Available online at www.sciencedirect.com
SCIENCE DIRECT®

NeuroImage

NeuroImage 20 (2003) S18–S29

www.elsevier.com/locate/ynimg

How the brain solves the binding problem for language: a neurocomputational model of syntactic processing

Peter Hagoort*

F.C. Donders Center for Cognitive Neuroimaging, P.O. Box 9101, NL-6500 HB Nijmegen, The Netherlands

Abstract

Syntax is one of the components in the architecture of language processing that allows the listener/reader to bind single-word information into a unified interpretation of multiword utterances. This paper discusses ERP effects that have been observed in relation to syntactic processing. The fact that these effects differ from the semantic N400 indicates that the brain honors the distinction between semantic and syntactic binding operations. Two models of syntactic processing attempt to account for syntax-related ERP effects. One type of model is serial, with a first phase that is purely syntactic in nature (syntax-first model). The other type of model is parallel and assumes that information immediately guides the interpretation process once it becomes available. This is referred to as the immediacy model. ERP evidence is presented in support of the latter model. Next, an explicit computational model is proposed to explain the ERP data. This Unification Model assumes that syntactic frames are stored in memory and retrieved on the basis of the spoken or written word form input. The syntactic frames associated with the individual lexical items are unified by a dynamic binding process into a structural representation that spans the whole utterance. On the basis of a meta-analysis of imaging studies on syntax, it is argued that the left posterior inferior frontal cortex is involved in binding syntactic frames together, whereas the left superior temporal cortex is involved in retrieval of the syntactic frames stored in memory. Lesion data that support the involvement of this left frontotemporal network in syntactic processing are discussed.

Introduction

Language comprehension is more than just the concatenation of individual lexical items that are retrieved from memory on the basis of spoken or written input. Higher order constraints at the phonological, semantic, and syntactic level are often necessary to bind lexical items together into a coherent overall interpretation of the utterance. Recent accounts of the human language system (Jackendoff, 1999, 2002; Levelt, 1999), therefore, assume a cognitive architecture, which consists of separate processing levels for conceptual/semantic information, orthographic/phonological information, and syntactic information. Based on this architecture, most current models of language processing agree that, in online sentence processing, different types of constraints are very quickly taken into consideration during speaking and listening/reading. Constraints on how words can be structurally combined operate alongside qualitatively distinct constraints on the combination of word meanings, on the grouping of words into phonological phrases, and on their referential binding into a discourse model. Together, these constraints solve the "binding problem" for language or, in other words, how speakers and writers, listeners and readers bind single-word information into multiword utterances and complex messages (Hagoort et al., 1999).

Despite fairly wide agreement on the types of constraints that are effective during the formulation and the interpretation of sentences, exactly how these are implemented in the overall sentence processing architecture is still a matter of considerable debate in psycholinguistics. One of the key issues is when and how the assignment of a syntactic structure to an incoming string of words and the semantic integration of single-word meanings interact during listening/reading. One view is that in sentence comprehension, the syntactic analysis is autonomous and initially not influenced by semantic variables (Frazier, 1987). Semantic integration can be influenced by syntactic analysis, but it does not contribute to the computation of syntactic structure. Some recent empirical evidence for this view is provided in a

---

* Fax: +31-24-3610652.
E-mail address: peter.hagoort@fcdonders.kun.nl

1053-8119/$ – see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2003.1796

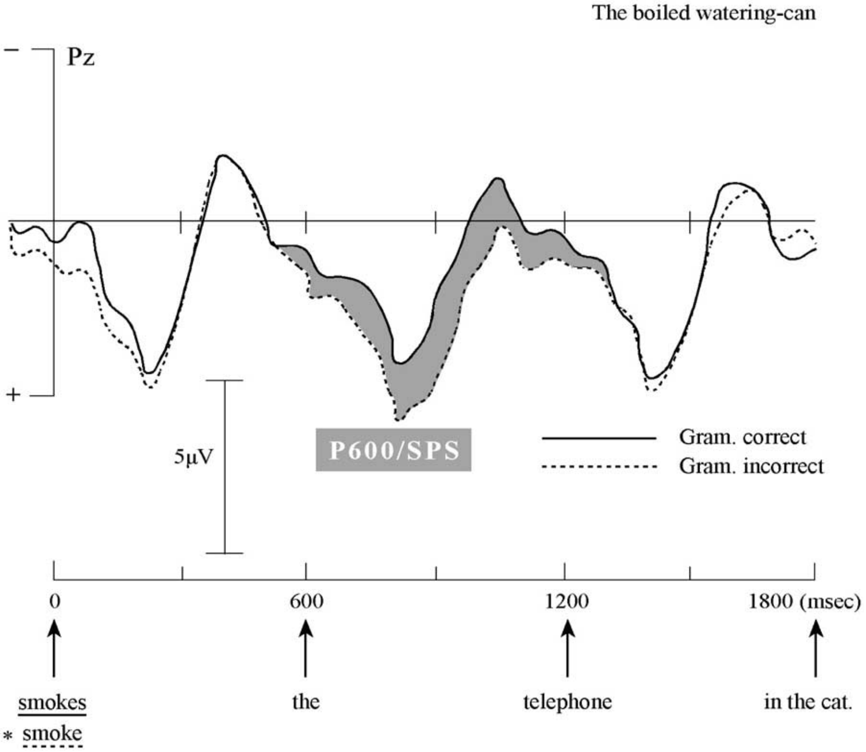

Fig. 1. ERPs to visually presented syntactic prose sentences. These are sentences without a coherent semantic interpretation. A P600/SPS is elicited by a violation of the required number agreement between the subject-noun phrase and the finite verb of the sentence. The averaged waveforms for the grammatically correct and the grammatically incorrect words are shown for electrode site Pz (parietal midline). The word that renders the sentence ungrammatical is presented at 0 ms on the time axis. The waveforms show the ERPs to this and the following two words. Words were presented word by word, with an interval of 600 ms. Negativity is plotted upward (adapted, with permission of the publisher, from Hagoort and Brown (1994), copyright 1994 Erlbaum).

study by O'Seaghdha (1997), in which naming latencies were measured. An alternative view maintains that lexical-semantic and discourse information can guide or contribute to the syntactic analysis of the utterance. This view is mainly supported by studies showing that the reading of syntactically ambiguous sentences is immediately influenced by lexical or more global semantic information (e.g., Altmann and Steedman, 1988; Trueswell et al., 1993, 1994; Tyler and Marslen-Wilson, 1977).

Some of the discrepancies between the different views on this topic are due to the fact that no clear distinction is made between cases in which the syntactic constraints are, at least temporarily, indeterminate with respect to the structural assignment (syntactic ambiguity) and cases in which these constraints are sufficient to determine the syntactic analysis. In the former case, there is a substantial body of evidence for an immediate influence of nonsyntactic context information on the structure that is assigned (Tanenhaus and Trueswell, 1995; Van Berkum et al., 1999b). However, for the latter case, although it has not been studied as intensely, the available evidence seems to provide support for a certain level of syntactic autonomy (Hagoort, 2003; O'Seaghdha, 1997).

