Brown, C., & Hagoort, P. (2000). On the electrophysiology of language comprehension: Implications for the human language system. In M. W. Crocker, M. Pickering, & C. Clifton Jr. (Eds.), Architectures and mechanisms for language processing (pp. 213-237). Cambridge, U.K.: Cambridge University.

9

On the Electrophysiology of Language Comprehension: Implications for the Human Language System

COLIN BROWN AND PETER HAGOORT

This chapter is on syntactic and semantic processes during on-line language comprehension. We present event-related brain potential (ERP) data from a series of experiments using both written and spoken input, focussing mainly on sentence processing. The data are discussed in terms of the constraints that they impose on the architecture of the human language system, in particular with respect to the separation of syntactic and semantic knowledge bases.

The issue of the separation of different linguistic knowledge sources is an underlying theme in much psycholinguistic research and is especially relevant for a number of contrasting models on parsing. Although linguistic theory has to a large extent postulated distinct representational systems for structure and meaning, psycholinguistic models of the parsing process differ in their assumptions about the separability of syntactic and semantic knowledge bases and of the processing operations that tap into these bases.

One class of models – referred to as garden-path models in the literature – posits a separate syntactic knowledge base. This linguistic knowledge is used for the computation of a separate intermediate level of representation for the syntactic structure of a sentence. Garden-path models claim that the construction of an intermediate syntactic level is a necessary and obligatory step during sentence processing, if not for all structural assignments, then at least for a significant subset of them (e.g., Clifton & Ferreira, 1989; Frazier, 1987; Frazier & Clifton, 1996; Friederici & Mecklinger, 1996; Rayner, Garrod, & Perfetti, 1992). In terms of the operational characteristics of the system, it is posited that nonsyntactic sources of information do not affect the parser's initial structurally based analysis. Instead, these sources serve to confirm or disconfirm the first parse.

Another class of models (known as interactionist models, originating in part from the connectionist tradition), does not contain any independent status for (intermediate) products of syntactic computation (e.g., Bates, McNew, MacWhinney, Devescovi, & Smith, 1982; Elman, 1990; McClelland, St. John, & Taraban, 1989; Taraban & McClelland, 1990). The McClelland et al. (1989) sentence processing model, for example, posits a single, undifferentiated representational network, in which syntactic and semantic constraints (among others) combine to influence a single representation. In line with this architectural assumption, the on-line comprehension process is characterized as fully interactional, with all sources of information immediately affecting the analysis process.

More recently, other kinds of interactive models have been proposed, known as constraint-satisfaction models (e.g., Boland, Tanenhaus, & Garnsey, 1990; MacDonald, Pearlmutter, & Seidenberg, 1994; Spivey-Knowlton & Sedivy, 1995). In these models, lexical factors play a central role, and the importance of lexically represented information such as verb frequency and conceptual-semantic knowledge is emphasized. Some proponents of these lexicalist models argue that "there is no need for either an initial category-based parsing stage or a separate revision stage" (Tanenhaus & Trueswell, 1995, p. 233), thereby not assigning any specific primacy to syntactic information and syntactic processes. Note that in this approach a separate syntactic stratum can be part of the model (although the architectural aspects of constraint-based models remain at present somewhat underspecified, cf. Tanenhaus & Trueswell, 1995; but see MacDonald et al., 1994). Therefore constraint-based models share at least some representational assumptions with garden-path models.

The clearest representational contrast, then, emerges between the garden-path and interactionist models. These two contrasting classes of models imply quite different architectures underlying sentence processing. One assumes some form of compartmental representation along the lines of basic linguistic distinctions. The other opts for a combined system that does not differentiate at the representational level. Although many aspects of language are involved in sentence processing, a major representational distinction between the two approaches concerns the disputed separation of syntax and semantics, and their related processes. It is this issue that we will address on the basis of the available ERP results. Constraint-based models sit somewhere in between the garden-path and interactionist models. On the one hand, the constraint-based approach proposes at least partly separate representations, but on the other hand, the processing operations are highly interactive.

It is important to reiterate that the representational dispute shares roots with an ongoing debate on the autonomous or interactive processing nature of the parser. Autonomists claim that in sentence comprehension a syntactic parse is first performed, based on syntactic principles only, before other kinds of information (such as that derived from semantics and pragmatics) are brought to bear on the comprehension process (e.g., Ferreira & Clifton, 1986; Frazier, 1987, 1990; Rayner, Carlson, & Frazier, 1983; Rayner et al., 1992). In contrast, interactionists state that nonsyntactic sources of information are used either to direct the parser's initial analysis (e.g., Bates et al., 1982; Holmes, 1987; McClelland et al., 1989), or to immediately evaluate the product of syntactic analysis on a word-by-word basis, as part of the process of constructing a semantic representation of the input (e.g., Altmann & Steedman, 1988; Crain & Steedman, 1985).

Clearly, anyone of the autonomist persuasion has to be committed to a separate level of syntactic representation during sentence processing. An interactionist account of parsing does not by necessity have to posit such a separate level. Given that under this view syntactic information is directly integrated with semantic and pragmatic information, there is little reason to presuppose separate representational tiers during processing. This is the position advocated by Bates, Elman, McClelland, and their coworkers. Constraint-based theorists such as MacDonald, Tanenhaus, Trueswell, and their colleagues are less radical in their architectural assumptions. They reserve an important role for syntactic representations, but at the same time emphasize the dynamic aspects of the constraint-satisfaction process. This process is characterized by competition among incompatible alternatives (e.g., multiple parses), based on both linguistic and nonlinguistic information. An important difference from the garden-path models is that constraint-based models do not assign any primacy to syntactic computations or to syntactic structure building.

In sum, the different architectural assumptions just outlined also relate to different positions concerning the nature of the processing mechanisms that are operative during on-line sentence processing. Therefore, data that bear on the representational debate on syntax and semantics can also have implications for the autonomous versus interactive processing debate. This is what makes the syntax-semantics interface such an important meeting ground for central issues in parsing research.

In the following, we will discuss several ERP experiments on on-line semantic and syntactic processing during language comprehension. Since extensive overviews of the literature on ERPs and language are already available (cf. Kutas & Van Petten, 1988, 1995; Osterhout & Holcomb, 1995; Van Petten & Kutas, 1991), we have chosen to focus on a set of results – primarily from our own laboratory – that provide a clear basis for the points about the architecture of the sentence processing system that we will be making further on in this chapter. Our main claim will be that distinct electrophysiological signatures can be found for semantic and syntactic processes, and, hence, that the ERP data provide evidence for the existence of different brain states for parsing and semantic processing.

We precede the discussion of the ERP results by a short section on the ERP method and the relevance of ERPs for sentence processing research in particular.

