16

The Binding Problem for Language, and Its Consequences for the Neurocognition of Comprehension

Peter Hagoort

For a linguist and a psycholinguist, this paper takes a slightly odd starting point. It considers the organization of sentence and discourse processing from the vantage point of brain sciences. This does not entail a change of the *explanans*, which remains the same, namely to provide an adequate account of the processing architecture of language processing. However, it assumes that useful additional constraints can be derived from our understanding of brain organization.

The domain of interest in this paper is not single-word processing, nor language production. The paper focuses on the interpretation processes beyond single-word recognition, at the level of the utterance and beyond (discourse). It is generally agreed among language-comprehension researchers that a characterization of interpretation as a concatenation of lexically stored single-word information is insufficient. In contrast, incoming sound or orthographic information triggers a cascade of memory-retrieval operations that make available the relevant basic ingredients for understanding. These include the morphophonological, semantic, and syntactic features of lexical items, which have to be combined in a principled way to bring about a coherent interpretation of the full input string. In analogy to the visual neurosciences, I will refer to the unification of different language-relevant feature types as the binding problem for language. Binding in this context refers to a problem that the brain has to solve, not to a concept from a particular linguistic theory.

The view that I develop in this chapter is strongly biased by research in which I was involved. I do not claim to do justice to the field as a whole. Nevertheless, I hope that I will succeed in sketching a picture of on-line language comprehension at the level of the sentence and beyond that is sufficiently motivated by the available empirical evidence.

The Binding Problem

One of the central questions in neuroscience is referred to as the binding problem. This problem is particularly well studied in the domain of vision. In short, it is the explanatory gap between the knowledge of relatively specialized brain areas for particular visual features (such as edges, color, motion, etc.) and the unified representation of the visual world that dominates awareness. How are the different attributes of an object, which are known to be processed in different cortical areas within visual cortex, brought together so that they result in a unified visual percept? One solution that has gained popularity in recent years, although it is still controversial, is that the mechanism of visual binding is related to the synchronicity of firing in the cell assemblies that code for the individual visual features (Varela et al. 2001).

A fair amount of data suggest that synchronicity of neuronal firing might be an important mechanism for visual binding. However, this does not guarantee that the same mechanism can solve the binding problem for language. In fact, I believe that synchronicity of firing cannot contribute to binding in the domain of language processing to the extent that it presumably does in visual perception. One major reason is that visual binding is more or less instantaneous. The relevant areas in visual cortex deliver their specific outputs (color information, motion information, etc.) within a very narrow time window. On the basis of the available experimental evidence, it is assumed that synchronous networks emerge and disappear at time scales between 100 and 300 msec (Varela et al. 2001). In contrast, one of the hallmarks of language processing is that information is spread out over relatively extended time periods. For instance, in parsing the auditory sentence "Noam thought of a couple of nice example sentences for his linguistics class but by accident wrote them down in his political diary," the information of Noam as the subject of the sentence still has to be available a second or so later when the acoustic information encoding the finite verb form 'wrote' has reached auditory cortex. In addition, the inherently hierarchical nature of language processing creates problems for a feature-binding account (such as when the same lexical features have to be bound into two different entities, as holds for the lexical features of "dog" in the phrase "the little, but not the big dog"). A feature-binding account does not seem to be able (at least in a straightforward way) to prevent the interpretation of this phrase as "the little big dog." This is known as the problem of 2 (Jackendoff 2002).

Crucially, the binding problem for language is how information that is processed not only in different parts of cortex, but also at different time scales and at relatively widely spaced parts of the time axis, can be unified into a coherent representation of a multi-word utterance.

One requirement for solving the binding problem for language is, therefore, the availability of cortical tissue that is particularly suited for maintaining information on-line, while binding operations take place. Prefrontal cortex (PFC) seems to be especially well suited for doing exactly this (Mesulam 2002). It has reciprocal connections to almost all cortical and subcortical structures, which puts it in a unique neuroanatomical position for binding operations across time, both within and across different domains of cognition.

In human evolution, PFC has shown a massive expansion. It occupies roughly one-third of the neocortical mantle in humans. Two major areas within PFC are lateral prefrontal cortex and orbitofrontal cortex. Lateral PFC includes portions of the inferior, middle, and superior frontal gyri. The posterior portions of the lateral PFC (roughly involving Brodmann's areas 9, 44, 45, and 46) are especially involved in various cognitive tasks (Knight and Stuss 2002). A core function of these areas is related to working memory; that is, to maintaining information over time and manipulating the contents during the maintenance period. Whether or not domain-specific subdivisions exist within lateral PFC is currently under debate. Another relevant area is orbitofrontal cortex, which is crucial for emotional and social control of cognitive function (Petrides and Pandya 2002). Activations related to sentence and discourse comprehension have been found in lateral prefrontal cortex, mainly in the left hemisphere. As I will argue below, this part of the brain is crucial for binding of phonological, syntactic, semantic, pragmatic, and presumably also non-linguistic contextual information (e.g. visuo-spatial, as in gestures) into a coherent discourse or situational model.

In addition, the left temporal cortex is suggested to play a critical role in storage and retrieval of linguistic information that language acquisition has laid down in memory. Thus, a major subdivision in the left-hemisphere temporofrontal language network is between the retrieval of lexically stored information (temporal cortex) and the on-line integration/binding of this information into the current context. How this constraint from considerations of brain organization fits to an explicit computational model, and to empirical data on language processing, will be discussed in more detail below for two crucial binding operations, namely syntactic binding and semantic binding.

Syntactic Binding

Recent accounts of the human language system (Jackendoff 1999, 2002; Levelt 1999) assume a cognitive architecture that 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 on-line 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.

Moreover, in recent linguistic theories, the distinction between lexical items and traditional rules of grammar is vanishing. For instance, Jackendoff (2002) proposes that the only remaining rule of grammar is UNIFY PIECES, "and all the pieces are stored in a common format that permits unification" (p. 180). The unification operation clips together lexicalized patterns with one or more variables in it. The operation MERGE in Chomsky's (1995) Minimalist Program has a similar flavor. Thus, phonological, syntactic, and semantic/pragmatic constraints determine how lexically available structures are glued together.

In models of language processing, there exists fairly wide agreement on the types of constraints that are effective during the formulation and the interpretation of sentences and beyond. However, disagreement prevails with respect to exactly how these are implemented in the overall sentence-processing architecture. One of the defining 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 or reading. The by-now-classical 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. An alternative view maintains that lexical-semantic information 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 information or by 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 non-syntactic context information on the structure that is assigned (Tanenhaus and Trueswell 1995; Van Berkum et al. 1999a). 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 that proposed by Friederici (2002). Based on the time course of different language-relevant ERP effects, Friederici 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 the second phase, lexical-semantic and morphosyntactic processes result in assignment of thematic roles. In the third phase, integration of the different types of information takes place, and the final interpretation results. This proposal is based mainly on findings in ERP studies on language processing. The last 15 years have seen an increasing number of ERP studies on syntactic processing, triggered by the discovery 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; figure 16.1). These studies have been followed up by a large number of ERP studies on syntactic processing that have provided a wealth of data. Here I will connect the known syntax-related ERP effects to a computational model of parsing (Vosse and Kempen 2000) that was developed to account for a large portion of behavioral findings in the parsing literature and for deficit patterns in aphasic patients. In the context of considerations based on brain organization, it makes the right distinction between lexicalized patterns and a unification component. However, before discussing the model, I will first discuss the relevant ERP results, then present some data that are incompatible with a syntax-first model. Later in this chapter, 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. Kutas and Hillyard observed a negative-going potential with an onset at about 250 msec and a peak around 400 msec (hence the N400), whose amplitude was increased when the semantics of the eliciting word (i.e., *socks*) mismatched with 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 and 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, 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 (figure 16.2; 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