A more recent version of the autonomous syntax view is proposed by Friederici (2002). Based on the time course of different language-relevant ERP effects, she proposes a three-phase model of sentence comprehension. The first phase is purely syntactic in nature. An initial syntactic structure is formed on the basis of information about the word category (noun, verb, etc.). During phase 2, lexical-semantic and morphosyntactic processes take place, which result in thematic role assignments (who is doing what to whom). In the third phase, integration of the different types of information takes place, and the final interpretation results. This proposal is mainly based on findings in ERP studies on language processing. The past decade has seen an increasing number of ERP studies on syntactic processing, triggered by the discovery some 10 years ago of an ERP effect to syntactic violations that was clearly different from the well-known N400 effect to semantic violations (Hagoort et al., 1993; Osterhout and Holcomb, 1992 see Fig. 1).

These studies have been followed up by a large number of ERP studies on syntactic processing that have provided a wealth of data. In light of this increasing amount of data, explicit models of sentence processing are needed that phrase the triggering conditions of different language-relevant ERP effects in terms of an explicit psycholinguistically motivated model. The model proposed by Friederici (2002) is the first attempt to do so. Here I will defend a different view, in which I will connect the known syntax-related ERP effects to a computational model of parsing (Vosse and Kempen, 2000) that has been able to account for a large portion of behavioral findings in the parsing literature and for deficit patterns in aphasic patients. In the remainder of this paper, I will first discuss the relevant ERP results and then present some data that are incompatible with a syntax-first model. Next, I will suggest an alternative model that can account for the ERP data more completely. Finally, I will indicate how the model connects to relevant brain areas for syntactic processing and to data from lesion studies.

Language-relevant ERP effects

The electrophysiology of language as a domain of study started with the discovery by Kutas and Hillyard (1980) of an ERP component that seemed especially sensitive to semantic manipulations. Their finding marked the beginning of a growing effort to find, interpret, and exploit language-relevant ERP components. Kutas and Hillyard observed a negative-going potential with an onset at about 250 ms and a peak around 400 ms (hence the N400), the amplitude of which was increased when the semantics of the eliciting word (e.g., *socks*) mismatched the semantics of the sentence context, as in *He spread his warm bread with socks.* Since 1980, much has been learned about the processing nature of the N400 (for extensive overviews, see Kutas and Van Petten, 1994; Osterhout and Holcomb, 1995). As Hagoort and Brown (1994) and many others have observed, the N400 effect does not depend on a semantic violation. Subtle differences in semantic expectancy such as between *mouth* and *pocket* in the sentence context *Jenny put the sweet in her mouth/pocket after the lesson* can modulate the N400 amplitude (see Fig. 2; Hagoort and Brown, 1994).

The amplitude of the N400 is most sensitive to the semantic relations between individual words or between words and their sentence and discourse context. The better the semantic fit between a word and its context, the more reduced the amplitude of the N400. Modulations of the N400 amplitude are generally viewed as directly or indirectly related to the processing costs of integrating the meaning of a word into the overall meaning representation that is built up on the basis of the preceding language input (Brown and Hagoort, 1993; Osterhout and Holcomb, 1992). This holds equally when the preceding language input consists of a single word, a sentence, or a discourse, indicating that semantic binding operations might be similar in word, sentence, and discourse contexts (Van Berkum et al., 1999b).

In recent years a number of ERP studies have been devoted to establishing ERP effects that can be related to the processing of syntactic information. These studies have found ERP effects to syntactic processing that are qualitatively different from the N400. Even though the generators of these effects are not yet well determined and not necessarily language specific (Osterhout and Hagoort, 1999), the existence of qualitatively distinct ERP effects to semantic and syntactic processing indicates that the brain honors the distinction between semantic and syntactic binding operations. Thus, the finding of qualitatively distinct ERP effects for semantic and syntactic processing operations supports the claim that these two levels of language processing are domain specific. However, domain specificity should not be confused with modularity (Fodor, 1983), which makes the much stronger claim that domain-specific levels of processing operate autonomously without interaction (informational encapsulation). Although domain specificity is widely assumed in models of language processing, there is much less agreement about the organization of the cross talk between different levels of sentence processing (cf. Boland and Cutler, 1996).

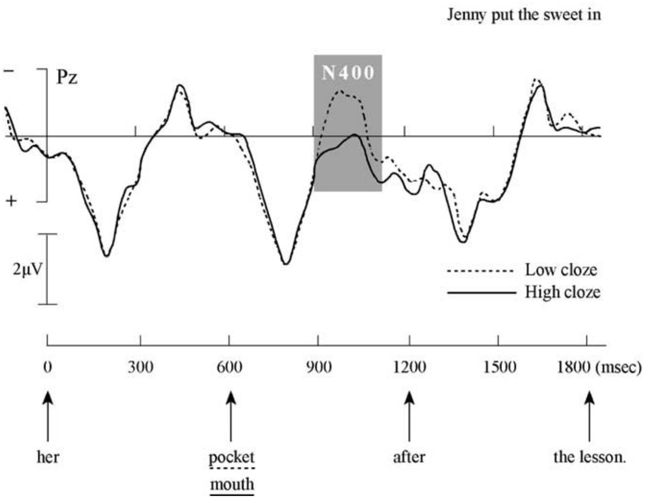

Fig. 2. Modulation of the N400 amplitude as a result of a manipulation of the semantic fit between a lexical item and its sentence context. The grand-average waveform is shown for electrode site Pz (parietal midline) for the best fitting word (High cloze) and a word that is less expected in the given sentence context (Low cloze). The sentences were visually presented word by word, every 600 ms. Here the critical word is preceded and followed by one word. The critical word is presented at 600 ms on the time axis. Negativity is up (adapted, with permission of the publisher, from Hagoort and Brown (1994), copyright 1994 Erlbaum).

ERP studies on syntactic processing have reported a number of ERP effects related to syntax (for an overview, see Hagoort et al., 1999). The two most salient syntax-related effects are an anterior negativity, also referred to as LAN, and a more posterior positivity, here referred to as P600/SPS.

LAN

A number of studies have reported negativities that are different from the N400, in that they usually show a more frontal maximum (but see Münte et al., 1997) and are sometimes larger over the left than the right hemisphere, although in many cases the distribution is bilateral (Hagoort et al., 2003a). Moreover, the conditions that elicit these frontal negative shifts seem to be more strongly related to syntactic processing than to semantic integration. Usually, LAN effects occur within the same latency range as the N400, that is, between 300 and 500 ms poststimulus (Friederici et al., 1996; Kluender and Kutas, 1993; Münte et al., 1993; Osterhout and Holcomb, 1992; Rösler et al., 1993). But in some cases the latency of a left-frontal negative effect is reported to be much earlier, somewhere between approximately 100 and 300 ms (Friederici, 2002; Friederici et al., 1993; Neville et al., 1991).

In some studies LAN effects in response to violations of word-category constraints have been reported (Friederici et al., 1996; Hagoort et al., 2003b; Münte et al., 1993). That is, if the syntactic context requires a word of a certain syntactic class (e.g., a noun in the context of a preceding article and adjective), but in fact a word of a different syntactic class is presented (e.g., a verb), early negativities are observed. Friederici and colleagues (e.g., Friederici, 1995; Friederici et al., 1996) have tied the early negativities specifically to the processing of word-category information. However, in other studies similar early negativities are observed with number, case, gender, and tense mismatches (Münte and Heinze, 1994; Münte et al., 1993). In these violations the word category is correct but the morphosyntactic features are wrong. Friederici (2002) has recently attributed the very early negativities that occur approximately between 100 and 300 ms (ELAN) to violations of word category and the negativities between 300 and 500 ms to morphosyntactic processing.