9.1 Event-related Brain Potentials and Language Research

Event-related brain potentials (or ERPs) are part of the brain's overall electrical activity, the electroencephalogram (or EEG). ERPs recorded from the scalp reflect the summation of postsynaptic activity in a group of synchronously firing neurons, all having approximately the same geometric (usually parallel) configuration, and the same orientation with respect to the scalp (for detailed discussion of the physiology of ERPs, see Allison, Wood, & McCarthy, 1986; Nunez, 1981, 1990; Picton, Lins, & Scherg, 1995; Wood, 1987). In contrast to the continuously fluctuating voltage variation over time that constitutes the EEG, ERPs represent a series of voltage changes within the EEG that are time-locked to the presentation of an external stimulus. Because the size of the voltage changes in ERPs is small in comparison with the fluctuation of the EEG (1 or 2 microvolts for ERPs, in comparison to between 10 and over a 100 microvolts for the EEG), it is a standard usage in cognitive ERP research to compute an averaged ERP. This average is calculated over a number of EEG epochs, each of which is time-locked to repetitions of the same event. What is repeated can be the exact same stimulus, but more often is different tokens of the same type. The assumption that motivates signal averaging is that the electrical activity that is not related to the processing of the external stimulus varies randomly over time across the individual epochs. The effect, then, of averaging individual time-locked ERP waveforms is that the randomly distributed voltage fluctuations will tend to average to zero, leaving a residual waveform that reflects the activity that is largely invariant over time and between separate eliciting events of the same type. In language research, an acceptable signal-to-noise ratio can be achieved by averaging over anything between roughly 25 and 60 or more trials, depending on the particular issue under investigation.

Within the series of voltage changes that make up an averaged ERP, a number of positive and negative polarity peaks, more commonly called components, can be identified. In the psychophysiological literature these components are broadly categorized into exogenous and endogenous ones. Exogenous components are particularly sensitive to the physical parameters of the external stimulation and are thought to be relatively insensitive to cognitive factors. In contrast, endogenous components vary as a function of specifically cognitive aspects such as task relevance and attention, and are not contingent upon aspects of the physical stimulation.[1]

Exogenous components are sometimes referred to as early components, reflecting the fact that they occur in the first part of the ERP waveform. To a rough first approximation, the endogenous components that are of interest for language researchers do not emerge until some 150 ms after stimulation.

Although the identification of endogenous components in particular remains a controversial issue (cf. Coles & Rugg, 1995), there are only three basic defining ingredients: polarity, latency, and distribution. Polarity has two values, negative or positive, and components are appropriately labelled by either an N or a P. Latency (in almost all cases of the peak of a component) is measured in milliseconds and can range from just a couple of milliseconds (as with very early auditory components) to several hundreds of milliseconds. Components are additionally labelled according to the latency in milliseconds at which their amplitude reaches its maximum. For example, N400 refers to a negative polarity component with a peak value at approximately 400 ms after stimulus onset. In a few cases, components have received functional labels, intended to reflect the process they are thought to be a manifestation of (e.g., the mismatch negativity, see Näätänen, 1992). In addition to their polarity and latency, components are also characterized by their distribution over the scalp, or scalp topography. This refers to the pattern of latency and peak amplitude values over electrode sites. As a rule, an ERP experiment entails the registration of activity from a number of electrodes, distributed in a more or less even fashion over the scalp. Many components (both exogenous and endogenous) show a graded distribution of amplitude values over electrode sites, and this distribution can serve as one of the ways in which to distinguish between components, especially if the temporal windows within which components occur partly overlap in time. It is, however, very important to note that the activity at any particular electrode site on the scalp cannot be taken to originate in the brain tissue directly underlying that site: Scalp topography does not provide a map of neuronal localization. This is due, among other things, to the volume conduction properties of the brain and its surrounding matter, which enable electrical activity generated in one area of the

[1] The exogenous-endogenous distinction is less clear than this brief description suggests. Many components that were first thought to be either strictly exogenous or endogenous have since been shown to be open to the influence of physical or cognitive factors. However, the distinction is still helpful, and under appropriately operationalized experimental conditions, the separation into exogenous and endogenous components can be largely upheld. See Rugg and Coles (1995) for an excellent discussion of this and related issues.

brain to be registered at locations at considerable distance from the generator site.

There are several advantages to using ERPs in the investigation of sentence processing. The first concerns the multidimensional nature of ERPs. Brain potentials can independently vary according to their polarity, latency, amplitude, and scalp topography. This provides a rich basis for distinguishing among separate processing events. One particularly interesting possibility is that ERPs can in principle establish a truly qualitative distinction among processes. On the basis of our knowledge of the physiological origins of ERPs, it is reasonable to assume that certain different types of ERP peaks (e.g., those of opposite polarity) are generated by separate, or at least not entirely overlapping, neural systems. This implies that under appropriately operationalized experimental conditions, we can infer from the presence of qualitatively different electrophysiological profiles that distinct language processes are operative (cf. Kutas & King, 1995; Osterhout & Holcomb, 1995).

A second advantage of ERPs is that they provide a continuous, real-time measure of the language comprehension process. This enables the uninterrupted measurement of activity throughout and even beyond the entire processing of a sentence or discourse. Therefore, it is possible to not only assess the immediate impact of an experimental manipulation (e.g., the presentation of a syntactic anomaly), but also the consequences of this manipulation for the processing of further incoming material.

A third appealing feature of ERPs for sentence processing research is that robust, statistically reliable, and replicable effects can be obtained in the absence of additional task requirements (a feature that ERPs share with eye-movement registration). In most of the experiments that we will report on here, subjects were only instructed to attentively read or listen to the sentences they were presented with. No extraneous task demands were imposed. The advantage here is that we are not plagued by the uncertainties that can accompany reaction-time research on sentence processing, where we always have to take into account the possibility of contamination due to task effects.[2] Of course, by not having the subjects perform any overt task during ERP registration, we are open to the criticism that we have no control over exactly what the subjects are doing during the experiment. However, we would argue that with attentive and cooperative subjects, who are instructed to comprehend the sentences they hear or

[2] A further advantage is that by avoiding an *overt* task it becomes possible to test subject groups who cannot cope with additional task demands at the same time as adequately processing linguistic stimuli. See Hagoort, Brown, and Swaab (1996) for an example of ERP registration in aphasic patients with severe comprehension deficits.

read, the main task that these subjects are engaged in is the normal process of comprehension. It is, after all, quite difficult if not impossible to not process and understand language when it is presented within the focus of attention (cf. Fodor, 1983). Moreover, in many cases the proof of the pudding is in the eating. That is, if in the absence of extraneous task demands we obtain clear effects as a function of our experimental manipulations, and if these effects can be sensibly interpreted in terms of a priori predictions based on a sufficiently articulated model, then it seems reasonable to claim that we have succeeded in obtaining meaningful insights into on-line sentence processing.

9.2 Semantic Processing and Electrophysiology

In 1980, Marta Kutas and Stephen Hillyard reported a finding that provided the starting point for what has since become a very active research area on the electrophysiology of language comprehension (Kutas & Hillyard, 1980). They presented subjects with written sentences that ended in a semantically congruous or incongruous word. The ERP elicited by the incongruous ending showed a monophasic component with a negative polarity that reached its maximal amplitude at approximately 400 ms after presentation of the sentence-final word. The congruous last word of the sentence elicited the same component, but its amplitude was significantly smaller than with the incongruous ending. In accordance with its polarity and peak latency, Kutas and Hillyard termed the component the N400. The difference in amplitude between the two conditions is referred to as the N400 effect. An extension of the seminal 1980 result is shown in Figure 9.1, from the work of Kutas, Lindamood, and Hillyard (1984). The figure shows data from one representative electrode site on the scalp, labelled Pz, which is located over the central midline on the back of the head. In the Kutas et al. experiment, the sentence ended in three different ways: (1) on a word that was entirely congruous with the preceding context (the best completion condition), (2) on a word that was semantically related to the congruous ending, but that was nevertheless anomalous with respect to the meaning of the preceding sentence (related anomaly), or (3) on a word that had no relation with either the best completion or the sentential context (unrelated anomaly).

