Syntax and the Brain

Jon Sprouse
Department of Cognitive Sciences
University of California, Irvine

Ellen F. Lau
Department of Linguistics
University of Maryland

1. The role of syntactic theory in a cognitive neuroscience of syntax

Marr (1982:19-29) argued that a complete description of any information processing device, including the human brain, will necessarily involve three levels of description: the computational level, the algorithmic level, and the implementational level. The computational level can best be described as an answer to the question "what problem is being solved by the device?" (and to some extent "why does the problem have the form that it does?"). The algorithmic level is a description of the specific operations necessary to solve the problem. And the implementational level is a description of how those operations are implemented in the hardware of the device itself. Although this partitioning of the problem into three neat levels of analysis can be abused (see Phillips and Lewis 2010 for discussion), it also provides a good starting point for understanding the relationship between syntactic theory, the theory of sentence processing within psycholinguistics, and the neural implementation of sentence processing within neurolinguistics. Therefore we will begin our discussion with Marr's three levels and how they map to the cognitive neuroscience of sentences.

1.1 The computational level – syntactic theory

Marr famously used a cash register to illustrate the three levels, so it seems appropriate to continue the tradition. The computational theory of a cash register (i.e. the problem that a cash register must solve) is the tallying of the cost of purchasing multiple items from a store, or more abstractly, the theory of addition. The properties of addition are familiar from arithmetic: (i) the order of addition for two numbers has no effect on the final sum (commutativity), (ii) the grouping of three or more items into pairs prior to addition also has no effect on the final sum (associativity), (iii) the addition of zero to a number has no effect on the number (identity), (iv) and the addition of a number and its inverse results in zero. As Marr (1982: 22) stated, "these properties are part of the fundamental *theory* of addition", which means that "they are true no matter how the numbers are written… and no matter how the addition is executed." In many ways, syntactic theory is the computational theory of the "human sentence processor", as syntactic theory seeks to describe the fundamental properties of the sentences that are the input to comprehension and the output of production – properties that are true of sentences no matter how they are actually constructed during processing. However, as we will see shortly, the three levels are highly interactive, and the difficulties inherent in studying the human brain mean that it is impossible to completely isolate the computational properties of sentences from the algorithmic and implementational properties the same way that one can with a cash register.

1.2 The algorithmic level – sentence processing

Once the computational theory is established, it is then possible to investigate how it is that the computational problem is solved. In the case of the cash register's need to solve the problem of addition, several possible algorithms are possible depending on the basic units that are used to compose the representation. For example, if the basic units of the representation are the base-10 Arabic numerals (0, 1, 2, 3…), then one possible algorithm is the "carry-over" procedure familiar to all school children: add the digits farthest to the right (least significant), and if the sum exceeds 9, carry over the leftmost digit of the sum to the next column to the left. Crucially, other units could be chosen for the representation, such as base-2 or binary units (0,1), which will necessitate a different algorithm for executing the addition. In other words, representational assumptions constrain algorithmic choices. Furthermore, there may be several different algorithms available for any given choice of units of representation, in which case the choice of algorithm may depend upon issues of efficiency or the desire to make certain types of information more easily available than others. Just as syntactic theory is in many ways a computational theory of sentences, the theories of sentence processing developed by psycholinguists are in many ways algorithmic theories. As we will see shortly, theories of sentence processing begin with different representational assumptions drawn from the computational theory (some assume that the basic units are words and syntactic rules, others that the basic units are syntactic frames) that constrain the types of algorithms that are proposed (structure-generation versus retrieval and unification).

1.3 The implementational level – neurolinguistics

Obviously there are multiple ways to physically implement a device like a cash register. One could use multiple spinning cylinders like the analog cash registers of the 20th century, or microchips as is common in computers today. However, the choice of implementation constrains both the units of the representation and the algorithms that are available for executing addition: analog cylinders tend to use base-10 numerals and an algorithm similar to the "carry-over" procedure, whereas computers use a binary representation and a bit-wise calculation. When it comes to cognitive tasks such as vision or sentence processing, the physical implementation is obviously the human brain. And although we know relatively little about the types of representations and algorithms that the human brain can (and cannot) implement, it is clear that the areas of the brain that form the network for any given cognitive task must be capable of carrying out at least one algorithm that solves the computational problem of the task. The task of developing an implementational theory of sentences rests with neurolinguists, as they leverage neuropsychological patient studies and various neuroimaging technologies (e.g., fMRI, PET, EEG, MEG) to identify the neural circuits involved in sentence processing.