P600/SPS

A second ERP effect that has been related to syntactic processing is a later positivity, nowadays referred to as P600/SPS (Coulson et al., 1998; Hagoort et al., 1999; Osterhout et al., 1997). One of the antecedent conditions of P600/SPS effects is a violation of a syntactic constraint. If, for instance, the syntactic requirement of number agreement between the grammatical subject of a sentence and its finite verb is violated (see (1) below, with the critical verb form in italic; the * indicates the nongrammaticality of the sentence), a positive-going shift is elicited by the word that renders the sentence nongrammatical (Hagoort et al., 1993). This positive shift starts at about 500 ms after the onset of the violation and usually lasts for at least 500 ms. Given the polarity and the latency of its maximal amplitude, this effect was originally referred to as the P600 (Osterhout and Holcomb, 1993) or, on the basis of its functional characteristics, as the syntactic positive shift (SPS; Hagoort et al., 1993).

(1) *The spoiled child *throw* the toy on the ground.

An argument for the independence of this effect from possibly confounding semantic factors is that it also occurs in sentences in which the usual semantic/pragmatic constraints have been removed (Hagoort and Brown, 1994). This results in sentences like (2a) and (2b), of which one is semantically odd but grammatically correct, whereas the other contains the same agreement violation as in (1):

(2a) The boiled watering-can *smokes* the telephone in the cat.
(2b) *The boiled watering-can *smoke* the telephone in the cat.

If one compares the ERPs to the italicized verbs in (2a) and (2b), a P600/SPS effect in response to the nongrammatical verb form is visible (see Fig. 1). Despite the fact that these sentences do not convey any conventional meaning, the ERP effect of the violation demonstrates that the language system is nevertheless able to parse the sentence into its constituent parts.

Similar P600/SPS effects have been reported for a broad range of syntactic violations in different languages (English, Dutch, German), including phrase-structure violations (Hagoort et al., 1993; Neville et al., 1991; Osterhout and Holcomb, 1992); subcategorization violations (Ainsworth-Darnell et al., 1998; Osterhout et al., 1994, 1997b); violations in the agreement of number, gender, and case (Coulson et al., 1998; Hagoort et al., 1993; Münte et al., 1997; Osterhout, 1997; Osterhout and Mobley, 1995); violations of subjacency (McKinnon and Osterhout, 1996; Neville et al., 1991); and violations of the empty-category principle (McKinnon and Osterhout, 1996). Moreover, they have been found with both written and spoken input (Friederici et al., 1993; Hagoort and Brown, 2000; Osterhout and Holcomb, 1993).

In summary, two classes of syntax-related ERP effects have been consistently reported. These two classes differ in their polarity, topographic distribution, and latency characteristics. In terms of latency, the first class of effects is an anterior negativity, which is seen only in response to syntactic violations. In a later latency range, positive shifts occur which are elicited not only by syntactic violations, but also in grammatically well-formed sentences that vary in complexity (Kaan et al., 2000) or as a function of the number of alternative syntactic structures that are compatible with the input at a particular position in the sentence (syntactic ambiguity) (Osterhout et al., 1994; Van Berkum et al., 1999a). Since these two classes of effects are now well established in the context of language processing, and are clearly different from the N400 effect, the need arises to account for these effects in terms of a well-defined model of language processing.

Models of sentence processing can be divided into two types. One type of model is serial in nature with a precedence of syntactic information. That is, an initial syntactic structure is constructed before other information (e.g., lexical-semantic, discourse information) is taken into account (Frazier, 1987). I will refer to this type of model as a syntax-first model. The alternative broad set of models claims that the different information types (lexical, syntactic, phonological, pragmatic) are processed in parallel and influence the interpretation process incrementally, that is, as soon as the relevant pieces of information are available (Marslen-Wilson, 1989). I will refer to this type of model as the immediacy model. Overall, the behavioral data, although not decisive, are more in favor of the second type of model than the first. I will first present some recent ERP data that are more compatible with the immediacy model.

Evidence against the syntax-first principle

The strong version of a serial syntax-first model of sentence processing assumes that the computation of an initial syntactic structure precedes semantic binding operations, because structural information is necessary as input for thematic role assignment. In other words, if no syntactic structure can be built up, semantic binding will be impaired. Recent electrophysiological evidence has been taken as evidence for this syntax-first principle (Friederici, 2002). Alternative models (Marslen-Wilson and Tyler, 1980; MacDonald et al., 1994) claim that semantic and syntactic information is processed in parallel and immediately used when it becomes available.

Friederici (2002) recently proposed a model for sentence comprehension that is largely based on the available ERP data. This model assumes a first autonomous phase in which a preliminary syntactic structure is formed on the basis of word category information. This latter information is purely syntactic in nature and part of a word's lexical specification in long-term memory. The empirical evidence for this autonomous first phase in sentence processing is derived from a series of studies in which Friederici and colleagues found an ELAN to auditorily presented words with a prefix indicative of a word category violation. For instance, Hahne and Jescheniak (2001) and Friederici et al. (1993) had their subjects listen to sentences such as "Die Birne wurde im *gepflückt*" ("The pear was being in-the *plucked*) or "Die Freund wurde im *besucht*" ("The friend was being in-the *visited*"), in which the prefixes "ge-" and "be-" in combination with the preceding auxiliary "wurde" indicate a past participle, while the preposition "im" requires a noun. In this case a very early (between 100 and 300 ms) left anterior negativity that preceded the N400 effect was observed.

Although this evidence is compatible with a syntax-first model, it is not necessarily incompatible with a parallel, interactive model of sentence processing. As long as word category information can be derived from the acoustic input earlier than semantic information, as was the case in the above-mentioned studies, the immediacy model predicts that it will be used as it comes in. The syntax-first model, however, predicts that even in cases in which word category information comes in later than semantic information, this syntactic information will nevertheless be used earlier than semantic information in sentence processing. We designed a strong test of the syntax-first model, in which semantic information precedes word category information. In many languages information about the word category is often encapsulated in the suffix rather than the prefix of a word. In contrast to parallel models, a syntax-first model would in such a case predict that semantic processing (more in particular, semantic binding) is postponed until after the information about the word category has become available.

Fig. 3. A waveform of an acoustic token of the Dutch verb form "kliederde" ("messed"). The suffix "-de" indicates past tense. The total duration of the acoustic token is approximately 450 ms. The onset of the suffix "-de" is at approximately 300 ms. Only after 300 ms of signal can the acoustic token be classified as a verb. Thus, in a context that does not allow a verb in that position, the category violation point is at 300 ms into the verb (see text).

In a recent study Van den Brink and Hagoort (2003) compared correct Dutch sentences (see (3a)) with their anomalous counterparts (see (3b)) in which the critical word (italicized in (3a/b)) was both a semantic violation in the context and of the incorrect word category. However, in contrast to the experiments by Friederici and colleagues, word category information was encoded in the suffix "-de".

(3a) Het vrouwtje veegde de vloer met een oude *bezem* gemaakt van twijgen (The woman wiped the floor with an old *broom* made of twigs).

(3b) *Het vrouwtje veegde de vloer met een oude *kliederde* gemaakt van twijgen (The woman wiped the floor with an old *messed* made of twigs).

Fig. 3 shows the waveform of the spoken verb form "kliederde" ("messed"). This verb form has a duration of approximately 450 ms. The stem already contains part of the semantic information. However, the onset of the suffix "-de" is at about 300 ms into the word. Only at this point will it be clear that the word category is a verb and not a noun as required by the context. We define this moment of deviation from the correct word category as the category violation point (CVP), because only at this time is information provided on the basis of which it can be recognized as a verb, which is the incorrect word category in the syntactic context. Although in this case semantic information can be extracted from the spoken signal before word category information, the syntax-first model predicts that this semantic information cannot be used for semantic binding until after the assignment of word category.

