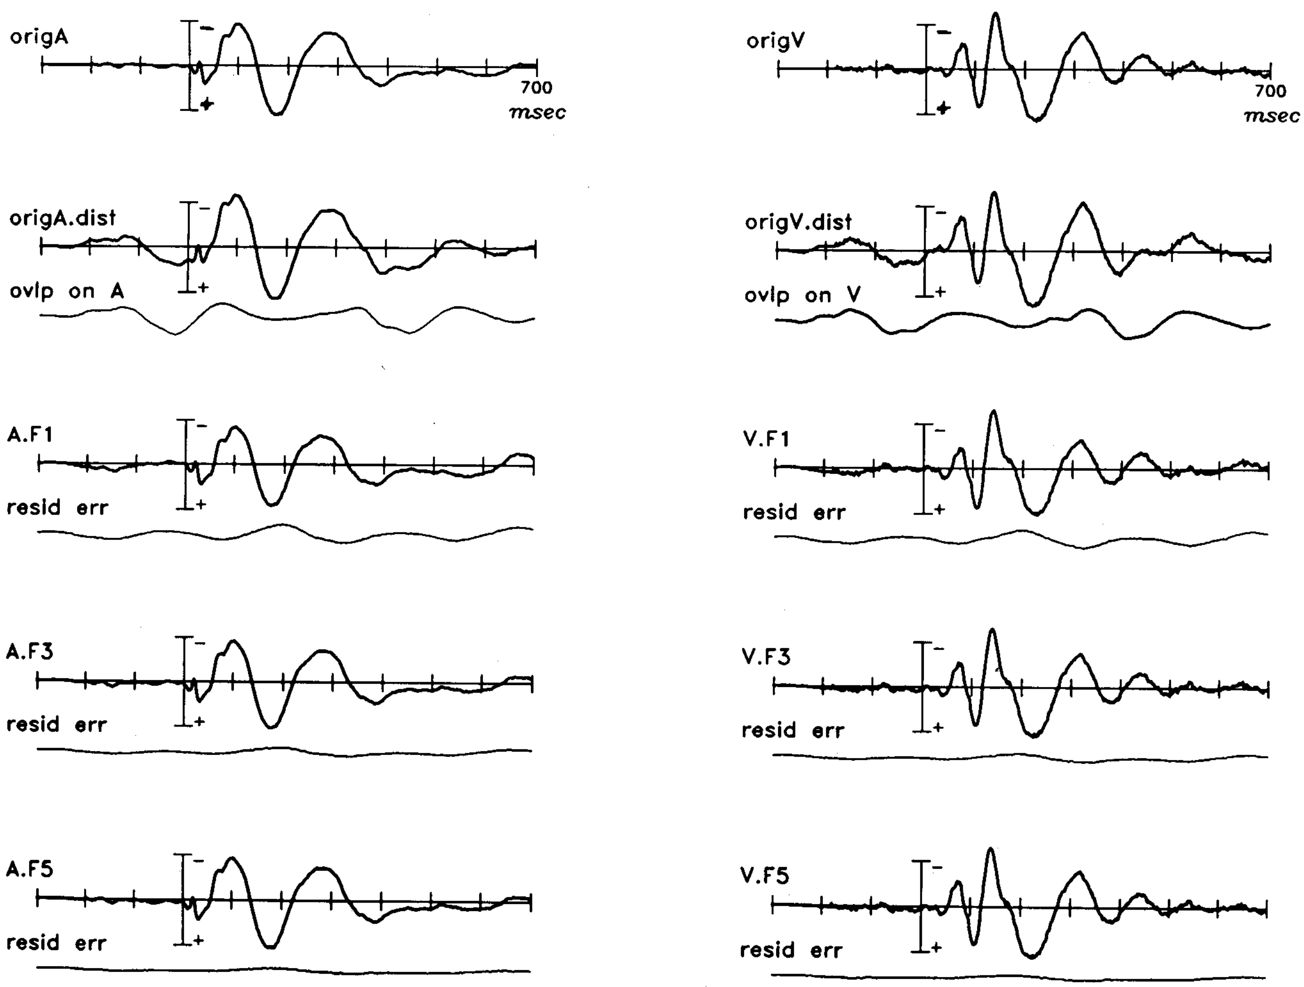

Figure 16. A simulation of Level 2 procedures with two types of ERP waveforms. Top row: an auditory ERP (left) and a visual ERP (right), which were used here as the original "true" undistorted Full Average waveforms. Row 2: the computer-generated, overlap-distorted averages that would be obtained if these waveforms were elicited in a cross-modal experiment at ISIs of 170–300 ms; the corresponding previous and subsequent overlap (first-order only) distorting these waveforms are shown immediately below. The distorted waveform averages were then taken as the recorded ERP averages that would be the starting point for application of the Level 2 Adjar iterations. Rows 3–5: the final waveform estimates that would be obtained (if Step 8 were completed each time), after the first, third, and fifth iterations, along with the corresponding residual error waveforms (resid err) derived by subtracting the original "true" waveforms (top row) from these waveform estimates. As the iterations proceed, one can also see the convergence process proceed, as the waveform estimates more and more closely approximate the original undistorted ones, and the residual error approaches a flat line.

computer generated the overlap-distorted averages that would be obtained if these waveforms were elicited in a cross-modal experiment at ISIs of 170–300 ms. The resultant distorted waveform averages were then assumed to be the recorded ERP averages that would be available to the experimenter for the Level 2 iterations. The waveform estimates that would be obtained at the first, third, and fifth iterations, along with the associated residual error, are shown. As the iterations proceed, so does the convergence process, as the waveform estimates more and more closely approximate the original undistorted ones and the residual error approaches a flat line.

A brief simplified analysis of the residual error in the Level 2 procedures can be found in Appendix 3. An important result to note from that analysis is that the convergence is frequency dependent, with the lowest frequencies in the residual overlap error taking the most iterations to remove. The analysis also suggests that the very lowest frequencies (less than, say, 0.25 of $1/T_{jw}$, where T_{jw} = the jitter width) would converge so slowly that attenuating them ahead of time would generally be advised. (In most cases, any such high-pass prefiltering should employ causal filters.)