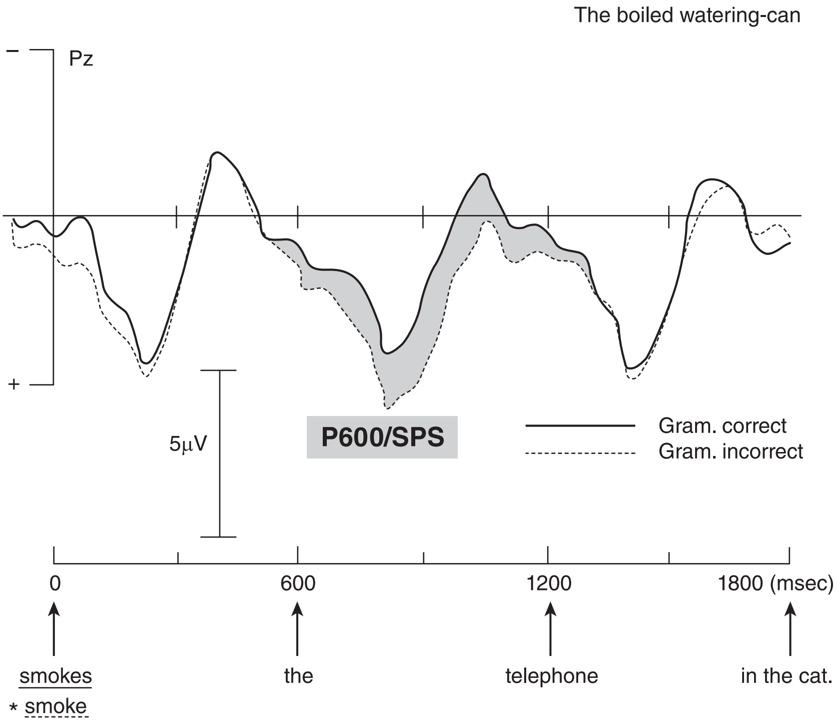

Figure 16.1
ERPs to visually presented syntactic prose sentences. These are sentences without a coherent semantic interpretation. A P600/SPS is elicited by a violation of the required number agreement between the subject-noun phrase and the finite verb of the sentence. The averaged waveforms for the 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 msec 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 (SOA) of 600 msec. Negativity is plotted upwards. (adapted from Hagoort and Brown 1994; copyright 1994 Erlbaum; reprinted by permission)

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 addition, recent evidence indi-

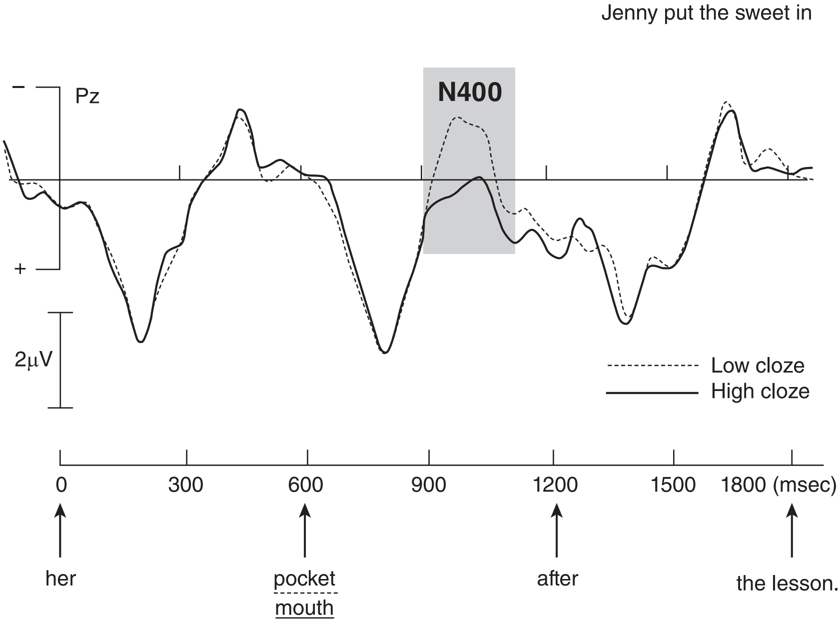

Figure 16.2
Modulation of the N400 amplitude as a result of a manipulation of the semantic fit between a lexical item and its sentence context. The grand-average waveform is shown for electrode site Pz (parietal midline), for the best-fitting word (high cloze), and a word that is less expected in the given sentence context (low cloze). The sentences were visually presented word by word, every 600 msec. In the figure the critical words are preceded and followed by one word. The critical word is presented at 600 msec on the time axis. Negativity is up. (adapted from Hagoort and Brown 1994; copyright 1994 Erlbaum; reprinted by permission)

cates that sentence verification against world knowledge in long-term memory modulates the N400 in the same way (Hagoort et al. 2004).

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). The modularity thesis 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 (e.g. Boland and Cutler 1996).

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 differ 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 hemisphere than over the right, although in many cases the distribution is bilateral (Hagoort et al. 2003b). 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 msec post-stimulus (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, between approximately 100 and 300 msec (Friederici 2002; Friederici et al. 1993; Neville et al. 1991).

In some studies, LAN effects have been reported to violations of word-category constraints (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 (e.g. a verb) is presented, early negativities are observed. Friederici (1995) and colleagues (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 attributed the very early negativities that occur approximately between 100 and 300 msec (labeled ELAN) to violations of word category, and the negativities between 300 and 500 msec to morphosyntactic processing.

LAN effects have also been related to verbal working memory in connection to filler-gap assignment (Kluender and Kutas 1993). This working-memory account of the LAN is compatible with the finding that lexical, syntactic, and referential ambiguities seem to elicit very similar frontal negativities (Hagoort and Brown 1994; Van Berkum et al. 1999a; Kaan and Swaab 2003b; King and Kutas 1995). Lexical and referential ambiguities are clearly not syntactic in nature, but can be argued to tax verbal working memory more heavily than sentences in which lexical and referential ambiguities are absent. Syntactic ambiguities may also tax working memory more strongly than their unambiguous counterparts. Future research should indicate whether or not these two functionally distinct classes of LAN effects can be dissociated at a finer grain of electrophysiological analysis.

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), with the critical verb form in italics; the * indicates the ungrammaticality of the sentence), a positive-going shift is elicited by the word that renders the sentence ungrammatical (Hagoort et al. 1993).

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

This positive shift starts about 500 msec after the onset of the violation and usually lasts for at least 500 msec. Because of 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 (Hagoort et al. 1993). 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), where one is semantically odd but grammatically correct and the other contains the same agreement violation as in (1).

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

If one compares the ERPs and the italicized verbs in (2a) and (2b), a P600/SPS effect is visible to the ungrammatical verb form (figure 16.1). Though 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 violations of phrase structure (Hagoort et al. 1993; Neville et al. 1991; Osterhout and Holcomb 1992), of subcategorization (Ainsworth-Darnell et al. 1998; Osterhout et al. 1997; Osterhout et al. 1994), of agreement of number, gender, and case (Coulson et al. 1998; Hagoort et al. 1993; Münte et al. 1997; Osterhout 1997; Osterhout and Mobley 1995), of subjacency (McKinnon and Osterhout 1996; Neville et al. 1991), and of the empty-category principle (McKinnon and Osterhout 1996). A P600/SPS has also been reported in relation to thematic-role animacy violations (Kuperberg, Sitnikova, Caplan, and Holcomb 2003). Moreover, a P600/SPS can be found with both written and spoken input (Friederici et al. 1993; Hagoort and Brown 2000a; Osterhout and Holcomb 1993).

In summary, two classes of syntax-related ERP effects have been consistently reported. These two classes differ in polarity, in topographic distribution, and in latency characteristics. In terms of latency, the first class of effects is an anterior negativity. Apart from LANs related to working memory, anterior negativities mainly appear in response to syntactic violations. In a later latency range, positive shifts occur that are elicited not only by syntactic violations, but also by complexity variation in grammatically well-formed sentences (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.