The waveforms show a series of voltage fluctuations. The effect of the physical stimulation can be seen in the exogenous potentials that are present in roughly the first 150 ms of the waveform. There is an early negative polarity peak at about 100 ms, immediately followed by a positive polarity peak. This so-called N1-P2 complex is a characteristic exogenous ERP profile that is invariably elicited by visual stimulation. The exogenous nature of this complex is underscored by the fact that it is identical in all three experimental conditions.

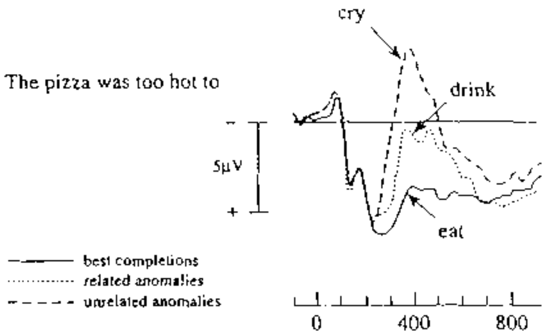

Figure 9.1. Grand average ERPs at electrode site Pz for three visually presented different sentence-final completions preceded by the same sentence. The solid line is the waveform for words that best completed the sentences. The dotted line is the waveform for words that did not fit the sentential contexts, but that were semantically related to the best completions. The dashed line is for words that did not fit the preceding contexts and that were not related to the best completions. In this figure, and all following figures, negative voltage values are plotted upwards, and the time axis is in milliseconds. Figure adapted from Kutas, Lindamood, and Hillyard, © 1984 The Psychonomic Society.

More important for our present purpose is the clear N400 component that is present in each condition. The onset of the N400 is at about 250 ms, with its characteristic peak at 400 ms. As can be clearly seen in the figure, the amplitude of the N400 is modulated as a function of the semantic match between the sentence-final word and the preceding context. The smallest N400 is elicited by the best completions, and the largest by the unrelated anomalies. The amplitude of the N400 to the related anomalies lies in between, reflecting the partial semantic match with the best completion. What these data demonstrate is that modulations in the amplitude of the N400 can be brought about by manipulating the semantic appropriateness of words with respect to the context in which they appear.

In the early research on the N400, semantic incongruities were used to assess the sensitivity of the N400 to contextual information. Subsequent research has shown that the N400 is also elicited in and modulated by more subtle semantic contexts. Figure 9.2 shows an example from our own work (Hagoort & Brown, 1994), in which we manipulated the cloze probability of words (that is, the extent to which a particular word is expected to occur in a sentence, based on the information conveyed by the sentence preceding that word). Compare the sentence *Jenny put the sweet in her mouth after the lesson* with *Jenny put the sweet in her pocket after the lesson*. Here, we contrast the ERP elicited by the word *mouth*, which is expected given the context (the high cloze condition), with

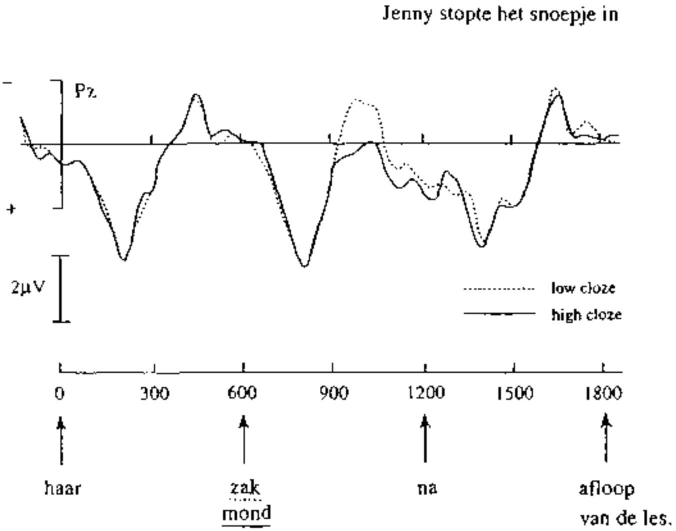

Figure 9.2. Grand average ERPs at electrode site Pz for visually presented less expected (Low Cloze, dotted line) and more expected (High Cloze, solid line) words in mid-sentence position. Presentation onset of the critical words is at 600 ms. In the figure, the critical words are preceded and followed by one word. The translation of the example sentence is *Jenny put the sweet in her pocket/mouth after the lesson*. Figure adapted from Hagoort and Brown, © 1994 Erlbaum.

the word *pocket*, which is less expected (low cloze). Note that in both cases the words are entirely acceptable continuations of the sentence. It just happens to be the case that when presented with the sentence preceding the high and low cloze words, more people choose *mouth* as a more likely continuation than *pocket*.

As can be seen in Figure 9.2, both the high and the low cloze words elicit the N400 component, but with a significantly larger amplitude in the low cloze condition. This demonstrates that the N400 is not a simple incongruity detector, but can reflect quite subtle aspects of ongoing semantic processing.

The majority of N400 research has used visual presentation. However, the elicitation of the N400 is not modality-dependent (cf. McCallum, Farmer, & Pocock, 1984; Connolly & Phillips, 1994), although its latency characteristics can differ depending on whether written or spoken language stimulation is used. In particular, it has been claimed that the onset of the N400 effect can be earlier with spoken input (cf. Holcomb & Neville, 1990; but see Connolly & Phillips, 1994, Hagoort & Brown, 1997). Figure 9.3 gives an example from our own work. Subjects heard naturally produced connected speech, in which in half of the sentences a semantic anomaly occurred in sentence-medial position.

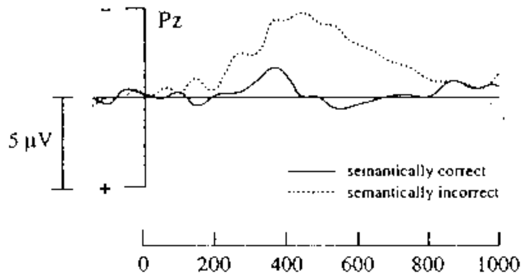

Figure 9.3. Grand average ERPs at electrode site Pz for semantically incorrect (dotted line) and correct (solid line) words occurring in mid-sentence position in naturally produced connected speech. The onset of the critical words is at 0 ms in the figure.

The waveforms in Figure 9.3 look quite different from those in the previous figures. The early exogenous stimulus components that can be seen in visually elicited ERPs are not discernable. This is due to the continuous physical stimulation by the speech signal, which gives rise to a series of temporally overlapping stimulus components. This overlap, in combination with the refractory period of auditory evoked potentials, leads to a smearing of the exogenous potentials, which in turn results in the relatively smooth morphology of the initial part of the waveform. Nevertheless, the waveforms clearly show an effect of the semantic anomaly, with a significant negative shift for the anomalous compared to the congruent condition. Based on its temporal properties and its distribution over the scalp (for expository purposes, we show only one electrode site here), it is clear that the difference between the anomalous and the incongruous condition can be classified as an N400 effect.

An additional observation on the modality independence of the N400 comes from work by Kutas, Neville, and Holcomb (1987). These researchers compared the N400 response to semantic anomalies during reading, listening, and signing. In the signing condition, the subjects were congenitally deaf users of American Sign Language (ASL). When presented with a semantic anomaly in ASL, these subjects showed a reliable N400 effect that was comparable to the N400 effects of the subjects who participated in the reading and listening conditions. This attests to the sensitivity and the validity of the N400 as an index of semantic processing during language comprehension (see Neville, Coffey, Lawson, Fischer, Emmorey, & Bellugi, 1997, for further, confirmatory evidence on ERP language effects in users of ASL).