Fig. 4 shows the averaged waveforms that are time locked to the CVP for two frontal sites at which usually the ELAN is observed and two posterior sites that are representative of N400 effects. As can be seen, the N400 effect clearly precedes the ELAN in time. Whereas the ELAN started at approximately 100 ms after the CVP, the N400 effect was already significant before the CVP. To my knowledge, this is the clearest evidence so far for the claim that semantic binding can start before word category infor-

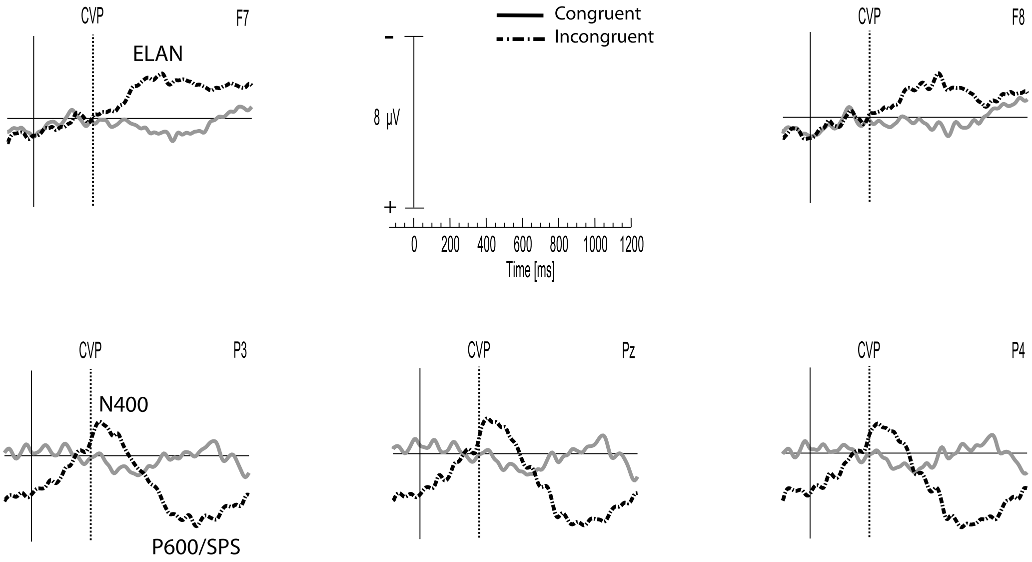

Fig. 4. Connected speech. Grand average ERPs from two frontal electrode sites (F7, F8) and three posterior electrode sites (Pz, P3, P4) to critical words that were semantically and syntactically congruent with the sentence context or semantically and syntactically incongruent. Grand average waveforms were computed after time locking on a trial-by-trial basis to the moment of word category violation (CVP). The baseline was determined by averaging in the 180 to 330 ms interval, corresponding to a 150-ms interval preceding the CVP in the incongruent condition. The time axis is in milliseconds, negativity is up. The ELAN is visible over the two frontal sites, the N400 and the P600/SPS are visible over the three posterior sites. The onset of the ELAN is at 100 ms after the CVP; the onset of the N400 effect precedes the CVP by approximately 10 ms (Van den Brink and Hagoort, 2003).

mation is provided. This is strong evidence for the immediacy assumption: information available in the signal is immediately used for further processing. In contrast to what a strong version of the syntax-first model predicts, semantic binding does not need to wait until an initial structure is built on the basis of word category information.

In serial models of sentence processing, the syntactic structure that gets built up incrementally as words come in sometimes needs to be revised. For instance, this is the case if later information in the sentence results in a garden-path or if semantic information is incompatible with the structure assigned on the basis of word category information. In the context of a serial syntax-first model, the P600/SPS is related to processes of reanalysis and repair (Friederici, 2002). The straightforward prediction of this account of the P600/SPS is that the harder reanalysis and repair are, the later the P600/SPS and/or the larger its amplitude. This prediction was tested in an experiment by Hagoort, Commandeur, and Kempen (manuscript in preparation). In this experiment we manipulated the strength of the syntactic repair options. The word category violation was either easily repaired or difficult to repair. An example of the materials is given below (the critical words are in italic):

(4a) Correct: De houthakker ontweek de *schroef* op dinsdag (The lumberjack dodged the *propellor* on Tuesday).

(4b) Weak violation: *De houthakker ontweek de *schroeft* op dinsdag (The lumberjack dodged the *propelled* on Tuesday).

(4c) Strong violation: *De houthakker ontweek de *omdat* op dinsdag (The lumberjack dodged the *because* on Tuesday).

Under the assumption that repair and revision are harder for strong than for relatively weak syntactic violations, the two-stage syntactic processing account of a syntax-first model predicts a later/larger P600/SPS for the word category violation in sentence (4c) than in sentence (4b). However, as can be seen in Fig. 5, the opposite was found. The most striking finding is that the stronger syntactic violation resulted in a substantially earlier P600/SPS than the weaker violation. The difference in onset latency between strong and weak violations is approximately 250 ms. This finding goes against the claim that the P600/SPS reflects mainly processes of reanalysis and repair, which are stages in a serial model, but have no place in a parallel model.

In summary, the evidence so far indicates that distinct ERP effects are observed for semantic integration (N400) and syntactic analysis ((E)LAN, P600/SPS). The ERP data presented are evidence against a syntax-first model of sentence processing. Rather, as soon as semantic or syntactic information is available, it is used for the purpose of interpretation. This is in line with the assumptions of the immediacy model. The triggering conditions of the syntax-related ERP effects are becoming clearer. So far, (E)LAN effects have been seen only in response to syntactic *violations.* These violations can be word category violations that are sometimes seen early (ELAN), but they can also be morphosyntactic violations that are usually observed within the same time frame as the N400 effects (300–500 ms). The