Broadly speaking, models of sentence processing can be divided into two types. One type of model assumes 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 (Jackendoff 2002; Marslen-Wilson 1989; Zwitserlood 1989). I will refer to this type of model as the *immediacy model*. Overall, the behavioral data, although not decisive, favor the second type of model more 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 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, semantic binding will be impaired if no syntactic structure can be built up. Certain 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 are immediately used when they become available without a priority for syntactic information over other information types.

ERP evidence for an autonomous syntax-first model for sentence processing is derived from a series of studies in which Friederici and colleagues found an ELAN in response to auditorily presented words whose prefix is indicative of a violation of word category. 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*"), where the prefixes "ge-" and "be-" in combination with the preceding auxiliary "wurde" indicate a past participle where the preposition "im" requires a noun. In this case a very early (between 100 and 300 msec) left anterior negativity is observed that precedes the N400 effect.

Although this evidence is compatible with a syntax-first model, it is not necessarily incompatible with an immediacy model of sentence processing. As long as word-category information can be derived earlier from the acoustic input 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 where word-category information comes in later than semantic information, the syntactic information will nevertheless be used earlier than semantic information in sentence processing. Van den Brink and Hagoort (2004) 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 an immediacy model, a syntax-first model would, in such a case, predict that semantic processing (more specifically, semantic binding) is postponed until after the information about the word category has become available.

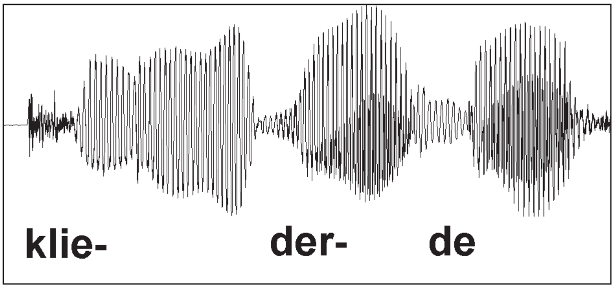

Figure 16.3
A waveform of an acoustic token of the Dutch verb form “kliederde” (messed). The suffix “-de” indicates past tense. The total duration of the acoustic token is approximately 450 msec. The onset of the suffix “-de” is at approximately 300 msec. After 300 msec of signal, the acoustic token can be classified as a verb. Thus, for a context that does not allow a verb in that position, the Category Violation Point (CVP) is at 300 msec into the verb.

Van den Brink and Hagoort (2004) compared correct Dutch sentences (see (3a)) with their anomalous counterparts (see (3b)) in which the critical word (italicized in (3)) was a semantic violation in the context and also had the incorrect word category. However, in contrast to the experiments by Friederici and colleagues, word-category information was encoded in the suffix ‘-de’.

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

Figure 16.3 shows the waveform of the spoken verb form ‘kliederde’ (messed). This verb form has a duration of approximately 450 msec. The stem already contains part of the semantic information. However, the onset of the suffix ‘-de’ is at about 300 msec into the word. Only at this point will it be clear that the word category is a verb, 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 (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.

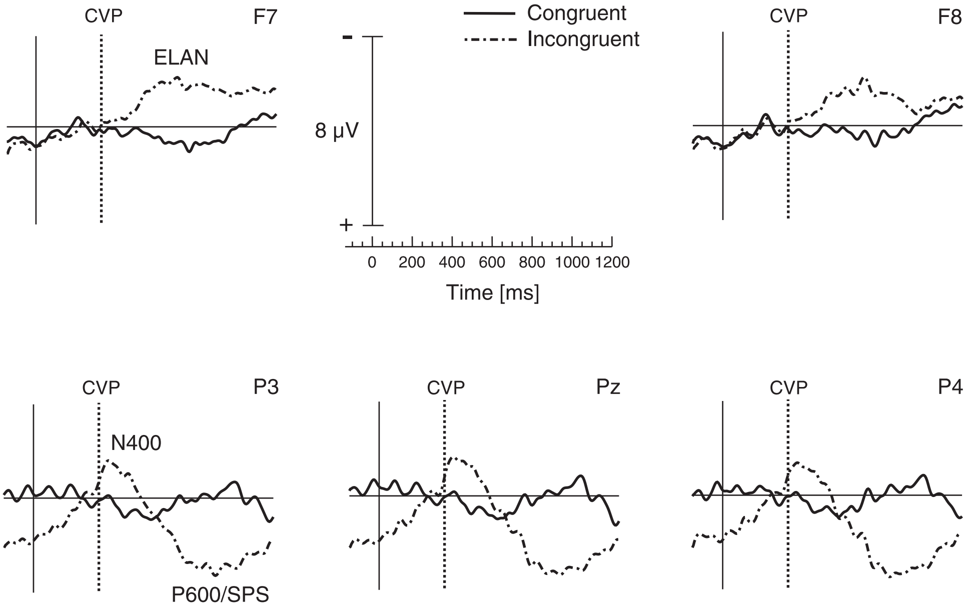

Figure 16.4
Connected speech. Grand-average ERPs from two frontal electrode sites (F7, F8) and three posterior electrode sites (Pz, P3, P4) to critical words that were semantically and syntactically congruent with the sentence context or semantically and syntactically incongruent. Grand-average waveforms were computed after time locking on a trial-by-trial basis to the moment of word-category violation (CVP: Category Violation Point). The baseline was determined by averaging in the 180–330-msec interval, corresponding to a 150-msec 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 over the three posterior sites. The onset of the ELAN is at 100 msec after the CVP; the onset of the N400 effect precedes the CVP by approximately 10 msec. (after Van den Brink and Hagoort 2004)

Figure 16.4 shows the averaged waveforms that are time-locked to the CVP for two frontal sites where the ELAN is usually 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 approximately 100 msec 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 information 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 need not wait until an initial structure is built on the basis of word-category information. A weaker syntax-first model, which allows prediction of word category, could claim that this prediction was only falsified at the CVP, and thus that semantic binding could be started in advance. However, this weaker version gives up the characteristic of bottom-up priority and assumes an interaction between syntactic context and lexical processing. One can then ask which feature of the processing architecture guarantees that interaction between context and lexical processing is restricted to syntax.

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. Apart from the LAN effects related to working memory, so far (E)LAN effects have mainly been seen 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 msec). The Anterior Negativities are normally followed by a P600/SPS. In contrast to the (E)LAN, the P600/SPS is not only seen in response 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 2003a,b).

The Unification Model

The increasing number of ERP studies on syntactic processing in the last 15 years has resulted in a substantial amount of data that are in need of a coherent overall account. 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. This proposal is certainly not the final version, but only a beginning. The model needs to be adapted, and the account of the ERP data needs to be refined. Nevertheless I believe that progress will be made only if we attempt to connect not only the behavioral data but also data from electrophysiology and neuroimaging to explicit computational accounts. I will first describe the general architecture of this model.

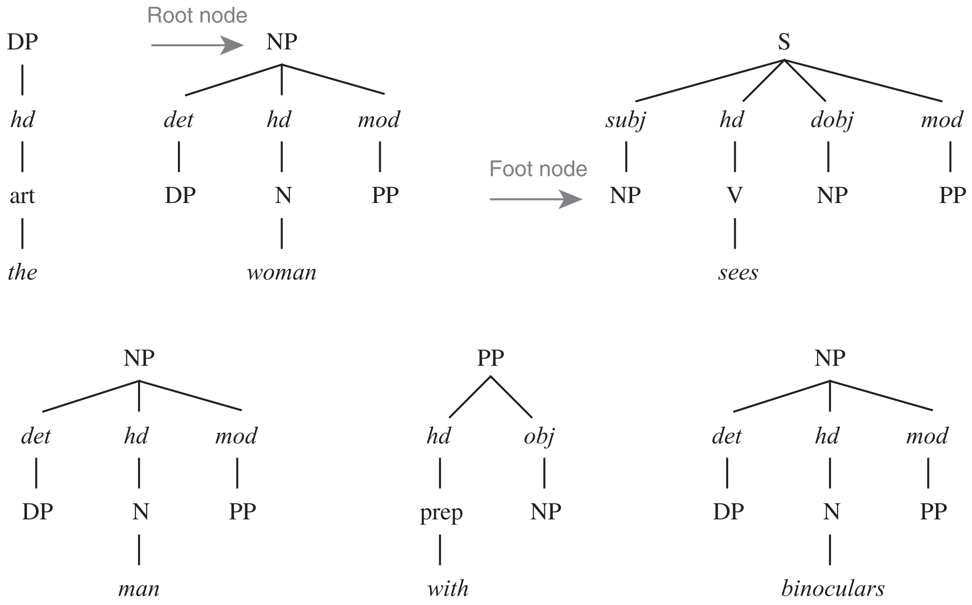

Figure 16.5
Syntactic frames in memory. These frames are retrieved on the basis of incoming word-form information for the example sentence "the woman sees the man with the binoculars." DP: determiner phrase. NP: noun phrase. S: sentence. PP: prepositional phrase. art: article. hd: head. det: determiner. mod: modifier. subj: subject. dobj: direct object.