Conclusions

Many potential pitfalls can result from adjacent-response overlap when using ERPs to probe the mechanisms involved in the processing of rapidly presented stimuli. The purpose of the development of the Adjar framework and technique was to enable a way to understand these problems and deal with them. In some cases, an effective approach could consist of careful consideration of these potential problems during the experimental design phase, with the intent of judiciously setting certain parameters (e.g., the ISI jitter) or making some other adjustment in the design. In other cases, the approach might entail greater caution in the interpretation of certain results, coupled with close scrutiny of the prestimulus baselines of the ERP averages, especially at certain sites. For a number of cases, however, precautions such as these might be inadequate in mitigating the potential overlap problems, and the application of the Adjar technique (or some other technique) for the actual removal of the residual overlap may be required.

The use of ERPs as probes in short-ISI experiments offers great potential for understanding the dynamic mechanisms of human information processing. Besides the selective attention examples discussed in detail in this paper, potential applications include language processing (e.g., priming effects), memory formation and decay, and perceptual/neurophysiological habituation. When pursuing such research, the possible problems incurred by response overlap must be addressed. The principles and techniques presented in this paper are intended to help in this regard, so that these avenues of research can be more effectively pursued.

REFERENCES

Aunon, J. I., McGillem, C. D., & Childers, D. G. (1981). Signal processing in evoked potential research: Averaging and modeling. *CRC Critical Reviews in Bioengineering*, *5*, 323–365.

Brillinger, D. R. (1981). The general linear model in the design and analysis of evoked response experiments. *Journal of Theoretical Neurobiology*, *1*, 105–119.

Eysholdt, U., & Schreiner, C. (1982). Maximum length sequences – A fast method for measuring brainstem evoked responses. *Audiology*, *21*, 242–250.

Galambos, R., Makeig, S., & Talmachoff, P. J. (1981). A 40-Hz auditory potential recorded from the human scalp. *Proceedings of the National Academy of Sciences of the United States of America*, *78*, 2643–2647.