Since the original report on the N400, a host of experiments has been performed on the conditions under which this component is elicited and modulated. We have discussed only a couple of experiments, focussing on the N400 in sentential contexts. In addition to the sentential work, a considerable amount of research has focussed on the N400 in single-word and word-word contexts. The net result of this wide-ranging research program is too extensive to discuss here (we refer the reader to the literature reviews mentioned earlier), but we conclude this section by briefly listing some of the major characteristics known to hold for the N400, focussing on those aspects that are particularly relevant for sentence processing research.

1. Each open class word as a rule elicits an N400.
2. The amplitude of the N400 is inversely related to the cloze probability of a word in sentence context. The better the fit between a word and its context, the smaller the amplitude of the N400.
3. The amplitude of the N400 varies with word position. The first open-class word in a sentence produces a larger negativity than open-class words in later positions. This reduction most likely reflects the increase in semantic constraints throughout the sentence.
4. The N400 is elicited by spoken, signed, and written language.
5. The N400 is not directly elicited by grammatical processes, although it can follow from them (see further on for more details on this point).

The overriding finding that emerges from the literature is that the amplitude fluctuations of the N400 are a function of ongoing semantic processing, ranging from outright anomalies to subtle variations in the goodness-of-fit of words in context. We have argued elsewhere (cf. Brown & Hagoort, 1993; Chwilla, Brown, & Hagoort, 1995; Hagoort et al., 1993) that the functional interpretation of the N400 effect relates to lexical integration processes. That is, following access to the mental lexicon, the activated word meaning has to be integrated into a message-level representation of the context within which that word occurs. It is this meaning integration process that is reflected by the N400 effect. The easier the integration process is, the smaller the amplitude of the N400 (for a similar position see Holcomb, 1993; Osterhout & Holcomb, 1992; Rugg, 1990; see also Kutas & King, 1995).

9.3 Syntactic Processing and Electrophysiology

Until some five years ago, almost no ERP results had been published on syntactic processing during language comprehension. Early work by Kutas and Hillyard (1983) included grammatical agreement errors as one of the conditions. No clear syntactic effects emerged that were dissociable from the standard N400 component, other than an enhanced negativity in the 300 to 500 ms range. With the upsurge in ERP and language research in the early 1990s, several groups have begun to focus on ERP manifestations of syntactic processing. As a result, there are now two basic ERP effects that are thought to reflect aspects of the parsing process.

The first effect concerns a relatively early negative shift with a peak latency at about 250 ms following relevant stimulation, which has been reported to be primarily observed at electrode sites over left anterior regions of the scalp. Hence, this effect is referred to as the left-anterior negativity, or LAN (Kluender & Kutas, 1993). The LAN has been observed following phrase structure violations (Neville, Nicol, Barss, Forster, & Garrett, 1991; Friederici, Pfeifer, & Hahne, 1993), subject-object relative sentences (King & Kutas, 1995),[3] long-distance dependencies (Kluender & Kutas, 1993), and word category violations (Münte, Heinze, & Mangun, 1993). At present, the functional interpretation of the LAN remains unclear. Some authors see the LAN as a reflection of differential working memory load related to thematic integration processes (King & Kutas, 1995; Kluender & Kutas, 1993). Others interpret it as an index of initial syntactic assignment on the basis of word category information (Friederici, 1995; Friederici & Mecklinger, 1996). These disparate views reflect the current lack of understanding about the functional nature of the LAN. Clearly, further research is called for. One issue that will need to be addressed is the extent to which we are actually dealing with one and the same effect. Some caution is called for here in the light of the variability in both the peak latency of the LAN and its distribution over the scalp. The peak latencies vary between 200 and 400 ms, with the range of the entire enhanced negativity running from just over 100 ms to beyond 500 ms. This leaves considerable scope for multiple processes to be operative. Furthermore, although the maximal effect is indeed observed over left anterior regions of the scalp, in most studies it is not restricted to these regions. The enhanced negativity has been reported to extend over left temporal, central, and parieto-temporal electrode sites, and it has even been observed over right anterior regions of the scalp.

The present uncertainties about the specificity of the enhanced negativity should not be taken to imply that the reported ERP effects do not reflect on-line comprehension processes. The number of published reports in which enhanced negativities have been reported is by now large enough to suggest that something interesting is emerging here. However, given the variety of linguistic manipulations that give rise to the effect(s), and given the variability in latency

[3] But see Mecklinger, Schriefers, Steinhauer, and Friederici (1995), who report a positivity with a peak latency of 345 ms for object-relative sentences.

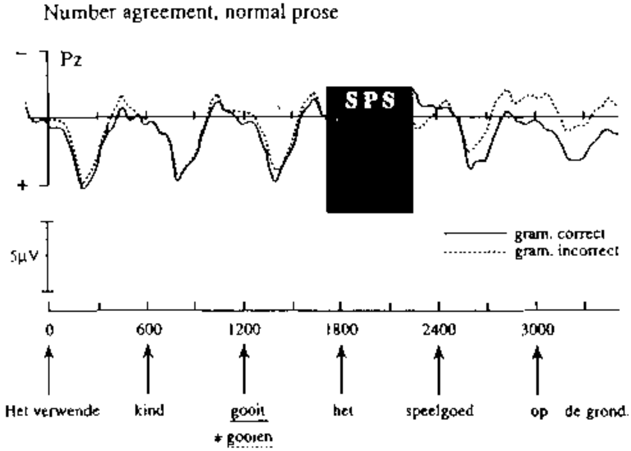

Figure 9.4. Grand average ERPs at electrode site Pz for visually presented grammatically incorrect (subject-verb number agreement, dotted line) and correct words (solid line) in mid-sentence position in normal prose. The critical words were presented at 1200 ms. The figure shows the critical words preceded by two and followed by three words. The translation of the example sentence is *The spoilt child throws/throw the toy on the ground.* The region within which the Syntactic Positive Shift developed is shaded, and labelled with SPS. Figure adapted from Hagoort and Brown, © 1994 Erlbaum.

and distribution, more detailed experimentation is required before the exact relationship with the comprehension process is elucidated.

The second effect that has been reported in the ERP literature on sentence processing is a late positivity. This effect starts at about 500 ms, extends for several hundred milliseconds, and has a broad distribution over the scalp, with a centro-parietal maximum. The effect was independently reported by two research groups, one working with American English (Osterhout & Holcomb, 1992), the other with Dutch (Hagoort, Brown, & Groothusen, 1993). Osterhout and Holcomb (1992) examined ERPs in response to violations of verb subcategorization and phrase structure constraints. Hagoort et al. (1993) additionally included agreement violations between a subject NP and the finite verb (e.g., **The spoilt child* throw *the toy on the ground*). An example from the latter work is shown in Figure 9.4, which depicts the ERP elicited by a number agreement, in comparison with its grammatically correct counterpart, using visual presentation.

The agreement error elicited a sustained positive shift, starting at 500 ms after the incorrect word (i.e., *throw*). This effect has been labelled the P600 (Osterhout & Holcomb, 1992), or the Syntactic Positive Shift (SPS; Hagoort et al., 1993). We will use the latter term in this chapter. As can be seen in Figure 9.4, the SPS is not the only processing effect that emerges in the ERP waveform. After the SPS in response to the grammatical violation, the waveform transforms into a negative shift for the next words in the incorrect condition, relative to the control condition. On the basis of its morphology and scalp topography, this negative shift can be classified as an N400 effect. We hypothesize that the origin of the N400 effect lies in the preceding syntactic violation. Due to the agreement error, the experimental subject is confronted with a problem in constructing a syntactic representation for the string of words. This structural problem has consequences for the ease with which further incoming words can be processed. In particular, an inadequate structural representation creates a problem for the integration of words into the overall message representation of the sentence. It is this integrational problem that emerges as an N400 effect.