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 figure 16.5; for details concerning the computation of word order, see Harbusch and Kempen 2002).

The top layer of the frame consists of a single phrasal 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 phrasal nodes to which lexical items or other frames can be attached.

This parsing account is "lexicalist" in the sense that all syntactic nodes (S, NP, VP, N, V, etc.) are retrieved from the mental lexicon. That is, chunks of syntactic structure are stored in memory. There are no syntactic rules that introduce additional nodes. In the on-line comprehension process, structural frames associated with the individual word forms incrementally enter the unification workspace. In this workspace, constituent structures spanning the whole utterance are formed by a unification operation. This operation consists

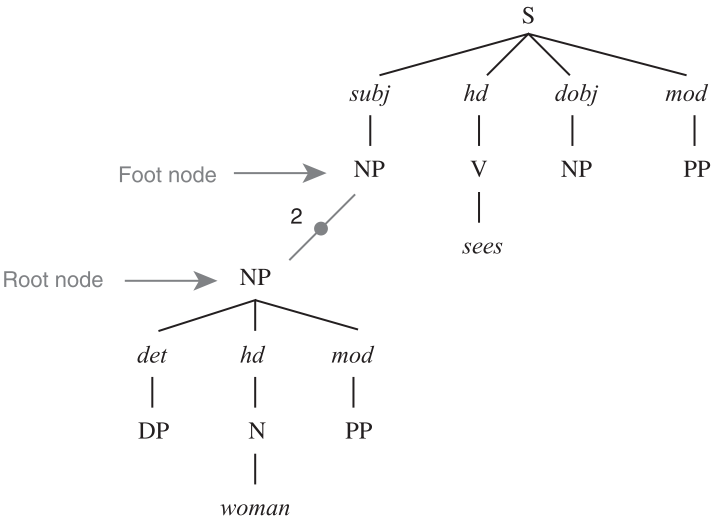

Figure 16.6
The unification operation of two lexically specified syntactic frames. The unification takes place by linking the root node NP to an available foot node of the same category. The number 2 indicates that this is the second link that is formed during on-line processing of the sentence "The woman sees the man with the binoculars."

of linking lexical frames that have matching root and foot nodes (see figure 16.6) and checking agreement features (number, gender, person, etc.). It specifies what Jackendoff (2002) refers to as the only remaining "grammatical rule": UNIFY PIECES.

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. Because of the ambiguity that is inherent 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 figure 16.7).

Ultimately, one phrasal configuration results. This requires that only one of the alternative binding candidates remain active. The required state of equilibrium is reached through a process of lateral inhibition between two or more alternative unification links. In general, owing to gradual decay of activation, more recent foot nodes will have a higher level of activation than those 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' (figure 16.7). In addition, the strengths of the

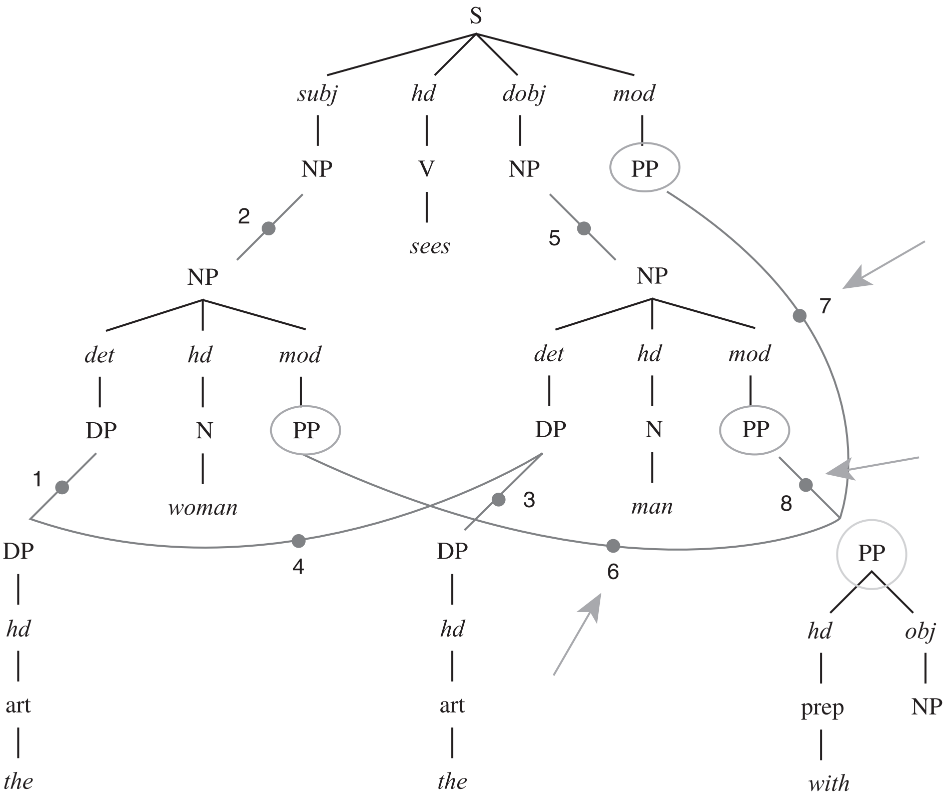

Figure 16.7
Lateral inhibition between three different PP-foot nodes that are candidate unification sites for the PP-root node of the preposition *with*. The three possible unification links are indicated by arrows. Lateral inhibition between these three possible unifications (6, 7, and 8) ultimately results in one unification that wins the competition and remains active.

unification links can vary as a 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 (figure 16.7) it means that the outcome of lateral inhibition is that the PP may be linked to the S-frame (Unification 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 time. In general, sentences are harder to analyze syntactically when more potential unification links of similar strength enter into competition with one another. Sentences are easy when the number of U links is small and the links are of unequal strength.

The Unification Model has these advantages: it is computationally explicit, it accounts for a large series of empirical findings in the parsing literature (but presumably not for all the locality phenomena in Gibson 1998) and in the neuropsychological literature on aphasia, and it belongs to the class of lexicalist parsing models that have found increasing support in recent years (Bresnan 2001; Jackendoff 2002; Joshi and Schabes 1997; MacDonald et al. 1994).

This model also nicely accounts for the two classes of syntax-related ERP effects reported in this paper and in and many others. In the Unification Model, binding (unification) is prevented in two cases: 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.e., NP) to bind to, and 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 (AN) 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 (see figure 16.8). As a result, unification fails.

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 build-up of U links will be affected by the partial mismatch in syntactic feature specification. Relative 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 preliminary account for the collective body of ERP data on syntactic processing.