Hansen, J. C. (1983). Separation of overlapping waveforms having known temporal distributions. *Journal of Neuroscience Methods*, *9*, 127–139.

Hansen, J. C., & Hillyard, S. A. (1984). Effects of stimulation rate and attribute cueing on event-related potentials during selective auditory attention. *Psychophysiology*, *21*, 394–405.

Hillyard, S. A., Hink, R. F., Schwent, V. L., & Picton, T. W. (1973). Electrical signs of selective attention in the human brain. *Science*, *182*, 177–180.

Mangun, G. R., & Hillyard, S. A. (1987). The spatial allocation of visual attention as indexed by event-related brain potentials. *Human Factors*, *29*, 195–212.

Mangun, G. R., & Hillyard, S. A. (1988). Spatial gradients of visual attention: Behavioral and electrophysiological evidence. *Electroencephalography and Clinical Neurophysiology*, *70*, 417–428.

Näätänen, R. (1967). Selective attention and evoked potentials. *Annals of the Finnish Academy of Science*, *151*, 1–226.

Näätänen, R. (1975). Selective attention and evoked potentials in humans – A critical review. *Biological Psychology*, *2*, 237–307.

Oppenheim, A. V., & Schafer, R. W. (1975). *Digital signal processing*. Englewood Cliffs, NJ: Prentice-Hall.

Parasuraman, R. (1978). Auditory evoked potentials and divided attention. *Psychophysiology*, *15*, 460–465.

Regan, D. (1972). *Evoked potentials in psychology, sensory physiology and clinical medicine*. London: Chapman and Hall.

Regan, D. (1982). Comparison of transient and steady-state methods. *Annals of the New York Academy of Sciences*, *388*, 45–71.

Schwent, V. L., Hillyard, S. A., & Galambos, R. (1976). Selective attention and the auditory vertex potential. I. Effects of stimulus delivery rate. *Electroencephalography and Clinical Neurophysiology*, *36*, 191–200.

Woldorff, M. G. (1989). Auditory selective attention in humans: Analysis of mechanisms using event-related brain potentials (Doctoral dissertation, University of California–San Diego, 1989). *Dissertation Abstracts International*, *50*, 5490B.

Woldorff, M., Hansen, J. C., & Hillyard, S. A. (1987). Evidence for effects of selective attention in the midlatency range of the human auditory event-related potential. In R. Johnson, Jr., J. W. Rohrbaugh, & R. Parasuraman (Eds.), *Current trends in event-related potential research. Electroencephalography and Clinical Neurophysiology* (Suppl. 40, pp. 146–154). Amsterdam: Elsevier.

Woldorff, M. G., & Hillyard, S. A. (1991). Modulation of early auditory processing during selective listening to rapidly presented tones. *Electroencephalography and Clinical Neurophysiology*, *79*, 170–191.

(Received April 30, 1990; Accepted October 21, 1991)

APPENDIX 1: WAVEFORM BOUNDARIES

The model of the averaged adjacent-response overlap in terms of convolutions of waveforms with adjacent event distributions implicitly assumes an infinite time axis. The Adjar technique estimates these convolutions by shifting waveforms either to the left or right along the time axis, weighting and summing them, and then subtracting the result from current ERP averages. Because in practice the epochs of these waveforms are finite in length, it is important to consider what happens at their boundaries when they are shifted and subtracted.

Shifting waveforms to the right to estimate subsequent responses poses less of a problem in this regard; in fact, useful information for the estimation process can be inferred and should be used. In particular, because the true subsequent response could not have begun until its own time zero, any activity before that point could not have belonged to it. Therefore, before the Full Average being used to estimate subsequent responses is shifted to the right, all activity in the prestimulus baseline should be removed (i.e., zeroed out). In addition, as the iterations proceed, this estimate of zero must be both accurate and stable, which again underscores the value of having long prestimulus baselines.

Shifting waveforms to the left to estimate previous response activity, however, has other considerations. In particular, the right end of a left-shifted waveform will be missing data points to subtract from the current epoch (considering that the waveform being shifted is not likely to have ended exactly on zero). Probably the best way to handle this discontinuity is to (a) have a fairly long epoch and (b) before shifting, taper the right end of the waveform down toward baseline – that is, down to the zero determined by the prestimulus baseline – by applying some sort of tapering or windowing function (e.g., a simple linear return to baseline over several hundred milliseconds).