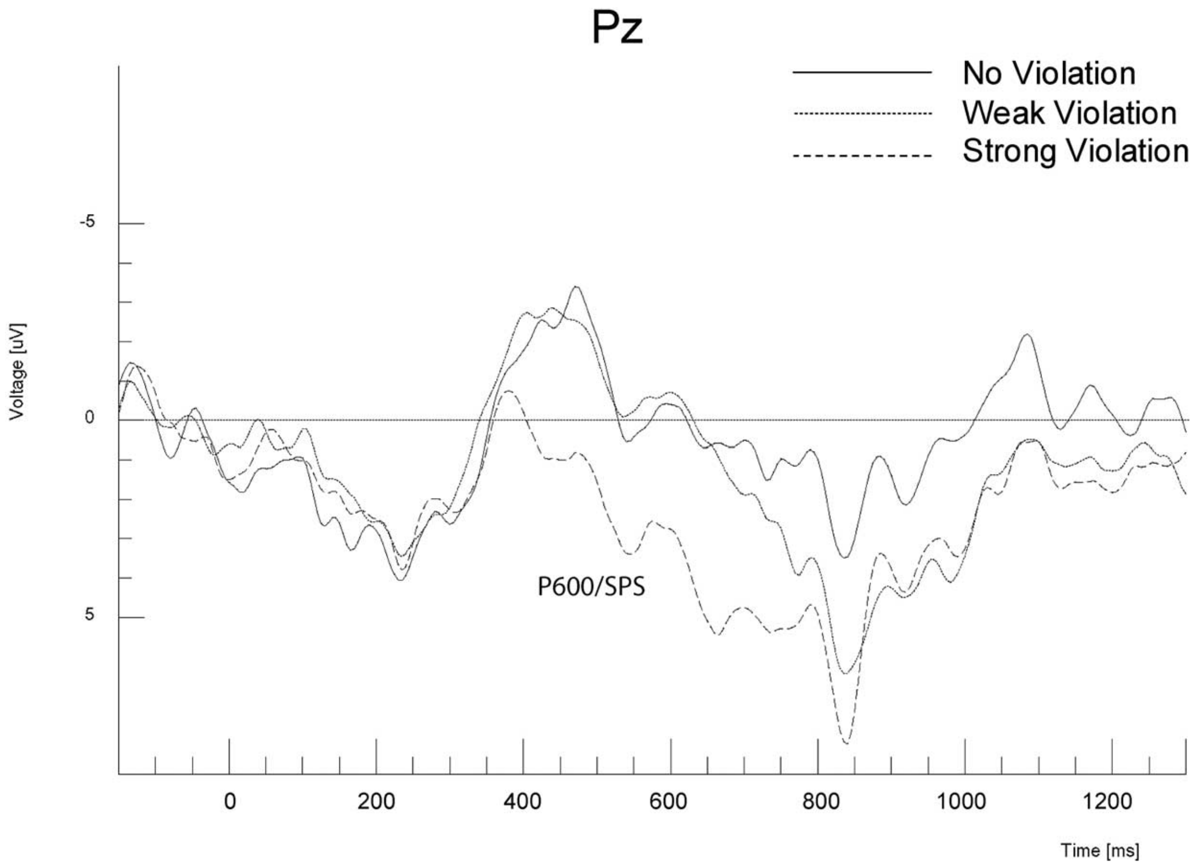

Fig. 5. The averaged waveform for a representative posterior electrode site (Pz), for grammatically well-formed sentences, sentences with a weak word category violation, and sentences with a strong word category violation. Onset of the critical words is at 0 ms. Negativity is up. The P600/SPS for the strong violation precedes the P600/SPS latency for the weak violation by about 250 ms.

anterior negativities are normally followed by a P600/SPS. In contrast to the (E)LAN, the P600/SPS is seen not only to syntactic violations, but also to syntactically less preferred structures (i.e., in the case of syntactic ambiguity; Van Berkum et al. 1999a; Osterhout et al., 1994) and to syntactically more complex sentences (Kaan et al., 2000). In many cases, the P600/SPS occurs without a concomitant early negativity. For straightforward syntactic violations the distribution of the P600/SPS seems to be more posterior than the P600/SPS reported in relation to syntactic ambiguity resolution and syntactic complexity (Hagoort et al., 1999; Kaan and Swaab, 2003).

In the rest of this paper I will discuss how a computational model of parsing that accounts for a substantial amount of behavioral data and data from aphasic patients can also handle the syntax-related ERP effects in quite a natural way. Before the model can be discussed, however, the issue of the language specificity of LAN and P600/SPS needs to be clarified. Although both the (left) anterior negativity and the P600/SPS seem to be syntax-related ERP effects within the domain of language processing, one cannot claim language specificity for either of these effects. Most likely they are not language specific. However, this does not have serious consequences for the processing architecture of sentence processing. As long as language input feeds into generator ensembles of different ERPs in accordance with different levels of processing within the cognitive architecture, it can be inferred that the brain honors the distinction between these processing levels.

The Unification Model

I will propose an explicit account of syntax-related ERP effects based on a computational model of parsing developed by Vosse and Kempen (2000), here referred to as the Unification Model. I will first describe the general architecture of this model. Following the model account of the ERP data, I will indicate how imaging and lesion data provide information about the neural implementation of the different components of the model.

According to the Unification Model each word form in the lexicon is associated with a structural frame. This structural frame consists of a three-tiered unordered tree, specifying the possible structural environment of the particular lexical item (see Fig. 6; for details concerning the computation of word order, see Harbusch and Kempen, 2002).

The top layer of the frame consists of a single phrasal

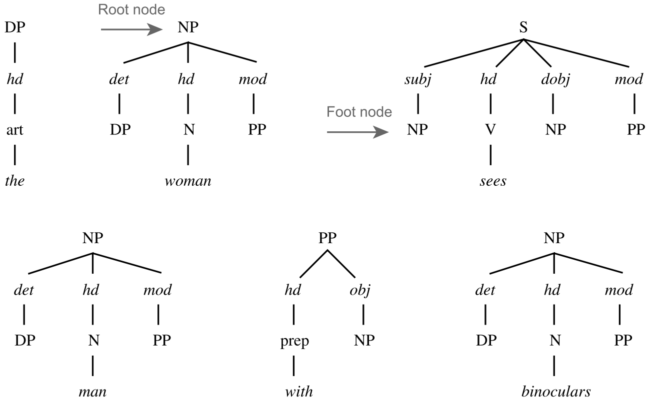

Fig. 6. Syntactic frames in memory. These frames are retrieved on the basis of incoming word form information. DP, determiner phrase; NP, noun phrase; S, sentence; PP, prepositional phrase; art, article; hd, head; det, determiner; mod, modifier; subj, subject; dobj, direct object.

node (e.g., NP). This so-called root node is connected to one or more functional nodes (e.g., subject, head, direct object) in the second layer of the frame. The third layer contains again phrasal nodes to which lexical items or other frames can be attached.

This parsing account is "lexicalist" in the sense that all syntactic nodes (e.g., S, NP, VP, N, V) are retrieved from the mental lexicon. In other words, chunks of syntactic structure are stored in memory. There are no syntactic rules that introduce additional nodes. In the online comprehension process, structural frames associated with the individual word forms incrementally enter the unification work space. In this work space constituent structures spanning the whole utterance are formed by a unification operation. This operation consists of linking up lexical frames with identical root and foot nodes (see Fig. 7) and checking agreement features (number, gender, person, etc.).

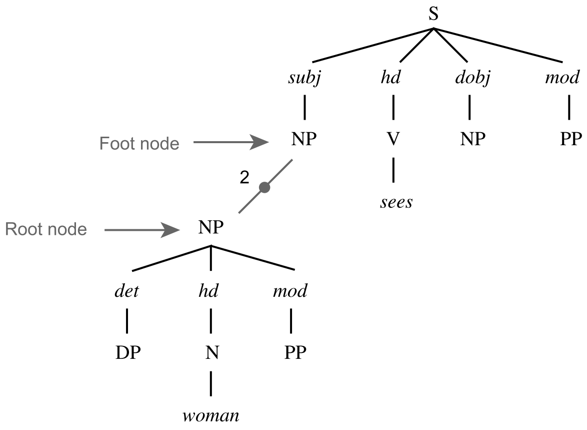

Fig. 7. The unification operation of two lexically specified syntactic frames. The unification takes place by linking the root node NP to an available foot node of the same category. The number 2 indicates that this is the second link that is formed during online processing of the sentence "The woman sees the man with the binoculars."