1.4 The interaction across levels

In principle, the three levels are only loosely related: for any given computational theory, there are any number of algorithms that can be adopted, and any number of physical devices that can be used to implement a given algorithm. In fact, the very definition of the computational level suggests that it can be investigated completely independently of the other two levels, as it is defined as the properties of the problem that are independent of how it is solved. However, in practice it is logically impossible to study human language at only one level. Any study of the implementational level must begin with an assumption about the set of algorithms (often called *processes* in the language literature) that one wants to localize. Any study of the algorithmic level must begin with an assumption about the syntactic structures (often called *representations*) that must be built, and the units that can be used in the construction (e.g., words and syntactic rules, or just syntactic frames). And any study of the computational level (i.e., syntactic theory) can only be constructed by observing human behavior. For example, the primary behavioral response used as evidence in generative syntax is the acceptability judgment, as it is assumed that one of the primary factors affecting acceptability judgments is the well-formedness of the syntactic structure. However, it is well known that acceptability judgments are also affected by factors that are more closely related to the algorithmic or implementational levels such as processing complexity (e.g., Chomsky and Miller 1963). The behaviors that form the empirical basis of the computational theory of language are rooted in algorithms implemented in the human brain.

The interaction of the three levels in language research makes the goal of a complete description of the cognitive neuroscience of language both challenging and exciting. The interaction of the three levels also suggests that progress toward that goal can only be achieved through a close interaction of syntacticians, psycholinguists, and neuroscientists. The field has made some strong advances in this respect over the past several decades, but the work is far from over. Our goal in this chapter is to illustrate the close, but often unstated, relationship between syntactic theory and brain-based studies of sentence processing (both electrophysiological and neuroimaging). Our hope is that by highlighting this relationship, we can draw attention to both the areas where there has been productive cross-fertilization, and the areas where the relationship could be closer. Along the way, we also intend to provide a primer for syntacticians on some of the driving questions underlying brain-based studies of sentence processing, as well as some of the primary results of those investigations[1].

2. The mentalistic commitments of syntactic theories

Syntactic theories are primarily concerned with properties of the full syntactic representation of a sentence, such as the structural configuration of the words in the sentence, and the various relationships that exist between those words (agreement, coreference, etc.). Decades of syntactic research have revealed that syntactic representations are complex cognitive objects containing a sophisticated set of relationships. Despite this complexity, there is generally broad agreement among the various syntactic theories about what the properties of the full syntactic representation are for any given sentence. However, syntactic theories disagree vehemently about the

[1] It should be noted from the outset that this is not intended to be, and indeed could not be, a complete review of the psycho- or neurolinguistics literature on sentence processing. It is unfortunate that many interesting areas (and results) of psycho/neurolinguistics research on sentence processing had to be left out of this chapter in the interest of presenting a coherent narrative in the space available. We hope that readers with psycho/neurolinguistics backgrounds will interpret these omissions as nothing more than the consequence of these limitations, and that readers who are new to psycho/neurolinguistics studies can use the (copious) references that we have supplied to find the many interesting studies that we could not cover here.

fundamental building blocks of syntactic representations, and the combinatorial mechanics necessary to combine those building blocks into all and only the syntactic representations that are licit in any given language. The different assumptions by the various syntactic theories necessarily entail different mentalistic commitments about the language units that must be stored in long-term memory in the brain, the cognitive processes necessary to retrieve those units from memory, and the cognitive processes that must be deployed to assemble those units into the relevant syntactic representations. And as we shall see in section 4, those mentalistic commitments also affect how the results of neuroimaging studies are interpreted to reveal the neural circuitry involved in sentence processing.

To be sure, it is not always straightforward to discern the mentalistic commitments of any given syntactic theory. In fact, given the delimited goal of syntactic theory to describe syntactic representations, and Marr's proviso that there are likely to be any number of algorithms to execute a given computational theory, there is no reason to believe that the mentalistic commitments of a given syntactic theory will fully constrain the space of possible sentence processes (indeed, they do not). However, a discussion of the potential mentalistic commitments of syntactic theories is an important component of any discussion of syntax and the brain, as these commitments play defining roles in the interpretation of electrophysiological and neuroimaging evidence. We believe that if progress in generative syntax is to benefit the broader cognitive neuroscience community, and vice versa, the relationship between the mentalistic commitments of syntactic theories and the evidence gathered in cognitive neuroscience experiments must be a central part of the discussion.

2.1 The generative framework

The following three questions provide a helpful starting point for the discussion of the mentalistic commitments of generative syntactic theories, and indeed, any syntactic theories.