APPENDIX 2: ADJAR LEVEL 1 – SOME ERROR CONSIDERATIONS

It was reasoned in the text that the Full Average response is generally the best waveform to use in the calculation of the averaged previous response, for both subaverages and slightly distorted Full Averages. The reasoning was that, despite variation in the individual previous responses that actually constitute the averaged previous response, their proportions are approximately the same as those constituting the Full Average. Strictly speaking, however, the Full Average in this reasoning is an idealized "true" Full Average, F_{tr}, one completely free of residual EEG noise and adjacent-response overlap. The best estimate available of F_{tr} is the recorded Full Average ERP, F_{rc}, which is therefore the best waveform actually available to use in the convolution calculation. However, the question remains as to just how well the summed average of all the individual previous responses actually would be estimated by the convolution of either of these two waveforms—that is, how much error is involved in this estimation?

First, consider how well the convolution of F_{tr} itself, if it were available, would approximate the total averaged previous response. In the example in which M_i attended left-ear responses occurred at latencies of $-t_i$ ms, it is the average of these M_i responses, after being weighted by M_i/N during the averaging process ($N =$ the total number of current trials), that actually ends up contributing to the total averaged previous response for all the previous time points. F_{tr} best approximates the average of these M_i responses because F_{tr} is the mean of all the attended left-ear responses and is therefore the "expected value" waveform of the attended left-ear response. However, the average of the M_i responses is unlikely to be *exactly* equal to F_{tr}. Therefore using even F_{tr} would introduce some error.

There are two main ways that the actual previous responses can vary around F_{tr} (i.e., around the mean waveform). One very straightforward way is simple random variation of the responses, unrelated to the temporal relationship with other stimuli. The other main way that these responses can vary is specifically as a result of what stimuli occurred *before them* and how long ago. However, as discussed earlier, assuming fairly well-randomized stimulus presentation, the M_i previous responses are still randomly likely to have been any of the different possibilities because the stimuli that occurred before them were randomly distributed. Thus, for both of these ways for the previous ERP responses to vary, the deviation of any single previous response from the true Full Average is basically random.

Therefore, to a first approximation, we can consider a model where each of the N previous attended left-tone responses consists of the *mean* of the attended left-tone responses in that experiment—namely, F_{tr}—plus a variation waveform, $V_n(t)$, that varies randomly from one previous attended left tone to the next. (These actual previous ERP responses by definition would be free of any distortion from background EEG or overlap.) Thus, every time F_{tr} is used to estimate one of these previous responses, it is likely to be incorrect by the average magnitude of the $V_n(t)$. When F_{tr} is used to estimate the average of the M_i responses that occurred at $-t_i$ ms, the expected amount of error would be approximately equal to the expected average amplitude of averaging together M_i independent $V_n(t)$ waveforms. In a manner similar to the reduction of background EEG noise with averaging, the expected average amplitude of this error will be reduced by a factor of the square root of M_i. Similarly, if there were a total of N previous responses across all the previous time points, the overall error in using F_{tr} in the calculation of the total average previous response would be equivalent to averaging together N of the randomly varying $V_n(t)$ waveforms. Thus, this final error will be random in nature, and its expected average magnitude, to a first approximation, will be equal to the average variation of the various previous-response waveforms around F_{tr}, divided by the square root of the number of previous (and therefore current) responses.

How does this compare with other sources of random error? The background EEG noise, which is also generally modeled as random in nature, would almost invariably be substantially larger than the average variation in the ERP response waveforms, especially at fast stimulus rates. Even if, for example, the average variation in the individual ERPs were as large as F_{tr} itself, the average background EEG amplitude and its variation is likely to be 10–30 times as great. Because the average amplitude of the EEG noise will have decreased by the same factor of the $\sqrt{N}$ during the averaging process, the contribution to the total random error from using F_{tr} is likely to be relatively negligible. Thus, this source of error will be ignored in the rest of the discussion.