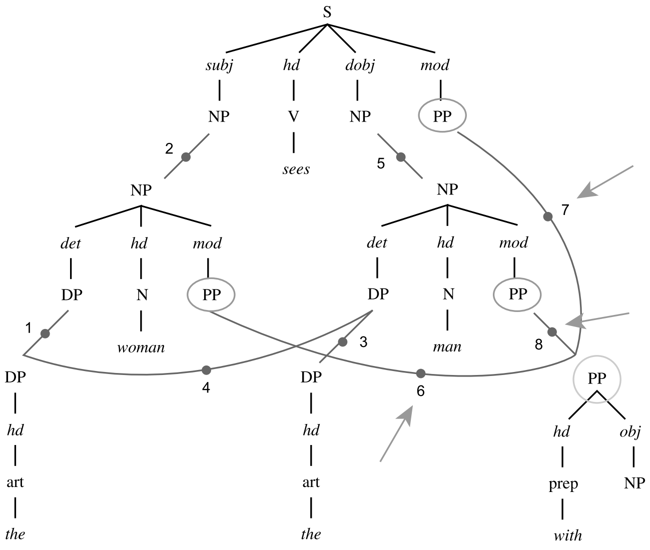

Fig. 8. Lateral inhibition between three different PP foot nodes that are candidate unification sites for the PP root node of the preposition "with." The three possible unification links are indicated by arrows. Lateral inhibition between these three possible unifications (6, 7, and 8) ultimately results in one unification that wins the competition and remains active.

The resulting unification links between lexical frames are formed dynamically, which implies that the strength of the unification links varies over time until a state of equilibrium is reached. Due to the inherent ambiguity in natural language, alternative binding candidates will usually be available at any point in the parsing process. That is, a particular root node (e.g., PP) often finds more than one matching foot node (i.e., PP) with which it can form a unification link (see Fig. 8).

Ultimately, one phrasal configuration results. This requires that among the alternative binding candidates only one remains active. The required state of equilibrium is reached through a process of lateral inhibition between two or more alternative unification links. In general, due to gradual decay of activation more recent foot nodes will have a higher level of activation than the ones that entered the unification space earlier. This is why the likelihood of an attachment of the PP into the syntactic frame of the verb "sees" is higher than into the syntactic frame for "woman" (see Fig. 8). In addition, strength levels of the unification links can vary in function of plausibility (semantic!) effects. For instance, if instrumental modifiers under S nodes have a slightly higher default activation than instrumental modifiers under an NP node, lateral inhibition can result in overriding the recency effect. For our example sentence (see Fig. 8) it means that the outcome of lateral inhibition is that the PP may be linked to the S frame (unification (U) link 7) rather than to the more recent NP node of "man" (U link 8) (for details, see Vosse and Kempen, 2000).

The Unification Model accounts for sentence complexity effects known from behavioral measures, such as reading times. In general, sentences are harder to analyze syntacti-

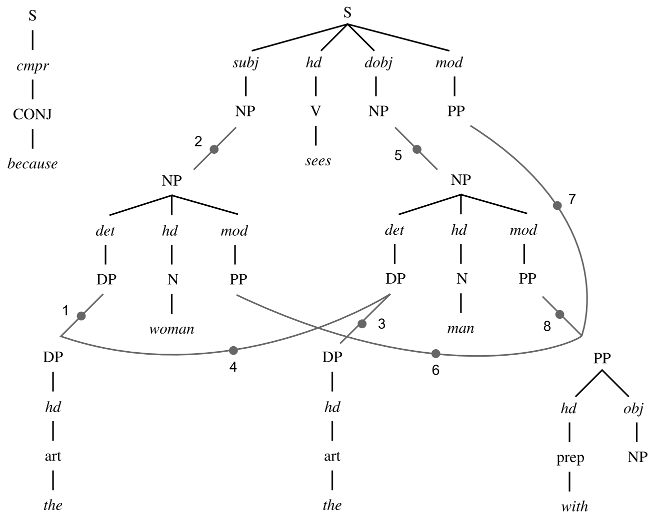

Fig. 9. A dangling syntactic frame for the conjunction element "because." This syntactic frame cannot be attached into the phrasal configuration for the remaining parts of sentence.

cally when more potential unification links of similar strength enter into competition with each other. Sentences are easy when the number of U links is small and of unequal strength.

The advantages of the Unification Model are that it (i) is computationally explicit, (ii) accounts for a large series of empirical findings in the parsing literature and in the neuropsychological literature on aphasia, and (iii) belongs to the class of lexicalist parsing models that have found increasing support in recent years (Bresnan, 2001; MacDonald et al., 1994).

This model also nicely accounts for the two classes of syntax-related ERP effects reported in this and many other studies. The P600/SPS is reported in relation to syntactic violations, syntactic ambiguities, and syntactic complexity. The (E)LAN, in contrast, has so far been observed only in response to syntactic violations. In the Unification Model, binding (unification) is prevented in two cases. One case is when the root node of a syntactic building block (e.g., NP) does not find another syntactic building block with an identical foot node (i.c., NP) to bind to. The other case is when the agreement check finds a serious mismatch in the grammatical feature specifications of the root and foot nodes. The claim is that a (left) anterior negativity results from a failure to bind, as a result of a negative outcome of the agreement check or a failure to find a matching category node. For instance, the sentence "The woman sees the man because with the binoculars" does not result in a completed parse, since the syntactic frame associated with "because" does not find unoccupied (embedded) S-root nodes that it can bind to. As a result, unification fails (Fig. 9).

In the context of the Unification Model, I propose that the P600/SPS is related to the time it takes to establish unification links of sufficient strength. The time it takes to build up the unification links until the required strength is reached is affected by ongoing competition between alternative unification options (syntactic ambiguity), by syntactic complexity, and by semantic influences. The amplitude of the P600/SPS is modulated by the amount of competition. Competition is reduced when the number of alternative binding options is smaller or when lexical, semantic, or discourse context biases the strengths of the unification links in a particular direction, thereby shortening the duration of the competition. Violations result in a P600/SPS as long as unification attempts are made. For instance, a mismatch in gender or agreement features might still result in weaker binding in the absence of alternative options. However, in such cases the strength and buildup of U links will be affected by the partial mismatch in syntactic feature specification. Compared to less complex or syntactically unambiguous sentences, in more complex and syntactically ambiguous sentences it takes longer to build up U links of sufficient strength. The latter sentences, therefore, result in a P600/SPS in comparison to the former ones.

In summary, it seems that the Unification Model provides an acceptable account for the collective body of ERP data on syntactic processing. Moreover, it does not assume a syntax-first architecture. It is, therefore, a better account of the empirical data, both behavioral and electrophysiological, than models that are serial and assume a syntax-first phase (e.g., Friederici, 2002).

The neural implementation of the Unification Model

In this section I will make some suggestions about how the different components of the Unification Model might be neurally instantiated. This proposal is not yet explicitly tested, but as I will argue, it makes good sense in light of our current knowledge about the contributions of the areas involved. In a recent meta-analysis of 28 neuroimaging studies, Indefrey (2003) found two areas that were critical for syntactic processing independent of the input modality (visual in reading, auditory in speech). These two supramodal areas for syntactic processing were the left posterior superior temporal gyrus and the left posterior inferior frontal cortex, substantially overlapping with left prefrontal cortex (see Fig. 10).

As is known from lesion studies in aphasic patients, lesions in different areas of the left perisylvian cortex can result in syntactic processing deficits in sentence comprehension (Caplan et al., 1996). The idea that modality-independent grammatical knowledge was mainly represented in Broca's area (Zurif, 1998) has thus been proven to be incorrect. At the same time, the left posterior temporal cortex is known to be involved in lexical processing (Indefrey and Cutler, 2003). In connection with the Unification Model, this part of the brain might be important for the retrieval of the syntactic frames that are stored in the lexicon.