In the Hagoort et al. (1993) experiment, the same pattern of positive and negative polarity effects was obtained for phrase structure violations.[4] Here, subjects read sentences in which the obligatory Dutch word order of adverb–adjective–noun was altered, such that the adjective preceded the adverb (e.g., **The man was startled by the emotional rather response of his partner*). An SPS was observed at the grammatical violation (i.e., *response*), in combination with an N400 effect on the words following the violation. An additional and important effect in terms of the functional characterization of the SPS was the occurrence of an SPS on the adverb preceding the noun (i.e. *rather*). At this point in the sentence, a grammatically correct continuation is still possible (e.g., *the emotional rather outspoken response*), but the syntactic structure of such a continuation is more complex than the more commonly occurring adverb–adjective–noun sequence. We propose that the SPS elicited by the adverb reflects a processing effect related to the fact that at this position a non-preferred syntactic structure (i.e., either a more complex, or a less frequent, or both) has to be assigned to the sentence (cf. Hagoort et al., 1993; Hagoort & Brown, 1994).

The sensitivity of the SPS to syntactic violations provides a potentially revealing finding on the nature of the syntax-semantic interface involved in sentence processing. However, before the significance of this finding can be fully assessed, more evidence is needed on the validity of the SPS. We will discuss two issues here. First, is the SPS specific to syntactic processing during on-line language comprehension? In particular, can we reliably isolate the SPS

[4] But not for subcategorization violations. See Hagoort et al. (1993) and Hagoort and Brown (1994) for discussion.

from the N400? Second, how sensitive is the SPS with respect to the on-line comprehension process? In particular, can the SPS also be observed during the processing of sentences that do not contain outright violations? Although this latter question has in part been answered by the SPS in response to the adverb in the ungrammatical phrase structure condition, further data on the SPS in non-violating contexts is required. We conclude this section on syntactic processing with a few remarks on the functional characterization of the SPS.

9.3.1 The Specificity of the SPS

Under normal stimulation conditions, it is notoriously difficult to factor out the contributions to understanding of possibly distinct sources of linguistic information. However, in the case of syntax, so-called syntactic prose offers a way to especially focus on syntactic processing. In standard syntactic prose experiments, subjects read or listen to sentences that are semantically uninterpretable, but accord with the grammatical rules of the language. For example, *The boiled watering can smokes the telephone in the cat.*[5] Despite the fact that this string of words does not convey any coherent meaning, subjects readily grasp the grammatical relations that hold between the words and are able to parse the sentence into its constituent parts. This demonstrates that in the absence of a message-level meaning representation, subjects do activate syntactic knowledge and use this knowledge to parse the sentence.

We used a syntactic prose manipulation to investigate the specificity of the SPS. The main issue was whether an SPS would be observed to grammatical violations in syntactic prose. For example, in the sentence *The boiled watering can* smoke *the telephone in the cat*, *smoke* creates an agreement error in combination with its preceding subject noun phrase. This grammatical error holds independently of the semantic uninterpretability of the sentence. The question then is, if we strip a sentence of meaning, and only vary its grammaticality, is an SPS still elicited by the word that instantiates the ungrammaticality, just as in normal prose?

A second and related question concerns the independence of SPS and N400 effects. In the normal prose experiment discussed previously, the SPS was followed by an N400 effect. We interpreted this effect as a reflection of problems with meaning integration, originating from the preceding grammatical

[5] Note that this kind of prose is different from so-called Jabberwocky, made famous by Lewis Carroll, which has been used by psycholinguists to investigate language comprehension. Jabberwocky contains a mixture of real words and pseudo-words. Syntactic prose contains only real words, so that the categorical and morphological lexical information is transparent.

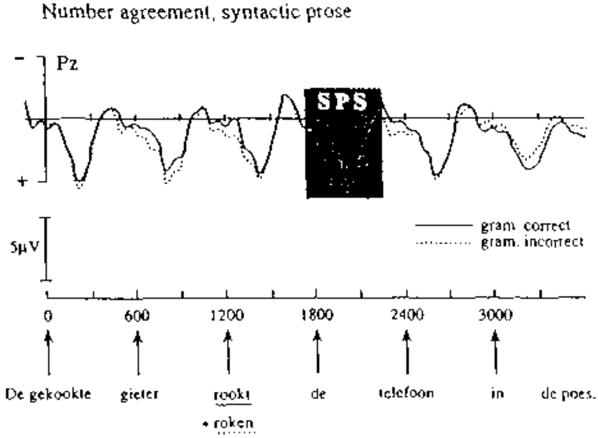

Figure 9.5. Grand average ERPs at electrode site Pz for visually presented grammatically incorrect (subject-verb number agreement, dotted line) and correct words (solid line) in mid-sentence position in syntactic prose. The critical words were presented at 1200 ms. The figure shows the critical words preceded by two and followed by three words. The translation of the example sentence is *The boiled watering can smokes/smoke the telephone in the cat*. The region within which the Syntactic Positive Shift developed is shaded, and labelled with SPS. Figure adapted from Hagoort and Brown, © 1994 Erlbaum.

violation. If this analysis is correct, and if the SPS and the N400 are related to different processing events, then we predict that no N400 effects should be observed in syntactic prose, since in this kind of context no higher-order meaning integration can occur.

Figure 9.5 shows the ERP waveform for the syntactic prose variant of the agreement errors that we investigated in the original normal prose experiment. The structure of the sentences in the syntactic prose experiment was identical to the normal prose sentences they were derived from, and the errors occurred in the same position.

Once again the agreement errors elicited an SPS. The onset of the effect and its morphology and topography are the same as for the normal prose experiment, though the size of the effect is slightly reduced in the syntactic prose experiment. Similar results were obtained for the syntactic prose variant of the phrase structure condition (e.g., *The heel tripped over the rather inhabited/inhabited rather cat on his pocket*). Here, an SPS was observed in response to both the adverb and the noun, replicating the findings of the normal prose experiment. These results provide a clear demonstration that the SPS is indeed a reflection of syntactic processing during language comprehension.

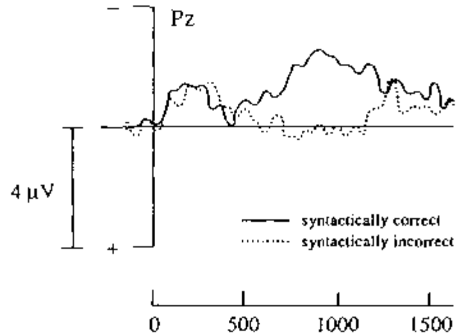

Figure 9.6. Grand average ERPs at electrode site Pz for grammatically incorrect (subject-verb number agreement, dotted line) and correct words (solid line) in mid-sentence position in normal prose with naturally produced connected speech. The onset of the critical words is at 0 ms in the figure.

In addition, the results provide evidence in favour of the independence of the SPS and N400 effects. As predicted, no N400 effects were observed in syntactic prose, neither for the agreement nor for the phrase structure condition.[6] Not only does this underscore the difference between the SPS and the N400, but it also provides supporting evidence for our functional interpretation of the N400 effects that we observed in the normal prose experiment.