A second, potentially more significant source of error arises when considering that practical application of the Adjar Level 1 technique requires recorded Full Averages that are less-than-perfect estimates of the F_{tr} responses. The magnitude and nature of such errors can be examined by first expressing the recorded Full Average, $F_{rc}(t)$, as the sum of $F_{tr}(t)$ and an error waveform, $E(t)$:

$$F_{rc}(t) = F_{tr}(t) + E(t).$$

The $E(t)$ waveform itself is actually made up of two parts, corresponding to the two main sources of error in a recorded Full Average waveform: (a) residual "unaveraged-out" EEG and (b) residual overlap from adjacent responses. Although the nature of these two error contributions is fundamentally different, the first generally being random across subjects and the second systematic, both will be included in the steps below. Thus,

$$F_{rc}(t) = F_{tr}(t) + E_{ra}(t) + E_{sy}(t),$$

where

$$E_{ra}(t) = \text{the random distortion from EEG noise}$$

and

$$E_{sy}(t) = \text{the systematic distortion from overlap.}$$

Now consider what happens when $F_{rc}(t)$ is convolved with the previous-event distribution, $D_p(t)$, to estimate the averaged previous response that is overlapping and distorting a particular current ERP:

$$\begin{aligned}(AvR_p)_{calc}(t) &= D_p(t) * F_{rc}(t)\\ &= D_p(t) * [F_{tr}(t) + E_{ra}(t) + E_{sy}(t)]\end{aligned}$$

or, more simply,

$$= D_p * (F_{tr} + E_{ra} + E_{sy}).$$

Using the distributive property of convolution, this can be rewritten as

$$(AvR_p)_{calc} = (D_p * F_{tr}) + (D_p * E_{ra}) + (D_p * E_{sy}).$$

Because F_{tr} is the "ideal" waveform to have convolved, the first term will be considered the true overlapping averaged previous response, $(AvR_p)_{tr}$, that should be subtracted from the original current ERP (ignoring the relatively negligible error discussed above that would result from using even F_{tr}). So,

$$(AvR_p)_{calc} = (AvR_p)_{tr} + (D_p * E_{ra}) + (D_p * E_{sy}).$$

Thus, subtracting $(AvR_p)_{calc}$ from the original ERP would be incorrect by an amount equal to $(D_p * E_{ra})$ plus $(D_p * E_{sy})$. Because E_{ra} is random across subjects, $(D_p * E_{ra})$ will be also and therefore will be of far less concern than any systematic distortion. In addition, because E_{ra} is the residual EEG noise in the Full Averages, which will usually have relatively many trials, it is likely to be of relatively low amplitude. Furthermore, the low-pass filtering resulting from the convolution with D_p will render it smaller still. Thus, although some additional random unaveraged-out EEG noise may be added onto the current ERPs during application of Level 1, it would generally be negligible and would be far outweighed by the ability to reduce systematic error due to overlap.

Therefore, the effectiveness of Level 1 in removing systematic error is mainly a function of the relative magnitudes of $(D_p * E_{sy})$ and $(AvR_p)_{calc}$. Because these functions were derived from convolutions of E_{sy} and F_{rc}, respectively, with the same event distribution D_p, the effectiveness of the overlap removal will also be closely related to the relative magnitudes of these two functions. It is therefore easy to see why F_{rc} does not have to be exactly correct to be useful. In particular, if F_{rc} is only slightly distorted by adjacent-response overlap, then, by definition, $|E_{sy}|$ must be small relative to F_{rc}. The convolution of E_{sy} with D_p will be smaller still (because of the low-pass filtering effects) and will generally be negligible, especially relative to $(AvR_p)_{calc} = D_p * F_{rc}$. Thus, even if F_{rc} is slightly distorted by overlap, using it for estimating the averaged previous response in the Level 1 procedures will still generally be effective.