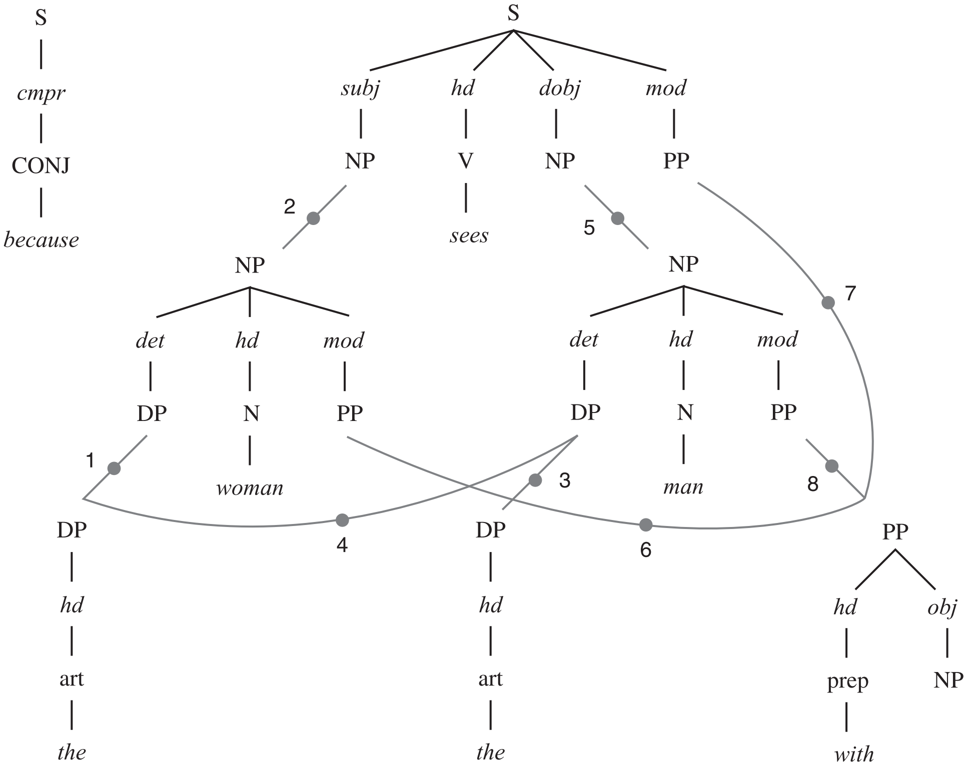

Figure 16.8
A dangling syntactic frame for the conjunction element *because*. This syntactic frame cannot be attached into the phrasal configuration for the remaining parts of sentence.

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 assume a syntax-first phase.

Semantic Binding

Along with syntactic binding, semantic binding operations have to take place. Studies of neuropsychological patients and data from neuroimaging studies suggest that semantic representations may be distributed, with the involvement of brain areas that support the most salient aspects of a concept (e.g., visual, kinesthetic, linguistic or propositional) (Allport 1985; Saffran and Sholl 1999). Context can differentially activate or select the saliency of meaning aspects (as in “The girl gave a wonderful performance on the old piano” vs. “Four men were needed to transport the old piano”). At the same time, the semantic aspects retrieved on the basis of lexical access have to be integrated into a coherent interpretation of a multi-word utterance. This I will refer to as *semantic binding*. It turns out that left lateral prefrontal cortex is also crucial for semantic binding (see below). Binding-relevant areas within the left prefrontal cortex (LPC) may overlap, at least to some degree, for syntactic and semantic binding. But there is also evidence that semantic binding may involve more ventral areas (especially Brodmann's area 47) than syntactic binding. More research is needed to determine commonalities and differences in LPC between areas involved in phonological, syntactic, and semantic binding. However, the qualitative differences between ERP effects of semantic (N400) and syntactic (LAN, P600) binding suggest that the brain honors the distinction between these two types of binding operations.

The Level of Semantic Binding: Sentence vs. Discourse

A central issue for semantic binding is whether or not a semantic representation at the sentence level is built up first, before semantic information is integrated into a discourse model in a second step. For instance, in their blueprint of the listener, Cutler and Clifton (1999) assume that utterance interpretation on the basis of syntactic analysis and thematic processing takes place first, before integration into a discourse model. Kintsch (1998; see also Ericsson and Kintsch 1995) has made similar claims. We conducted an ERP study to investigate how and when the language-comprehension system relates an incoming word to semantic representations of the unfolding local sentence and the wider discourse (Van Berkum et al. 1999b). In the first experiment, subjects were presented with short stories, of which the last sentence sometimes contained a critical word that was semantically anomalous with respect to the wider discourse (e.g., "Jane told the brother that he was exceptionally slow" in a discourse context where he had in fact been very quick). Relative to a discourse-coherent counterpart (e.g., 'quick'), these discourse-anomalous words elicited a large N400 effect (i.e., a negative shift in the ERP that began about 200 to 250 msec after word onset and peaked around 400 msec).

In addition to the discourse-related anomalies, sentence-semantic anomaly effects were elicited under comparable experimental conditions. We found that the ERP effects elicited by both types of anomalies were highly similar. Relative to their coherent counterparts, discourse-anomalous and sentence-anomalous critical words elicited an N400 effect with an identical time course and identical scalp topography (figure 16.9). The similarity of these effects, particularly in polarity and scalp distribution, is compatible with the claim that they reflect the activity of a largely overlapping or identical set of underlying neural generators, indicating similar functional processes.

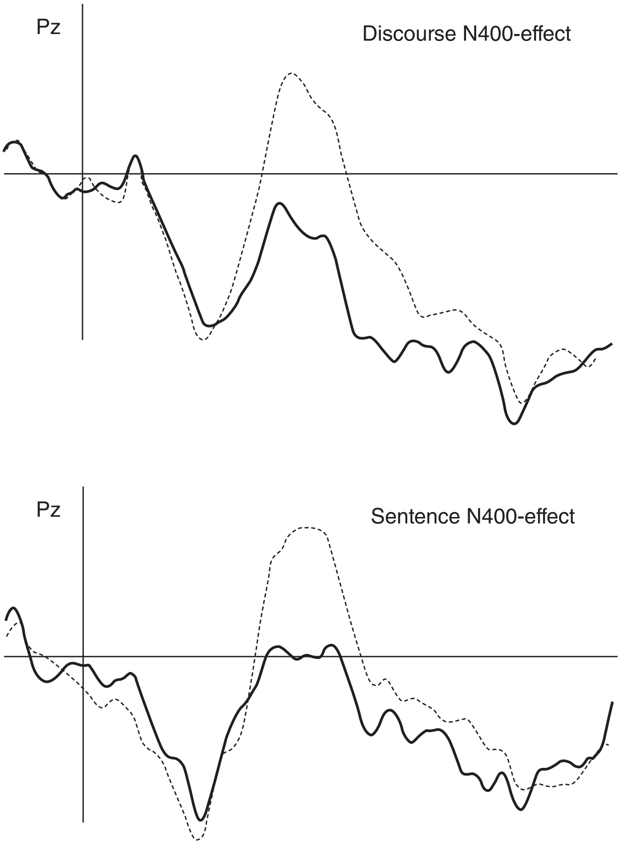

Figure 16.9
N400 effects triggered by discourse-related and sentence-related anomalies. Waveforms are presented for a representative electrode site. The latencies of the N400 effect in discourse and sentence contexts (both onset and peak latencies) are the same. (after Van Berkum et al. 1999b)

In summary, there is no indication that the language-comprehension system is slower in relating a new word to the semantics of the wider discourse than in relating it to local sentence context. Our data clearly do not support the idea that new words are related to the discourse model *after* they have been evaluated in terms of their contribution to the semantics of the sentence. The speed with which discourse context affects processing of the current sentence appears to be at odds with recent estimates of how long it would take to retrieve information about preceding discourse from long-term memory. In the materials of Van Berkum et al., the relative coherence of a critical word usually hinged on rather subtle information that was implicit in the discourse and that required considerable inferencing about the discourse topic and the situation it described. Kintsch (1998; see also Ericsson and Kintsch 1995) has suggested that during on-line text comprehension such subtle discourse information is not immediately available and must be retrieved from "long term working memory" when needed. This is estimated to take some 300–400 msec at least. However, the results of our experiments suggest that the relevant discourse information can be brought to bear on local processing within at most 200–250 msec.