All of the results that we have presented so far were obtained on the basis of visual stimulation. Just as with the N400, it is important to verify that the SPS is a modality-independent effect. If the SPS is indeed an electrophysiological signature of syntactic processing, then it should be obtained in both the visual and auditory modality. That this is indeed the case can be seen in Figure 9.6 (which shows data from Hagoort and Brown, 1997). These waveforms are for the correct and incorrect agreement conditions of the normal prose sentences, presented as naturally produced connected speech.

Just as with the connected speech data of the N400 congruity experiment that we presented earlier, here too the overall shape of the waveform does not show readily discernable stimulus components. However, a clear condition effect is observed for the agreement errors, showing a positivity with an onset latency of 500 ms, and a centro-parietal distribution. Although not shown in this figure,

[6] The difference following the agreement error did not begin to approach statistical significance, and was absent at other electrode sites. Note also that the N400 component as such is visible in the waveforms, with clear N400s in response to the individual words in both conditions. This is in accordance with the literature. The crucial point is the absence of an N400 effect (i.e., a difference within the N400 time domain between the two conditions).

just as with visual stimulation, an N400 effect emerges in the ungrammatical condition for words following the syntactic error. These findings are in accordance with the work of Osterhout and Holcomb (1993). We can conclude from these data that the SPS is a modality-independent effect.

9.3.2 The Sensitivity of the SPS

In the experiments discussed so far, most of the manipulations concerned grammatical violations. Although the SPS in response to the adverb in the phrase structure manipulation already indicated that the presence of violations is not a necessary condition for eliciting the SPS, this effect was observed in the context of materials containing a series of grammatical violations. Therefore, just as was the case for the N400 component, the question needs to be addressed whether the SPS is more than a violation detector. We tackled this issue by presenting subjects with structurally ambiguous sentences.

The ambiguity that we focussed on was one of the so-called attachment ambiguities. In our manipulation the sentence ultimately could be assigned only one structural analysis, but the initial part could be assigned two different analyses. For example: *The sheriff saw the cowboy and the Indian spotted the horse in the bushes*. Until the second verb (i.e., *spotted*) it is unclear whether the two nouns *cowboy* and *Indian* are part of one, conjoined noun phrase, or whether the second noun (i.e., *Indian*) is the subject of a second clause. Obviously, these two structural analyses are not trivially different, and they have consequences for the overall meaning of the sentence. In the example sentence, if readers opt for the conjoined NP analysis they will be confronted with a parsing problem on the verb *spotted*.

This kind of attachment ambiguity has been under investigation in the literature on the autonomous or interactive nature of the parser. It has been claimed that the conjoined NP analysis is a less complex syntactic structure than the conjoined S analysis (cf. Frazier, 1987). If the real-time operation of the parser is based in part on principles of economy and efficiency, such that less complex structures are preferred over more complex ones, then it follows that in the case of the example sentence, the conjoined NP structure will be considered the more viable analysis.[7] Therefore, at the moment that the parser is confronted with a word that refutes this preferred analysis (i.e., at the verb *spotted*), a parsing problem should occur. Note that we are not dealing with an overt grammatical

[7] In the example that we present here, we have not included an additional semantic manipulation, biasing for one of the two readings. In this chapter, we will not discuss ERP data on semantic effects during on-line parsing. For current purposes, we need to first discover whether the SPS is at all sensitive to structural ambiguities.

Attachment ambiguity

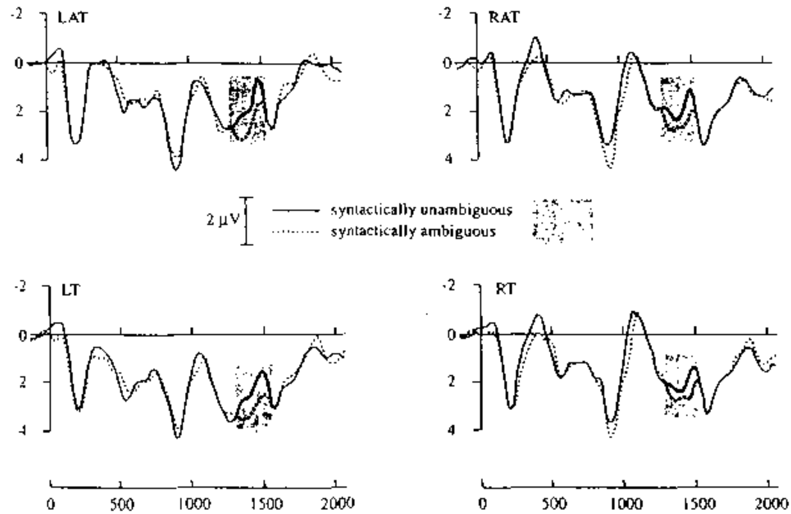

Figure 9.7. Grand average ERPs over left and right anterior temporal (LAT, RAT) and left and right temporal (LT, RT) electrode sites for visually presented sentences. In the ambiguous condition (dotted line) the sentences were initially syntactically ambiguous. At the point of disambiguation (at 686 ms in the figure), the sentence continued with a grammatically correct but non-preferred reading. In the control condition (solid line), unambiguous versions of the same non-preferred structures were presented. In the figure, the disambiguating words are preceded and followed by one word. The region within which the Syntactic Positive Shift (SPS) developed is shaded.

violation here. The sentence is entirely grammatical. The origin of the (putative) problem lies in the proposed operating characteristics of the parser.

Figure 9 7 shows the waveform for the syntactically ambiguous sentence and its control. The zero time point marks the presentation of the noun preceding the second verb. The verb was presented at 686 ms. The control sentence was identical to the ambiguous one, with the exception of the inclusion of a comma after the second noun (i.e., *The sheriff saw the cowboy, and the Indian spotted the horse in the bushes*). In Dutch, a comma in this position is a normal, though not obligatory part of written language. When included, the comma rules out the possibility of conjoining the nouns preceding and following it.

The waveforms for the ambiguous condition deviate from the control condition following the presentation of the second verb (i.e., *spotted*). In the ambiguous condition, this is the position at which the ambiguity is resolved in favour of the second clause reading, which is hypothesized to be the non-preferred structural analysis. Although the scalp topography of the effect is somewhat different from that observed for violations (in the present case the effect has a more anterior distribution), the effect is a positive shift, with an onset latency of 500 ms relative to the onset of the presentation of the critical verb. Given its similarity to the previously observed syntactic effects, we classify this effect as an SPS. Our conclusion, therefore, is that the SPS is not only elicited by syntactic violations, but is also sensitive to processing operations related to parsing preferences.

9.3.3 The Functional Characterization of the SPS

We have provided evidence that the SPS is elicited by a variety of syntactic phenomena. Although we have focussed on our own work, the sensitivity of the SPS (also known as the P600) to syntactic processes has by now been reported by several research groups, working with various aspects of word order, agreement, and subcategorization (cf. Neville et al., 1991; Osterhout & Holcomb, 1992; Osterhout & Mobley, 1995; Osterhout, Holcomb, & Swinney, 1994; see also Friederici & Mecklinger, 1996; Mecklinger et al., 1995). The SPS is obtained in both the visual and the auditory modality, with a remarkably invariant onset latency of 500 ms, and with a broad distribution over the scalp, showing slight topographical variations as a function of whether the SPS is elicited in violating or in ambiguous contexts.