However, if F_{rc} is more than just slightly distorted by overlap, then by definition E_{sy} is not small relative to F_{rc}. Therefore, $(D_p * E_{sy})$ will probably not be as negligible, and the method will be proportionately less accurate. However, it is important to note that in almost all cases (in which the ISIs have been jittered), E_{sy} will still be considerably smaller than F_{tr} (and therefore than F_{rc}). This is because E_{sy} consists of the average of jittered adjacent responses that are "smeared-out" and attenuated (or, equivalently, low-pass filtered), whereas F_{tr} is the average of the current responses that are time-locked to the averaging epoch and therefore not attenuated. Accordingly, $(D_p * F_{tr}) = (AvR_p)_{tr}$ will generally be larger than $(D_p * E_{sy})$, and subtracting the sum of these [by subtracting $(AvR_p)_{calc}$] will therefore provide more gain than loss in removing previous-response overlap distortion from F_{rc}. However, in cases where this process does not provide adequate removal of the previous-response overlap, from either the Full Averages or the subaverages, an iterative approach, such as that of Level 2, may be necessary.

APPENDIX 3: ADJAR LEVEL 2—SOME CONVERGENCE CONSIDERATIONS

A full analysis of the various approaches and convergence criteria for Level 2 is beyond the intended scope of this article. However, some insight can be gained into the convergence process by carrying out a simplified analysis of the propagation of the residual overlap distortion error during a series of Level 2 iterations.

Consider an ERP experiment with essentially one type of stimulus and response. Assume that because of the ISIs used, the waveforms elicited, and the amplifier bandpass, essentially all of the overlap distortion results from the ERPs elicited by the first-order (i.e., immediately) adjacent stimuli in the sequence. That is, assume that the second- and higher order adjacent responses are essentially too remote in time (as well as too highly jittered, perhaps) to contribute much overlap. Also assume that any effects of response variation are relatively negligible, as described in Appendix 2. (See Woldorff [1989, Appendix 2c], however, for some additional discussion on this point.)

First, define the following terms:

D_p = previous-event distribution

D_s = subsequent-event distribution

F_t = "true" Full Average

F_0 = original distorted Full Average ($=F_{rc}$)

ovp_t = the true (actual) previous-response overlap

ovs_t = the true (actual) subsequent-response overlap.

By definition,

$$F_0 = F_t + ovp_t + ovs_t.$$

According to the convolution model (and assuming first-order overlap only),

$$ovp_t = D_p * F_t,$$

$$ovs_t = D_s * F_t,$$

so,

$$F_0 = F_t + (D_p * F_t) + (D_s * F_t).$$

Now, let

ovp_1 = 1st estimate of previous-response overlap,

ovs_1 = 1st estimate of subsequent-response overlap.

Following the steps outlined in the section on Level 2, first F_0 is used to calculate an initial estimate of the subsequent overlap,

$$\begin{aligned} ovs_1 &= D_s * F_0 = D_s * [F_t + ovp_t + ovs_t] \\ &= (D_s * F_t) + (D_s * ovp_t) + (D_s * ovs_t) \\ &= (D_s * F_t) + [D_s * (D_p * F_t)] + [D_s * (D_s * F_t)]. \end{aligned}$$

The last term here is essentially the overlap from the second-order subsequent event, so it would be activity that is highly shifted to the right on the time axis. Because in this analysis such activity is assumed to be relatively remote and negligible, it will be dropped. Thus,

$$ovs_1 = (D_s * F_t) + [D_s * (D_p * F_t)].$$

Subtracting from F_0 to obtain the first new template, F_{1s},

$$F_{1s} = F_0 - ovs_1$$
$$= [F_t + (D_p * F_t) + (D_s * F_t)]$$
$$- \{(D_s * F_t) + [D_s * (D_p * F_t)]\}.$$

Cancelling terms,

$$= [F_t + (D_p * F_t)] - [D_s * (D_p * F_t)].$$

Using this waveform as the template to obtain the first estimate of the previous-response overlap, convolve with D_p,

$$ovp_1 = D_p * F_{1s}$$
$$= \{(D_p * F_t) + [D_p * (D_p * F_t)]\} - \{D_p * [D_s * (D_p * F_t)]\}.$$