The unification space where individual frames are connected into a phrasal configuration for the whole utterance might be localized in the left frontal part of the syntax-relevant network of brain areas. One of the main specializations of the prefrontal cortex is the holding online and binding of information (Mesulam, 2002). It might be the right area for providing the computational resources for binding together the lexical-syntactic frames through the dynamics of creating unification links between them (cf. Duncan and Miller, 2002). It thus seems that the components of the Unification Model and the areas known to be crucial for syntactic processing can be connected in a relatively natural way, with the left superior temporal cortex relevant for storage and retrieval of syntactic frames, and the left prefrontal cortex important for binding these frames together.

Lesioning syntactic binding

Further support for the neural implementation of the model comes from agrammatic aphasics with a lesion in the left frontal and left temporal areas involved in syntactic processing. In a recent study (Hagoort et al., 2003a), we found that these patients scored at chance level in a syntactic comprehension test for structures that contained a relative clause. Moreover, in contrast to elderly controls and Broca's aphasics with an above-chance performance on the syntactic comprehension test, the agrammatic aphasics failed to show any of the syntax-relevant ERP effects in sentences that violated the agreement between the subject of the sentence and one of the finite verbs in the sentence (e.g., "The girls pay the baker and *takes* the bread home"; the violation is italicized). As can be seen in Fig. 11, for a representative electrode site over the left temporal cortex, neither a LAN nor a P600/SPS is observed in the agrammatic aphasics. These patients were severely impaired in exploiting the morphosyntactic markers on verbs and nouns to unify the individual lexical items into an phrasal configuration that spans the whole utterance.

Earlier lesion studies in patients with syntactic deficits (e.g., Caplan et al., 1996) have shown that both anterior and posterior lesions can result in the same syntactic comprehension problems. This supports the idea that syntactic processing is not restricted to Broca's area, but requires the recruitment of a distributed network, in which both left frontal and left temporal areas are important nodes.

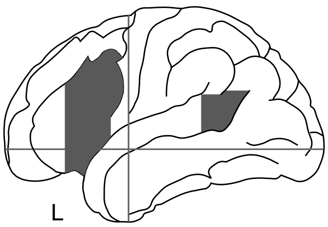

Fig. 10. Common areas of activation in a meta-analysis of 28 imaging studies on the processing of syntactic information during language comprehension. Areas were restricted to the temporal and frontal lobes of the left hemisphere (after Indefrey, 2003 reproduced by permission of the publisher).

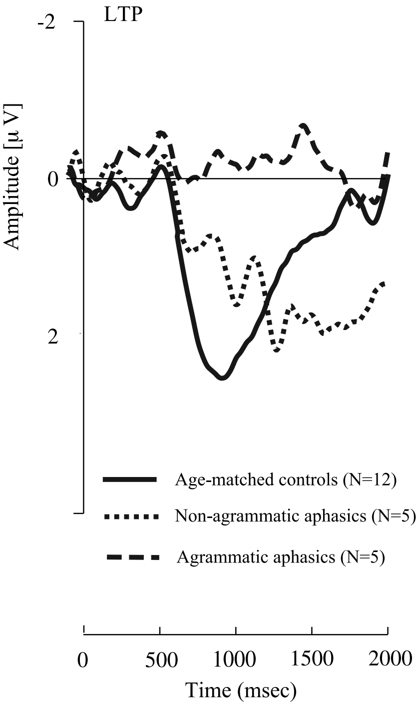

Fig. 11. Difference waveform (obtained by subtracting the waveforms for the correct from those for the incorrect condition) at a left temporal electrode site. Waveforms are averaged over participants for the group of age-matched controls ($n = 12$), nonagrammatic aphasics ($n = 5$), and agrammatic aphasics ($n = 5$). The waveforms are aligned at zero on the onset of the word that violates number agreement (after Hagoort et al., 2003b reproduced by permission of the publisher).

Acknowledgments

This research was supported by Grant 400-56-384 from the Netherlands Organization for Scientific Research. I thank Michael Coles, Peter Indefrey, Gerard Kempen, and two reviewers for their helpful comments on an earlier version of this paper.

References

Ainsworth-Darnell, K., Shulman, H., Boland, J., 1998. Dissociating brain responses to syntactic and semantic anomalies: evidence from event-related potentials. J. Memory Lang. 38, 112–130.

Altmann, G.T.M., Steedman, M., 1988. Interaction with context during human sentence processing. Cognition 30, 191–238.

Boland, J.E., Cutler, A., 1996. Interaction with autonomy: multiple output models and the inadequacy of the great divide. Cognition 58, 309–320.

Bresnan, J.W., 2001. Lexical–functional syntax. Blackwell, Oxford.

Brown, C., Hagoort, P., 1993. The processing nature of the N400: evidence from masked priming. J. Cognit. Neurosci. 5, 34–44.

Caplan, D., Hildebrandt, N., Makris, N., 1996. Location of lesions in stroke patients with deficits in syntactic processing in sentence comprehension. Brain 119, 933–949.

Coulson, S., King, J.W., Kutas, M., 1998. Expect the unexpected: event-related brain response to morphosyntactic violations. Lang. Cognit. Processes 13, 21–58.

Duncan, J., Miller, E.K., 2002. Cognitive focus through adaptive neural coding in the primate prefrontal cortex, in: Stuss, D.T., Knight, R.T. (Eds.), Principles of frontal lobe function. Oxford Univ. Press, Oxford.

Fodor, J.D., 1983. The modularity of mind. MIT Press, Cambridge, MA.

Frazier, L., 1987. Sentence processing: A tutorial review, in: Coltheart, M (Ed.), Attention and performance, Vol. XII. Erlbaum, London, pp. 559–585.

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

Friederici, A.D., 2002. Towards a neural basis of auditory sentence processing. Trends Cognit. Sci. 6, 78–84.

Friederici, A.D., Hahne, A., Mecklinger, A., 1996. Temporal structure of syntactic parsing: early and late event-related brain potential effects. J. Exp. Psychol. Learn. Memory Cognit. 22, 1219–1248.

Friederici, A.D., Pfeifer, E., Hahne, A., 1993. Event-related brain potentials during natural speech processing: effects of semantic, morphological and syntactic violations. Cognit. Brain Res. 1, 183–192.

Hagoort, P., 2003. The interplay between syntax and semantics during sentence comprehension: ERP effects of combining syntactic and semantic violations. J. Cognit. Neurosci. 15, 883–899.

Hagoort, P., Brown, C., 1994. Brain responses to lexical ambiguity resolution and parsing, in: Clifton Jr.C., Frazier, L., Rayner, K. (Eds.), Perspectives on sentence processing. Erlbaum, Hillsdale, NJ, pp. 45–81.

Hagoort, P., Brown, C.M., Groothusen, J., 1993. The syntactic positive shift (SPS) as an ERP measure of syntactic processing. Lang. Cognit. Processes 8, 439–483.

Hagoort, P., Brown, C., Osterhout, L., 1999. The neurocognition of syntactic processing, in: Brown, C.M., Hagoort, P. (Eds.), Neurocognition of language. Oxford Univ. Press, Oxford, pp. 273–317.

Hagoort, P., Wassenaar, M., Brown, C.M., 2003a. Syntax-related ERP-effects in Dutch. Cognit. Brain Res. 16, 38–50.

Hagoort, P., Wassenaar, M., Brown, C.M., 2003b. Real-time semantic compensation in patients with agrammatic comprehension: electrophysiological evidence for multiple-route plasticity. Proc. Natl. Acad. Sci. USA 100, 4340–4345.