The observed identity of discourse-level and sentence-level N400 effects is most parsimoniously accounted for by a processing model that abandons the distinction between sentence-level and discourse-level semantic binding. This is compatible with the notion of *common ground* (Stalnaker 1978; Clark 1996). Clark's analysis clearly demonstrates that the meaning of linguistic utterances cannot be determined without taking into account the knowledge that the speaker and the listener share and mutually believe they share. This common ground includes a model of the discourse itself, which is continually updated as the discourse unfolds. If listeners and readers always immediately evaluate new words relative to the discourse model and the associated information in common ground (i.e., immediately compute "contextual meaning"), the identity of the ERP effects generated by sentence and discourse anomalies has a natural explanation. With a single sentence, the relevant common ground includes only whatever discourse and world knowledge has just been activated by the sentence fragment presented so far. With a sentence presented in discourse context, the relevant common ground will be somewhat richer, now also including information elicited by the specific earlier discourse. But the process that maps incoming words onto the relevant common ground can run into trouble either way. The N400 effects observed by Van Berkum et al. (1999b) reflect the activity of this unified binding process. Of course, this is not to deny the relevance of sentential structure for semantic interpretation. In particular, how the incoming words are related to the discourse model is co-constrained by sentence-level syntactic devices (such as word order, case marking, local phrase structure, or agreement), and by the associated mapping onto thematic roles. However, this is fully compatible with the claim that there is no separate stage during which word meaning is exclusively evaluated with respect to "local sentence meaning," independent of the discourse context in which that sentence occurs.

Binding Plasticity

It is often assumed in language-comprehension research that all information has to be available at the right moment for binding operations to occur. However, in the reality of daily communication it is not uncommon that the system works under noisy conditions. Therefore, it might well be that the system works with what it gets, which is very often non-optimal. With a noisy signal or an underspecified context, one type of binding operation might be more easily achieved than another. The comprehension system adapts to changing circumstances, and can change the weights assigned to the different binding operations that run in parallel accordingly. Some evidence for this idea came from a recent ERP study with agrammatic aphasics in which Hagoort, Wassenaar, and Brown (2003a) investigated the ERP effects of syntactic violations in aphasic patients and their elderly controls.

The most interesting finding was that the ERP response to one type of syntactic violation (a violation of word order in adverb-adjective-noun sequences) was qualitatively different from the ERP effect in the non-agrammatic subjects. In these latter subjects (elderly controls and a group of non-agrammatic aphasics), the word-order violation resulted in a P600/SPS. In contrast, the ERP of the agrammatic aphasics was dominated by the N400 effect that is usually observed to semantic binding operations during on-line language processing. Thus, whereas word-order violations triggered a syntax-related ERP response in normal controls and non-agrammatic comprehenders, the same violations triggered an ERP response related to semantic binding in Broca's aphasics with agrammatic comprehension. Interestingly, a similar shift can be seen in early second-language learners (Osterhout, personal communication).

We offered the following explanation for this semantic ERP response in agrammatic aphasics: The absence of a P600/SPS suggests that the agrammatic aphasics are no longer able to exploit syntactic information during sentence comprehension. The N400 effect for the word-order violations suggests that these sentences were processed through another (compensatory) processing route. The lack of a syntax-related ERP effect suggests that the agrammatic comprehenders did not interpret these sentences through a hierarchically organized phrase-structure representation. Instead, word meanings were incrementally integrated in the semantic representation of the linear string of preceding words, where the interpretation process was more difficult when the adjective preceded the adverb (*thief steal expensive very* . . .) than in the reverse order (*thief steal very expensive* . . .). In the adjective-before-adverb word order, the internal event structure is less coherent than in the correct order, owing to the reversal of the semantic arguments of the denoted event. That is, in the semantic context of *thief steal expensive*, the canonical structure of events is better matched by mentioning what is being stolen than by further expanding the meaning of *expensive* (as in *thief steal expensive very*). The results indicate that agrammatic patients still have access to this level of semantic information and are able to use this during real-time processing. This is not to say that their use of semantic information is optimal, but it is certainly less affected than syntactic processing operations. As such, the relative preservation of a semantic processing route presumably results in the N400 effect.

The data therefore address the real-time functioning of the language system under impairment: the way in which different sources of linguistic information are combined to derive an interpretation seems to be tailored to the processing options that are still available to the impaired language-comprehension system. The results we obtained suggest that a semantic processing stream provides an optimization of language comprehension within the limitations imposed by a syntactic deficit resulting from brain damage. Although this does not imply that semantic processing is fully optimal in agrammatic aphasics, it is relatively preserved compared to syntactic processing. Under impairment, the comprehension system seems to weigh the remaining information differently or more strongly. This multiple-route plasticity instantiates the potential for on-line adaptation to impairments in the language-comprehension system.

What holds for the system under impairment might also hold for the intact system working under conditions of some external noise. It just works with what it gets, and is, to some degree, capable of flexibility and adaptation to what is the most salient information under the current circumstances. Language comprehension is characterized by processing that is "opportunistic" rather than rigidly regulated (Jackendoff 2003).

The Neural Implementation of Binding in Language

In the context of the language system, the binding problem refers to the following question: How is information that is incrementally retrieved from the mental lexicon unified into a coherent overall interpretation of a multi-word utterance? Most likely, unification must take place at the conceptual, syntactic, and phonological levels, as well as across these levels (Jackendoff 2002). So far I have discussed the features of the cognitive architecture for syntactic and semantic binding. In this section I will argue that the left inferior prefrontal cortex may have the characteristics necessary for performing the unification operations at the different levels of the language system.

One requirement for solving the binding problem for language is the availability of cortical tissue that is particularly suited for maintaining information on-line while binding operations take place. Prefrontal cortex seems to be especially well suited for doing exactly this. Areas in prefrontal cortex are able to hold information on-line (Mesulam 2002) and to select among competing alternatives (Thompson-Schill, D'Esposito, and Kan 1999). Electrophysiological recordings in the macaque have shown that this area is important for sus-

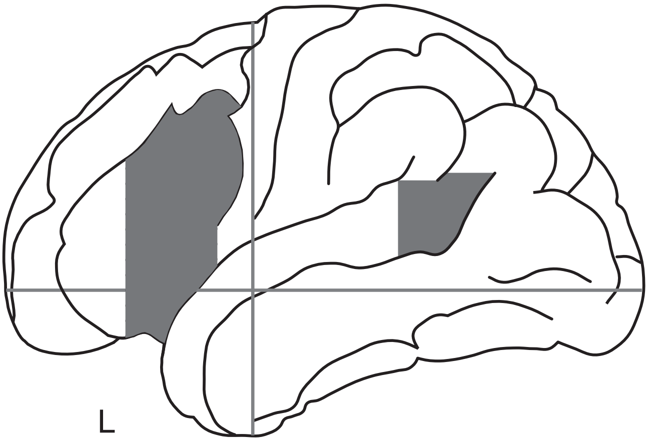

Figure 16.10
Common areas of activation (shaded) in a meta-analysis of 28 imaging studies on the processing of syntactic information during language comprehension. The activated areas are shown on a lateral view of the left hemisphere. They were restricted to the temporal and frontal lobes of that hemisphere. (after Indefrey 2003)

taining information triggered by a transient event for many seconds (Miller 2000). This allows prefrontal cortex to establish unifications between pieces of information that are perceived or retrieved from memory at different moments in time.

I will make some tentative suggestions about how the different components of the Unification Model for syntactic binding that I discussed above could be connected to our knowledge about the neural architecture. This proposal is not yet explicitly tested, but, as I will argue, it makes good sense in the 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 (see figure 16.10).