What are the basic units of the syntactic representation?

What types of operations are necessary to build the syntactic representation from the basic units?

How are the syntactic operations influenced by non-syntactic information that may be available during sentence comprehension (e.g., semantic information)?

As we shall see in the following sections, the answers to these questions are highly dependent on each other. And given the difficulties inherent in studying the human brain, it is not always possible to obtain empirical evidence that unambiguously supports one answer over another.

2.1.1 The units of the syntactic representation in the generative framework

Even within the generative framework there is healthy debate and active research on the basic units of the syntactic representation (Chomsky 1995, Halle and Marantz 1993, Pollard and Sag 1994, Bresnan 2001). However, we believe it is fair to say that there is some degree of consensus that the basic units are *bundles of features*. These features can encode morphosyntactic properties of the bundle (e.g. phi-features: Person, Number, Gender), morphosyntactic requirements of the bundle (e.g., Case features), discourse properties (e.g., focus), meaning (e.g., lexical semantic features), and even pronunciation (e.g., phonetic features). The bundles themselves can map to the canonical idea of a word, to a morpheme, or even to a silent element such as a functional head or covert pronoun. The overarching architectural idea in the generative framework is that if certain features are present in a sentence (e.g., the *uninterpretable features* in the minimalist program), they must be arranged in very specific syntactic configurations in order for that sentence to be well formed. Although the rules for combining the features into well-formed configurations take various forms depending on the precise version of generative theory, it is generally true that the combinatory rules (e.g., phrase structure rules, transformations) are stored separately from the basic units themselves (which are hypothesized to be stored in the mental lexicon).

2.1.2 The operations that build the syntactic representation in the generative framework

The generative commitment to (bundles of) features as the basic unit of the syntactic representation has direct consequences for the generative commitment to syntactic structure building operations during sentence processing. For example, in order to combine bundles of features, the minimalist program posits a grammatical structure-building operation *merge*, which concatenates two bundles of features into a single unit. Because sentence processing is necessarily incremental (word-by-word) and directional (the order of the words is "left to right"), the parsing equivalent of merge must (at least superficially) be subdivided into the options made available by the syntactic rules of the language. Presented with the following partial sentence, the parser can potentially concatenate the preceding word in any one of four structural positions:

(1) Bill spread the rumor that Mary left…
(i) the complement of the immediately preceding head (e.g., *the band*)
(ii) the complement of a non-immediately previous head (e.g., *to John*)
(iii) the adjunct of the most recent phrase (e.g., *suddenly*)
(iv) the adjunct of a previous phrase (e.g., *maliciously*)

The choice of concatenation location is licensed by the syntactic rules governing feature combinatorics, and presumably mediated by the parsing strategies that the human parser utilizes to rapidly and efficiently assign structure to incoming sentences (e.g. minimal attachment and late closure: Frazier 1978, Frazier and Fodor 1978, Frazier 1987). Whether there are multiple concatenation operations (one for each location) or just a single operation is an open question, and a good example of how syntactic theories will not completely constrain the space of possible parsing processes.

In addition to *merge*, many generative theories also include the operation *move*, which displaces words or phrases from their canonical position to a position elsewhere in the sentence. There is quite a bit of research in the sentence processing literature devoted to the cognitive operations necessary to parse displaced elements, a full review of which is well beyond the scope of this chapter (Crain and Fodor 1985, Stowe 1986, Traxler and Pickering 1996, Nicol and Swinney 1989, Nicol, Fodor, and Swinney 1994; but cf. McKoon, Ratcliff, and Ward 1994, Garnsey, Tanenhaus, and Chapman 1989, Kaan, Harris, Gibson, and Holcomb 2000, Phillips, Kazanina, and Abada, 2005, Felser, Clahsen, and Münte 2003, Sussman and Sedivy 2003, Wagers and Phillips 2009). One relatively popular view of the parsing of movement dependencies maps the grammatical operation *movement* to (at least) three parsing operations: (i) encoding of the displaced element (called the *filler*) in working memory (for interpretation later in the sentence), (ii) an active search for potential gap locations (verbs, prepositions, etc.), (iii) integration of the filler with the appropriate gap location. From a parsing perspective, syntactic constraints on the grammatical operation *move* are realized as constraints on the distribution of gap locations, and thus potential surface as constraints on the gap-searching operation. Taken together with the discussion of *merge* above, we hope it is clear that generative theories appear to be committed to the assumption that (bundles of) features are the basic unit of syntactic structure, and to the assumption that the combinatory rules for combining those features are stored separately from the features themselves. This leads to the postulation of several (syntactic) parsing operations, and the idea that the constraints on grammatical operations that are familiar from generative syntax surface as constraints on the various (syntactic) parsing operations (e.g., Stowe 1986, Clifton and Frazier 1989, Neville et al. 1991, Phillips 2007, Wagers and Phillips 2009).