Dropping the second-order adjacent-event term,

$$ovp_1 = (D_p * F_t) - \{D_p * [D_s * (D_p * F_t)]\}.$$

Subtracting from F_0 to obtain a new template, F_{1p}, which this time will have the previous- (but not the subsequent-) overlap estimate removed,

$$F_{1p} = F_0 - ovp_1$$
$$= [F_t + (D_p * F_t) + (D_s * F_t)]$$
$$- \big((D_p * F_t) - \{D_p * [D_s * (D_p * F_t)]\}\big).$$

Cancelling terms,

$$F_{1p} = F_t + (D_s * F_t) + \{D_p * [D_s * (D_p * F_t)]\}.$$

The next step is to iterate, now using F_{1p} (rather than F_0) as the new template for subsequent responses,

$$ovs_2 = D_s * F_{1p}$$
$$= (D_s * F_t) + [D_s * (D_s * F_t)]$$
$$+ \big(D_s * \{D_p * [D_s * (D_p * F_t)]\}\big).$$

Dropping the second-order adjacent-event term,

$$ovs_2 = (D_s * F_t) + \big(D_s * \{D_p * [D_s * (D_p * F_t)]\}\big),$$

and subtracting from F_0 and then cancelling terms,

$$F_{2s} = F_0 - ovs_2$$
$$= [F_t + (D_p * F_t) + (D_s * F_t)]$$
$$- \big[\{D_s * F_t\} + (D_s * \{D_p * [D_s * (D_p * F_t)]\})\big]$$
$$= [F_t + (D_p * F_t)] - \big(D_s * \{D_p * [D_s * (D_p * F_t)]\}\big).$$

F_{2s} would then be convolved with D_p to calculate ovp_2, and so on, but the pattern should now be clear. The general forms for the subsequent and previous overlap estimates are, respectively,

$$ovs_i = (D_s * F_t) + \big(D_s * \{D_p * [D_s * (\ldots F_t)]\}\big),$$

$$ovp_i = (D_p * F_t) - \big[D_p * \big(D_s * \{D_p * [D_s * (\ldots F_t)]\}\big)\big].$$

Thus, the error in these overlap estimates at the ith iteration can be expressed as

$$ovs_i - ovs_t = D_s * \{D_p * [D_s * (\ldots F_t)]\},$$

$$ovp_i - ovs_t = -D_p * \big(D_s * \{D_p * [D_s * (\ldots F_t)]\}\big).$$

Thus, the residual error in the overlap estimates at each iteration is, to a first approximation, equal to the repeated convolutions of F_t alternatively with D_s and D_p. Such repeated convolutions are equivalent to repeated filtering of the adjacent responses with the filter that corresponds to (convolution with) these event distributions. Assuming these distributions are the rectangular ones typically used in ERP experiments, this would be equivalent to repeated application of the corresponding low-pass (sinc function) filters, such as shown in Figures 4c, 5a, or 5b. Because the gain of such filters is less than 1.0 at essentially all frequencies (except DC), and substantially less than 1.0 at all but the lowest frequencies, repeated application of such a filter would tend to eventually attenuate all but the lowest frequencies in the residual error [i.e., (gain ** i) $\rightarrow$ 0, if gain $<$ 1.0]. Therefore, the magnitude of the residual error will tend to become smaller and smaller with each iteration, and the iterative process will tend to converge.

This analysis also reveals the basic frequency dependence of the convergence—namely, that the low-frequency overlap error would be slowest to converge due to the gain value of the filter(s) being closer to 1.0 at such frequencies. Because the rate at which these filter gain functions approach 1.0 at low frequencies (see Figure 5) is inversely related to the ISI jitter width, T_{jw}, the use of wider jitter not only will tend to reduce the initial distortion of the ERP averages due to overlap, it will also facilitate the convergence process of Level 2 iterations. The present analysis also suggests, however, that to facilitate convergence some initial high-pass filtering to attenuate the very lowest frequencies (those less than, say, 0.25 of $1/T_{jw}$) may often be worthwhile. If such filtering is done digitally after data collection, the filters should be causal because of some assumptions of causality in the technique (e.g., the zeroing of the prestimulus baseline).

Future research might focus on analyzing in more detail the convergence properties of this sequence of convolutions, as well as on pursuing a more complete and rigorous analysis of this process wherein the second- and higher-order overlap terms are explicitly included.