As is known from lesion studies in aphasic patients, lesions in different areas of left perisylvian cortex can result in deficits in syntactic processing in sentence comprehension (Caplan, Hildebrandt, and Makris 1996). The idea that modality-independent grammatical knowledge was mainly represented in Broca's area (Zurif 1998) has thus been proved incorrect. At the same time, the left posterior temporal cortex is known to be involved in lexical processing (Indefrey and Cutler 2004). In connection to the Unification Model, this part of the brain may 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, may be localized in the left frontal part of the syntax-relevant network of brain areas. One of the main specializations of prefrontal cortex is the holding on-line and binding of information (Mesulam 2002). It may be the right area for providing the computational resources for binding together 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 left superior temporal cortex relevant for storage and retrieval of syntactic frames, and with the left prefrontal cortex important for binding these frames together. The need for combining independent bits and pieces into a single coherent percept is not unique for syntax. Models for semantic/conceptual unification and phonological unification could be worked out along similar lines as the Unification Model for syntax. Recent neuroimaging studies suggest that parts of prefrontal cortex in and around Broca's area may be involved in conceptual and phonological unification, with Brodmann Areas (BA) 47 and 45 involved in semantic binding, BA 45 and 44 in syntactic binding, and BA 44 and ventral BA 6 in phonological binding (see figure 16.11).

Six Principles of the Processing Architecture

In analogy to other domains of cognitive neuroscience, for language comprehension I have made the distinction between memory retrieval and unification or binding. I have discussed features of the processing architecture for syntactic and semantic binding. Evidence from neuroimaging studies seems to support the distinction between brain areas recruited for memory retrieval and brain areas crucial for binding. Based on the evidence discussed in the preceding sections, I propose the following six general architectural principles for comprehension beyond the single-word level:

(i) The brain honors the distinction between syntactic and semantic binding. However, both involve contributions from the left prefrontal cortex (in and around Broca's area), it being the workspace where unification operations take place. It is very well possible that this area is not language-specific but also subserves other functions (e.g. binding in music; see Patel 2003). Left prefrontal cortex is suggested to maintain the activation state of representational structures retrieved from memory (the mental lexicon), and to provide the necessary neuroanatomical space for binding operations.

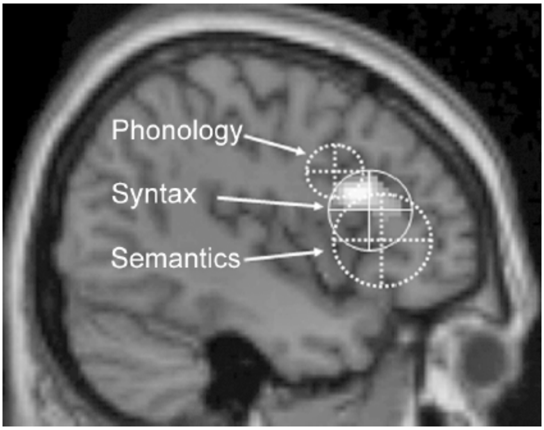

Figure 16.11
The gradient in left inferior frontal cortex for activations and their distribution, related to semantic, syntactic, and phonological processing, based on the meta-analysis in Bookheimer 2002. Centers represent the mean coordinates of the local maxima; radii represent the standard deviations of the distance between the local maxima and their means (courtesy of Karl Magnus Petersson). The activation shown is from artificial grammar violations (Petersson et al. 2004).

(ii) Immediacy is the general processing principle of binding. Semantic binding does not wait until relevant syntactic information (such as word-class information) is available, but starts immediately with what it derives on the basis of the bottom-up input and the left context. The corollary of immediacy is incrementality: output representations are built up from left to right in close temporal contiguity to the input signal.

(iii) There does not seem to be a separate stage during which word meaning is exclusively integrated at the sentence level. Incremental interpretation is, for the most part, done by an immediate mapping onto a discourse model (Clark 1996).

(iv) In parsing, lexically specified structures enter the unification space. Lexical information (e.g. animacy), discourse information, and (recent data suggest) inputs from other modalities (e.g., visual world, gesture) immediately influence the competition between alternative binding options, and can change the binding links. However, in the absence of competing binding sites, assignment of structure is not influenced by non-syntactic information.

(v) There is no evidence for a privileged position of syntax and/or a processing priority for syntax, as is assumed in syntax-first models. The different processing levels (phonological, syntactic, semantic/pragmatic) operate in parallel, and to some degree independently. Where necessary, cross-talk takes

place, which is again characterized by the immediacy principle. That is, cross-talk takes place more or less moment-to-moment.
(vi) Within certain limitations, the language-comprehension system can adapt the weight of evidence in the light of system-internal or system-external noise. The degrees of freedom in language comprehension are much greater than in language production.

References

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

Allport, D. A. 1985. Distributed memory, modular subsystems and dysphasia. In S. K. Newman and R. Epstein (eds.), *Current Perspectives in Dysphasia*. Churchill Livingstone.

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

Boland, J. E., and 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.

Bookheimer, S. 2002. Functional MRI of Language: New approaches to understanding the cortical organization of semantic processing. *Annual Review of Neuroscience* 25, 151–188.

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

Brown, C., and Hagoort, P. (eds.). 1999. *The Neurocognition of Language*. Oxford University Press.

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

Chomsky, N. 1995. *The Minimalist Program.* MIT Press.

Cutler, A., and Clifton, C. E. 1999. Comprehending spoken language: A blueprint of the listener. In C. M. Brown and P. Hagoort (eds.), *The Neurocognition of Language*. Oxford University Press.

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

Clark, H. H. 1996. *Using Language*. Cambridge University Press.

Coulson, S., King, J. W., and Kutas, M. 1998. Expect the unexpected: Event-related brain response to morphosyntactic violations. *Language and Cognitive Processes* 13, 21–58.

Duncan, J., and Miller, E. K. 2002. Cognitive focus through adaptive neural coding in the primate prefrontal cortex. In R. T. Knight (ed.), *Principles of Frontal Lobe Function*. Oxford University Press.

Ericsson, K. A., and Kintsch, W. 1995. Long-term working memory. *Psychological Review* 102, 211–245.

Fodor, J. D. 1983. *The Modularity of Mind*. MIT Press.

Frazier, L. 1987. Sentence processing: A tutorial review. In M. Coltheart (ed.), *Attention and Performance XII*. Erlbaum.

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. 2002. Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences* 6, 78–84.

Friederici, A. D., Hahne, A., and Mecklinger, A. 1996. Temporal structure of syntactic parsing: Early and late event-related brain potential effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22, 1219–1248.

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

Gibson, E. 1998. Linguistic complexity: Locality of syntactic dependencies. *Cognition* 68, 1–76.

Hagoort, P. 2003. Interplay between syntax and semantics during sentence comprehension: ERP effects of combining syntactic and semantic violations. *Journal of Cognitive Neuroscience* 15, 883–899.

Hagoort, P., and Brown, C. M. 1994. Brain responses to lexical ambiguity resolution and parsing. In K. Rayner (ed.), *Perspectives on Sentence Processing*. Erlbaum.

Hagoort, P., and Brown, C. M. 2000a. ERP effects of listening to speech compared to reading: the P600/SPS to syntactic violations in spoken sentences and rapid serial visual presentation. *Neuropsychologia* 38, 1531–1549.

Hagoort, P., and Brown, C. M. 2000b. ERP effects of listening to speech: Semantic ERP effects. *Neuropsychologia* 38, 1518–1530.

Hagoort, P., Brown, C. M., and 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., and Osterhout, L. 1999. The neurocognition of syntactic processing. In P. Hagoort (ed.), *Neurocognition of Language*. Oxford University Press.

Hagoort, P., Hald, L., Bastiaansen, M., and Petersson, K. M. 2004. Integration of Word Meaning and World Knowledge in Language Comprehension. *Science* 304, 438–441.

Hagoort, P., Wassenaar, M., and Brown, C. 2003a. Real-time semantic compensation in patients with agrammatic comprehension: electrophysiological evidence for multiple-route plasticity. *Proceedings of the National Academy of Sciences of the United States of America* 100, 4340–4345.