Clearly, then, with the SPS we have an ERP effect in hand that can be used as a tool with which to probe on-line parsing. What remains to be elucidated, and what lies beyond the scope of the currently available data, is the functional nature of the SPS. Exactly which aspect of syntactic processing is reflected by the SPS? On the basis of the work that has been reported, we can confidently claim that the SPS is elicited by the word in a sentence that indicates that the current structural assignment is an incorrect or non-preferred syntactic analysis for the incoming string of words. What still has to be established is whether the elicitation of the SPS is a direct consequence of a failing first parse, or whether the SPS is related to a process of syntactic reanalysis, occurring after a first-pass structural assignment has resulted in a misanalysis. A further issue that needs to be addressed concerns the topographical differences that we have observed between outright syntactic violations and parsing preferences. It is possible that we are seeing a family resemblance among ERP effects related to different aspects of syntactic processing.

9.4 Conclusions

In this chapter, we have presented ERP data on different aspects of language comprehension, comparing electrophysiological manifestations of semantic and syntactic processes. What the data clearly show is that the brain response to semantic processing is distinct from the response to syntactic processing. In particular, the N400 component has proven to be an especially sensitive index of meaning integration processes, whereas within the domain of language processing the SPS is only observed in the context of specifically syntactic processes. It is important to note that we are not claiming that either the N400 or the SPS is unique to the language domain. What we do claim is that during language processing the N400 and SPS are separate components with separate sensitivities. The two effects are dissociable by experimental manipulation: a semantic cloze manipulation elicits only an N400, whereas a syntactic prose manipulation elicits only an SPS. At the same time, they can be observed during the processing of one and the same sentence (e.g., the succession of SPS and N400 effects in the agreement and phrase structure violation conditions), in a manner that makes sense in terms of the ongoing comprehension process. Moreover, the N400 and the SPS are qualitatively entirely different effects in terms of their electrophysiological characteristics: Semantic processing emerges as a negative-going shift in the ERP waveform, whereas syntactic processing emerges as a positive-going shift.

With this evidence in hand, we can return to the issue that we raised at the beginning of this chapter, namely the separation of linguistic knowledge sources during on-line comprehension. Based on the separate identity and sensitivity of the SPS and the N400, our claim is that separable, non-identical brain processes underlie syntactic and semantic processing. If this claim is correct, then it provides a boundary condition for models of language processing. At the very least, these models will have to allow for a qualitative distinction between syntactic and semantic processing effects. This boundary condition is not compatible with interactionist models that argue against the existence of intermediate products of syntactic computation (e.g., Bates et al., 1982; McClelland et al., 1989). If no distinction is made at the representational level, it becomes difficult to account for the different brain responses elicited by syntactic and semantic constraints. In this respect, garden-path models fit neatly with the ERP data that we have discussed in this chapter. Constraint-based lexicalist models, although denying an initial category-based parsing stage (cf. Tanenhaus & Trueswell, 1995), can also accommodate the ERP data, since these models incorporate elaborate lexical-syntactic representations.

In sum, the evidence on the differential sensitivity of the SPS and the N400 is more compatible with models that include a separate level of syntactic computation during the process of language understanding. This should not be taken to imply that we can now distinguish between autonomous or interactive processing accounts of the parser. The mere existence of the N400 and the SPS is insufficient evidence in this respect. However, given these two separate reflections of the brain's electrical activity during language comprehension, we have good tools in hand with which to attempt to obtain further insights into the architecture and mechanisms of the human language system.

References

Allison, T., Wood, C.C., & McCarthy, G.M. (1986). The central nervous system. In M.G.H. Coles, E. Donchin, & S.W. Porges (Eds.), *Psychophysiology: systems, processes, and applications*. (pp. 5–25) New York: Guilford.

Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition* 30, 191–238.

Altmann, G.T.M., Garnham, A., & Dennis, Y. (1992). Avoiding the garden path: Eye movements in context. *Journal of Memory and Language* 31, 685–712.

Bates, E., McNew, S., MacWhinney, B., Devescovi, A., & Smith, S. (1982). Functional constraints on sentence processing: A cross-linguistic study. *Cognition*, 11 245–299.

Boland, J.E., Tanenhaus, M.K., & Garnsey, S.M. (1990). Evidence for the immediate use of verb control information in sentence processing. *Journal of Memory and Language* 29, 413–432.

Brown, C.M., & Hagoort, P. (1993). The processing nature of the N400: Evidence from masked priming. *Journal of Cognitive Neuroscience* 5, 34–44.

Chwilla, D. J., Brown, C. M., & Hagoort, P. (1995). The N400 as a function of the level of processing. *Psychophysiology* 32, 274–285.

Clifton, C., & Ferreira, F. (1989). Ambiguity in context. *Language and Cognitive Processes* 4, 77–104.

Coles, M.G.H., & Rugg, M.D. (1995). Event-related brain potentials: An introduction. In M.D. Rugg, & M.G.H. Coles (Eds.), *Electrophysiology of mind: Event-related brain potentials and cognition*. (pp. 1–26) New York: Oxford University Press.

Connolly, J.F., & Phillips, N.A. (1994). Event-related potential components reflect phonological and semantic processing of the terminal word of spoken sentences. *Journal of Cognitive Neuroscience* 6, 256–266.

Crain, S., & Steedman, M. (1985). On not being led up the garden-path: The use of context by the psychological parser. In D. Dowty, L. Karttunen, & H. Zwicky (Eds.), *Natural language parsing*.(pp. 320–358) Cambridge: Cambridge University Press.

Elman, J.L. (1990). Representation and structure in connectionist models. In G.T.M. Altmann (Ed.), *Cognitive models of speech processing: Psycholinguistic and computational perspectives*. (pp. 345–382) Cambridge, MA: MIT Press.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language* 25, 348–368.

Fodor, J.A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.) *Attention and performance XII. The psychology of reading*. (pp. 559–586) London: Erlbaum.

Frazier, L. (1990). Parsing modifiers: Special purpose routines in the HSPM? In G.B. Flores d'Arcais, K. Rayner, & D. Balota (Eds.), *Comprehension processes in reading*. Hillsdale, NJ: Erlbaum.

Frazier, L., & Clifton, C. (1996). *Construal*. Cambridge, MA: MIT Press.

Friederici, A.D. (1995). The time course of syntactic activation during language processing: A model based on neuropsychological and neurophysiological data. *Brain and Language* 50, 259–281.

Friederici, A.D., & Mecklinger, A. (1996). Syntactic parsing as revealed by brain responses: First-pass and second-pass parsing processes. *Journal of Psycholinguistic Research* 25, 157–176.

Friederici, A.D., Pfeifer, E., & Hahne, A. (1993). Event-related brain potentials during natural speech processing: Effects of semantic, morphological and syntactic violations. *Cognitive Brain Research* 1, 183–192.

Hagoort, P., & Brown, C.M. (1994). Brain responses to lexical-ambiguity resolution and parsing. In C. Clifton, L. Frazier, & K. Rayner (Eds.), *Perspectives on sentence processing*. (pp. 45–80) Hillsdale, NJ: Erlbaum.

Hagoort, P., & Brown, C.M. (1997). Semantic and syntactic ERP effects of listening to speech compared to reading. Submitted for publication.

Hagoort, P., Brown, C.M., & Groothusen, J. (1993) The Syntactic Positive Shift (SPS) as an ERP-measure of syntactic processing. *Language and Cognitive Processes* 8, 439–483.