Hahne, A., Jescheniak, J.D., 2001. What's left if the Jabberwock gets the semantics? An ERP investigation into semantic and syntactic processes during auditory sentence comprehension. Cognit. Brain Res. 11, 199–212.

Harbusch, K., Kempen, G., 2002. A quantitative model of word order and movement in English, Dutch and German complement constructions, in: Proceedings of the 19th International Conference on Computational Linguistics (COLING-2002). Morgan Kaufmann, San Francisco, pp. 328–334.

Indefrey, P. 2003. Hirnaktivierungen bei syntaktischer Sprachverarbeitung: eine Meta-Analyse. Submitted for publication.

Indefrey, P., Cutler, A., 2003. Prelexical and lexical processing in listening, in: Gazzaniga, M.S. (Ed.), The new cognitive neurosciences III. MIT Press, Cambridge, MA (in press).

Jackendoff, R., 1999. The representational structures of the language faculty and their interactions, in: Brown, C.M., Hagoort, P. (Eds.), The neurocognition of language. Oxford Univ. Press, Oxford, pp. 37–79.

Jackendoff, R., 2002. Foundations of language: brain, meaning, grammar, evolution. Oxford Univ. Press, Oxford.

Kaan, E., Harris, A., Gibson, E., Holcomb, P., 2000. The P600 as an index of syntactic integration difficulty. Lang. Cognit. Processes 15, 159–201.

Kaan, E., Swaab, T.Y., 2003. Repair, revision, and complexity in syntactic analysis: an electrophysiological differentiation. J. Cognit. Neurosci. 15, 98–110.

Kluender, R., Kutas, M., 1993. Subjacency as a processing phenomenon. Lang. Cognit. Processes 8, 573–633.

Kutas, M., Hillyard, S.A., 1980. Reading senseless sentences: brain potentials reflect semantic anomaly. Science 207, 203–205.

Kutas, M., Van Petten, C.K., 1994. Psycholinguistics electrified: event-related brain potential investigations, in: Gernsbacher, M.A. (Ed.), Handbook of psycholinguistics. Academic Press, San Diego, pp. 83–143.

Levelt, W.J.M., 1999. Producing spoken language: a blueprint of the speaker. in: Brown, C.M., Hagoort, P. (Eds.), The neurocognition of language. Oxford Univ. Press, Oxford, pp. 83–122.

MacDonald, M.C., Pearlmutter, N.J., Seidenberg, M.S., 1994. Lexical nature of syntactic ambiguity resolution. Psychol. Rev. 101, 676–703.

Marslen-Wilson, W., 1989. Access and integration: projecting sound onto meaning, in: Marslen-Wilson, W. (Ed.), Lexical representation and process. MIT Press, Cambridge, MA, pp. 3–24.

Marslen-Wilson, W., Tyler, L.K., 1980. The temporal structure of spoken language understanding. Cognition 8, 1–71.

McKinnon, R., Osterhout, L., 1996. Constraints on movement phenomena in sentence processing: evidence from event-related brain potentials. Lang. Cognit. Processes 11, 495–523.

Mesulam, M.-M., 2002. The human frontal lobes: transcending the default mode through contingent encoding, in: Stuss, D.T., Knight, R.T. (Eds.), Principles of frontal lobe function. Oxford Univ. Press, Oxford.

Münte, T.F., Heinze, H.J., 1994. ERP negativities during syntactic processing of written words, in: Heinze, H.J., Münte, T.F., Mangun, G.R. (Eds.), Cognitive electrophysiology. Birkhauser, Boston.

Münte, T.F., Heinze, H.J., Mangun, G.R., 1993. Dissociation of brain activity related to syntactic and semantic aspects of language. J. Cognit. Neurosci. 5, 335–344.

Münte, T.F., Matzke, M., Johannes, S., 1997. Brain activity associated with syntactic incongruities in words and pseudo-words. J. Cognit. Neurosci. 9, 300–311.

Neville, H., Nicol, J.L., Barss, A., Forster, K.I., Garrett, M.F., 1991. Syntactically based sentence processing classes: evidence from event-related brain potentials. J. Cognit. Neurosci. 3, 151–165.

O'Seaghdha, P.G.O., 1997. Conjoint and dissociable effects of syntactic and semantic context. J. Exp. Psychol. Learn. Memory Cognit. 23, 807–828.

Osterhout, L., 1997. On the brain response to syntactic anomalies: manipulations of word position and word class reveal individual differences. Brain Lang. 59, 494–522.

Osterhout, L., Hagoort, P., 1999. A superficial resemblance doesn't necessarily mean you're part of the family: counterarguments to Coulson, King, and Kutas (1998) in the P600/SPS debate. Lang. Cognit. Processes 14, 1–14.

Osterhout, L., Holcomb, P.J., 1992. Event-related brain potentials elicited by syntactic anomaly. J. Memory Lang. 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. Lang. Cognit. Processes 8, 413–438.

Osterhout, L., Holcomb, P.J., 1995. Event-related potentials and language comprehension, in: Rugg, M.D., Coles, M.G.H. (Eds.), Electrophysiology of mind. Oxford Univ. Press, Oxford, pp. 171–215.

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. J. Exp. Psychol. Learn. Memory Cognit. 20, 786–803.

Osterhout, L., McLaughlin, J., Bersick, M., 1997b. Event-related brain potentials and human language. Trends Cognit. Sci. 1, 203–209.

Osterhout, L., Mobley, L.A., 1995. Event-related brain potentials elicited by failure to agree. J. Memory Lang. 34, 739–773.

Rösler, F., Friederici, A.D., Pütz, P., Hahne, A., 1993. Event-related brain potentials while encountering semantic and syntactic constraint violations. J. Cognit. Neurosci. 5, 345–362.

Tanenhaus, M.K., Trueswell, C., 1995. Sentence comprehension, in: Miller, J.L., Eimas, P.D. (Eds.), Speech, language, and communication. Academic Press, San Diego, pp. 217–262.

Trueswell, J.C., Tanenhaus, M.K., Garnsey, S.M., 1994. Semantic influences on parsing: use of thematic role information in syntactic ambiguity resolution. J. Memory Lang. 33, 285–318.

Trueswell, J.C., Tanenhaus, M.K., Kello, C., 1993. Verb-specific constraints in sentence processing: separating effects of lexical preference from garden-paths. J. Exp. Psychol. Learn. Memory Cognit. 19, 528–553.

Tyler, L.K., Marslen-Wilson, W.D., 1977. The on-line effects of semantic context on syntactic processing. J. Verbal Learn. Verbal Behav. 16, 683–692.

Van Berkum, J.J., Brown, C.M., Hagoort, P., 1999a. Early referential context effects in sentence processing: evidence from event-related brain potentials. J. Memory Lang. 41, 147–182.

Van Berkum, J.J., Hagoort, P., Brown, C.M., 1999b. Semantic integration in sentences and discourse: evidence from the N400. J. Cognit. Neurosci. 11, 657–671.

Van den Brink, D., Hagoort, P., 2003. The influence of semantic and syntactic context constraints on lexical selection and integration in spoken-word comprehension as revealed by ERPs. J. Cognit. Neurosci. Submitted for publication.

Vosse, T., Kempen, G.A.M., 2000. Syntactic structure assembly in human parsing: a computational model based on competitive inhibition and lexicalist grammar. Cognition 75, 105–143.

Zurif, E.B., 1998. The neurological organization of some aspects of sentence comprehension. J. Psycholing. Res. 27, 181–190.