Hagoort, P., Wassenaar, M., and Brown, C. M. 2003b. Syntax-related ERP-effects in Dutch. *Cognitive Brain Research* 16, 38–50.

Hahne, A., and 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. *Cognitive Brain Research* 11, 199–212.

Harbusch, K., and Kempen, G. 2002. A quantitative model of word order and movement in English, Dutch and German complement constructions. Paper presented at 19th International Conference on Computational Linguistics, San Francisco.

Indefrey, P. 2003. Hirnaktivierungen bei syntaktischer Sprachverarbeitung: Eine Meta-Analyse. In G. Rickheit and H. M. Mueller (eds.), *Neurokognition in der Sprache*. Stauffenburg.

Indefrey, P., and Cutler, A. 2004. Prelexical and lexical processing in listening. In M. D. Gazzaniga (ed.), *The Cognitive Neurosciences*, third edition. MIT Press.

Jackendoff, R. 1999. The representational structures of the language faculty and their interactions. In P. Hagoort (ed.), *The Neurocognition of Language*. Oxford University Press.

Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press.

Jackendoff, R. 2003. Précis of foundations of language: Brain, meaning, grammar, evolution. *Behavioral and Brain Sciences* 26, 651–707.

Joshi, A. K., and Schabes, Y. 1997. Tree-adjoining grammars. In A. Salomma and G. Rosenberg (eds.), *Handbook of Formal Languages and Automata*, volume 3. Springer-Verlag.

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

Kaan, E., and Swaab, T. Y. 2003a. Electrophysiological evidence for serial sentence processing: A comparison between non-preferred and ungrammatical continuations. *Cognitive Brain Research* 17, 621–635.

Kaan, E., and Swaab, T. Y. 2003b. Repair, revision and complexity in syntactic analysis: An electrophysiological differentiation. *Journal of Cognitive Neuroscience* 15, 98–110.

King, J. W., and Kutas, M. 1995a. A brain potential whose latency indexes the length and frequency of words. *Newsletter of the Center for Research in Language* 10, 3–9.

King, J. W., and Kutas, M. 1995b. 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.

Kintsch, W. 1998. *Comprehension: A Paradigm for Cognition*. Cambridge University Press.

Kluender, R., and Kutas, M. 1993. Subjacency as a processing phenomenon. *Language and Cognitive Processes* 8, 573–633.

Knight, R. T., and Stuss, D. T. 2002. Prefrontal cortex: The present and the future. In R. T. Knight (ed.), *Principles of Frontal Lobe Function*. Oxford University Press.

Kuperberg, G. R., Sitnikova, T., Caplan, D., and Holcomb, P. J. 2003. Electrophysiological distinctions in processing conceptual relationships within simple sentences. *Cognitive Brain Research* 17, 117–129.

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

Kutas, M., and Van Petten, C. K. 1994. Psycholinguistics electrified: Event-related brain potential investigations. In M. A. Gernsbacher (ed.), *Handbook of Psycholinguistics*. Academic Press.

Levelt, W. J. M. 1999. Producing spoken language: A blueprint of the speaker. In C. M. B. P. Hagoort (ed.), *The Neurocognition of Language*. Oxford University Press.

MacDonald, M. C., Pearlmutter, N. J., and Seidenberg, M. S. 1994. Lexical nature of syntactic ambiguity resolution. *Psychological Review* 101, 676–703.

Marslen-Wilson, W. 1989. Access and integration: Projecting sound onto meaning. In W. Marslen-Wilson (ed.), *Lexical Representation and Process*. MIT Press.

Marslen-Wilson, W., Brown, C. M., and Tyler, L. K. 1988. Lexical representations and language comprehension. *Language and Cognitive Processes* 3, 1–16.

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

McKinnon, R., and Osterhout, L. 1996. Constraints on movement phenomena in sentence processing: Evidence from event-related brain potentials. *Language and Cognitive Processes* 11, 495–523.

Mesulam, M.-M. 2002. The human frontal lobes: Transcending the default mode through contingent encoding. In R. T. Knight (ed.), *Principles of Frontal Lobe Function*. Oxford University Press.

Miller, E. K. 2000. The prefrontal cortex and cognitive control. *Nature Review Neuroscience* 1, 59–65.

Münte, T. F., and Heinze, H. J. 1994. ERP negativities during syntactic processing of written words. In H. J. Heinze, T. F. Münte, and G. R. Mangun (eds.), *Cognitive Electrophysiology*. Birkhäuser.

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

Münte, T. F., Matzke, M., and Johannes, S. 1997. Brain activity associated with syntactic incongruities in words and pseudo-words. *Journal of Cognitive Neuroscience* 9, 300–311.

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

O'Seaghdha, P. G. O. 1997. Conjoint and dissociable effects of syntactic and semantic context. *Journal of Experimental Psychology* 23, 807–828.

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

Osterhout, L., Bersick, M., and McKinnon, R. 1997. Brain potentials elicited by words: Word length and frequency predict the latency of an early negativity. *Biological Psychology* 46, 143–168.

Osterhout, L., and 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. *Language and Cognitive Processes* 14, 1–14.

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

Osterhout, L., and 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, 413–438.

Osterhout, L., and Holcomb, P. J. 1995. Event-related potentials and language comprehension. In M. G. H. Coles (ed.), *Electrophysiology of Mind*. Oxford University Press.

Osterhout, L., Holcomb, P. J., and 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.

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

Patel, A. D. 2003. Language, music, syntax and the brain. *Nature Neuroscience* 6, 674–681.

Petersson, K. M., Forkstam, C., and Ingvar, M. 2004. Artificial syntactic violations activates Broca's region. *Cognitive Science* 28, 383–407.

Petrides, M., and Pandya, D. N. 2002. Association pathways of the prefrontal cortex and functional observations. In R. T. Knight (ed.), *Principles of Frontal Lobe Function*. Oxford University Press.

Rösler, F., Friederici, A. D., Pütz, P., and Hahne, A. 1993. Event-related brain potentials while encountering semantic and syntactic constraint violations. *Journal of Cognitive Neuroscience* 5, 345–362.

Saffran, E., and Sholl, A. 1999. Clues to the functional and neural architecture of word meaning. In C. M. Brown and P. Hagoort (eds.), *The Neurocognition of Language*. Oxford University Press.

Stalnaker, R. C. 1978. Assertion. In P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*. Academic Press.

Tanenhaus, M. K., and Trueswell, C. 1995. Sentence comprehension. In P. D. Eimas (ed.), *Speech, Language, and Communication*. Academic Press.

Thompson-Schill, S. L., D'Esposito, M., and Kan, E. P. 1999. Effects of repetition and competition on activity in left prefrontal cortex during word generation. *Neuron* 23, 513–522.

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

Trueswell, J. C., Tanenhaus, M. K., and Kello, C. 1993. Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19, 528–553.

Tyler, L. K., and Marslen-Wilson, W. D. 1977. The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behavior* 16, 683–692.

Van Berkum, J. J., Brown, C. M., and Hagoort, P. 1999a. Early referential context effects in sentence processing: Evidence from event-related brain potentials. *Journal of Memory and Language* 41, 147–182.

Van Berkum, J. J., Hagoort, P., and Brown, C. M. 1999b. Semantic integration in sentences and discourse: Evidence from the N400. *Journal of Cognitive Neuroscience* 11, 657–671.

Van den Brink, D., and Hagoort, P. 2004. The influence of semantic and syntactic context constraints on lexical selection and integration in spoken-word comprehension as revealed by ERPs. *Journal of Cognitive Neuroscience* 16, 1068–1084.

Varela, F., Lachaux, J.-P., Rodriguez, E., and Martinerie, J. 2001. The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience* 2, 229–239.

Vosse, T., and 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. *Journal of Psycholinguistic Research* 27, 181–190.

Zwitserlood, P. 1989. The locus of the effects of sentential-semantic context in spoken-word processing. *Cognition* 32, 25–64.