Hagoort, P., Brown, C.M., & Swaab, T.S. (1996). Lexical-semantic event-related potential effects in patients with left hemisphere lesions and aphasia, and patients with right hemisphere lesions without aphasia. *Brain* 119, 627–649.

Holcomb, P.J. (1993). Semantic priming and stimulus degradation: Implications for the role of the N400 in language processing. *Psychophysiology* 30, 47–61.

Holcomb, P.J., & Neville, H.J. (1990). Semantic priming in visual and auditory lexical decision: A between modality comparison. *Language and Cognitive Processes* 5, 281–312.

Holmes, V.M. (1987). Syntactic parsing: In search of the garden path. In M. Coltheart (Ed.), *Attention and performance XII. The psychology of reading*. (pp. 587–599) London: Erlbaum.

King, J.W., & Kutas, M. (1995). Who did what and when? Using word- and clause-level ERPs to monitor working memory usage in reading. *Journal of Cognitive Neuroscience* 7, 376–395.

Kluender, R., & Kutas, M. (1993). The interaction of lexical and syntactic effects in the processing of unbounded dependencies. *Language and Cognitive Processes* 8, 573–633.

Kutas, M., & Hillyard, S.A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science* 207, 203–205.

Kutas, M., & Hillyard, S.A. (1983). Event-related brain potentials to grammatical errors and semantic anomalies. *Memory & Cognition* 11, 539–550.

Kutas, M., & King, J.W. (1995). The potentials for basic sentence processing: Differentiating integrative processes. In T. Inui & J. McClelland (Eds.), *Attention and performance XVI: Information integration in perception and communication*. Cambridge, MA: MIT Press.

Kutas, M., & Van Petten, C. (1988). Event-related brain potential studies of language. In P.K. Ackles, J.R. Jennings, & M.G.H. Coles (Eds.), *Advances in Psychophysiology, Volume 3*. Greenwich, CT: JAI Press.

Kutas, M., & Van Petten, C. (1995). Psycholinguistics electrified: Event-related brain potential investigations. In M. Gernsbacher (Ed.), *Handbook of psycholinguistics*. New York: Academic Press.

Kutas, M., Lindamood, T., & Hillyard, S.A. (1984). Word expectancy and event-related brain potentials during sentence processing. In S. Kornblum & J. Requin (Eds.), *Preparatory states and processes*. (pp. 217–238) Hillsdale, NJ: Erlbaum.

Kutas, M., Neville, H.J., & Holcomb, P.J. (1987). A preliminary comparison of the N400 response to semantic anomalies during reading, listening and signing. In R.J. Ellingson, N.M.F. Murray, & A.M. Halliday (Eds.), *The London Symposium (EEG Suppl. 39)*. Amsterdam: Elsevier.

MacDonald, M.A., Pearlmutter, N.J., & Seidenberg, M.S (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review* 101, 676–703.

McCallum, W.C., Farmer, S.F., & Pocock, P.K. (1984). The effects of physical and semantic incongruities on auditory event-related potentials. *Electroencephalography and Clinical Neurophysiology* 59, 447–488.

McClelland, J.L., St. John, M., & Taraban, R. (1989). Sentence comprehension: A parallel distributed processing approach. Language and Cognitive Processes, 4, 287–335.

Mecklinger, A., Schriefers, H., Steinhauer, K., & Friederici, A.D. (1995). Processing relative clauses varying on syntactic and semantic dimensions: An analysis with event-related potentials. *Memory & Cognition* 23, 477–494.

Münte, T.F., Heinze, H., & Mangun, G. (1993). Dissociation of brain activity related to syntactic and semantic aspects of language. *Journal of Cognitive Neuroscience* 5, 335–344.

Näätänen, R. (1992). *Attention and brain function*. Hillsdale, NJ: Erlbaum.

Neville, H.J., Coffey, S.A., Lawson, D.S., Fischer, A., Emmorey, K., & Bellugi, U. (1997). Neural systems mediating American Sign Language: Effects of sensory experience and age of acquisition. *Brain and Language* 57, 285–308.

Neville, H.J., Nicol, J.L., Barss, A., Forster, K.I., & Garrett, M.F. (1991). Syntactically based sentence processing classes: Evidence from event-related brain potentials. *Journal of Cognitive Neuroscience* 3, 151–165.

Nunez, P.L. (1981). *Electric fields of the brain*. New York: Oxford University Press.

Nunez, P.L. (1990). Physical principles and neurophysiological mechanisms underlying event-related potentials. In J.W. Rohrbaugh, R. Parasuramen, & R. Johnson (Eds.), *Event-related brain potentials*. (pp. 19–36) New York: Oxford University Press.

Osterhout, L., & Holcomb, P.J. (1992). Event-related brain potentials elicited by syntactic anomaly. *Journal of Memory and Language* 31, 785–806.

Osterhout, L., & Holcomb, P.J. (1993). Event-related potentials and syntactic anomaly: Evidence of anomaly detection during the perception of continuous speech. *Language and Cognitive Processes* 8, 413–437.

Osterhout, L., & Holcomb, P.J. (1995). Event-related potentials and language comprehension. In M.D. Rugg, & M.G.H. Coles (Eds.), *Electrophysiology of mind: Event-related brain potentials and cognition*. (pp. 171–215) New York: Oxford University Press.

Osterhout, L., & Mobley, L.A. (1995). Event-related brain potentials elicited by the failure to agree. *Journal of Memory and Language* 34, 739–773.

Osterhout, L., Holcomb, P.J., & Swinney, D.A. (1994). Brain potentials elicited by garden path sentences: Evidence of the application of verb information during parsing. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20, 786–803.

Picton, T.W., Lins, O.G., & Scherg, M. (1995). The recording and analysis of event-related potentials. In F. Boller & J. Grafman (Eds.), *The handbook of neuropsychology, Volume 10*. Amsterdam: Elsevier.

Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior* 22, 358–374.

Rayner, K., Garrod, S., & Perfetti, C.A. (1992). Discourse influences during parsing are delayed. *Cognition*, 45, 109–139.

Rugg, M.D. (1990). Event-related brain potentials dissociate repetition effects of high- and low-frequency words. *Memory & Cognition*, 18, 367–379.

Rugg, M.D., & Coles, M.G.H. (Eds.). (1995). *Electrophysiology of mind: Event-related brain potentials and cognition*. New York: Oxford University Press.

Spivey-Knowlton, M.J., & Sedivy, J.C. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition* 55, 227–267.

Tanenhaus, M.K., & Trueswell, J.C. (1995). Sentence comprehension. In J. Miller, & P. Eimas (Eds.), *Handbook of perception and cognition, Volume 11: Speech, language, and communication*. (pp. 217–262) New York: Academic Press.

Taraban, R., & McClelland, J.L. (1990). Parsing and comprehension: A multiple-constraint view. In D. Balota, G.B. Flores d'Arcais, & K. Rayner (Eds.), *Comprehension processes in reading*. (pp. 231–263) Hillsdale, NJ: Erlbaum.

Van Petten, C., & Kutas, M. (1991). Electrophysiological evidence for the flexibility of lexical processing. In G.B. Simpson (Ed.), *Understanding word and sentence*. (pp. 129–174). North-Holland: Elsevier Science.

Wood, C.C. (1987) Generators of event-related potentials. In A.M. Halliday, S.R. Butler, & R. Paul (Eds.), *A textbook of clinical neurophysiology*. (pp. 535–567) New York: Wiley.