2.1.3 The interaction of syntactic and semantic information in the generative framework

In addition to the quantity and quality of parsing operations, syntactic theories have constrained theories of the architecture of the human sentence parser in other ways. One prominent example of this can be seen in the long-standing debate in the sentence processing literature about the temporal dynamics of "syntactic" and "semantic" operations during sentence comprehension. A priori, there are at least three logical possibilities:

(i) Syntactic operations are computed independently of, and functionally prior to, semantic operations (often called *syntax-first* models)
(ii) Syntactic and semantic operations interact from the earliest stages of sentence processing (highly interactive models are often called *constraint-based*, as both syntactic and semantic constraints have equal importance and precedence)
(iii) Syntactic and semantic operations can operate independently of the other

As we will see in section 3.3 and again in section 4, some care is necessary when interpreting macroscopic labels such as "syntax" and "semantics", as different authors can have very different conceptions of what phenomena fit under each label, especially when it comes to the fuzzy boundaries between morphology and syntax, and between lexical and compositional semantics. It is also important to note the distinction between computing the syntactic structure for a string and choosing one from multiple licit syntactic structures; for syntactically ambiguous strings, all parsing theories are likely to agree that the parser may use semantic information at some point to arbitrate between the possibilities.

Classic generative syntax assumes that syntactic rules alone are sufficient to characterize the set of well-formed sentences in a given language, but this assumption does not, strictly speaking, constrain the temporal relationship between syntactic and semantic operations in parsing. For example, nothing about the representational commitments of the generative theory rules out a processing module that rapidly generates likely propositions from a non-syntactic 'bag-of-words' representation, which is then used by the parser to decide between alternative syntactic parses (see Townsend and Bever 2001 for another example). However, since the generative theory already posits syntactic representations that are sufficient to uniquely determine argument structure, it might seem more parsimonious to assume a parser that simply computes the syntactic representation first and then uses the syntactic representation to compute the interpretation, as in (i). This kind of consideration has led to a historical association between generative theories of syntax and syntax-first theories of parsing.

2.2. The unification framework

The parallel architecture proposed by Jackendoff (1999, 2002) may provide a good second case study in the role of syntactic theory in a broader cognitive neuroscience of sentences. The parallel architecture has been adopted by Hagoort and colleagues (e.g., Hagoort 2003, 2005) to interpret a range of electrophysiological and neuroimaging findings. As we will describe in this section, the parallel architecture has led Hagoort to propose a radically different conception of parsing in which there is only one parsing process, *unification*, which acts on syntactic, semantic, and phonological representations simultaneously (i.e., in parallel).

2.2.1 The units of the representation in the unification framework

The basic units within the unification framework are similar to the basic units in generative syntax in that they contain three types of information: syntactic, semantic, and phonological. However, whereas generative syntax makes a distinction between the units and the rules that combine them, the unification framework collapses the syntactic rules into the units themselves by assuming that each unit (or word) is stored in the lexicon as part of a syntactic frame that specifies the structure(s) that the unit can appear in (as well as a semantic frame and a phonological frame). This is similar in some respects to other frame-based theories such as construction grammar within cognitive linguistics (e.g., Goldberg 1995, 2007, Croft and Cruse 2004), and Tree Adjoining Grammar within computational linguistics (Joshi and Schabes 1997, Vosse and Kempen 2000; Frank, *this volume*). In other words, in the unification framework, the basic units are small chunks of syntactic structure that are stored in the mental lexicon.

2.2.2 The operations that build the representation in the unification framework

Because the basic units of the unification framework are lexicalized chunks of syntactic structure, the unification framework does not require several different parsing operations to compose the variety of syntactic structures made available by each language. Instead, there are just two operations: lexical *retrieval*, which we assume that all theories must posit since all theories must retrieve lexical material from the mental lexicon, and *unification*, which is the process of integrating two syntactic frames together. In cases of ambiguity, the assumption is that two (or more) frames will be accessed from the lexicon, and a winner-takes-all competitive process based on various factors (plausibility, frequency, temporal decay, etc) will yield a single phrasal configuration.

2.2.3 The interaction of syntactic and semantic information in the unification framework

Unification is assumed to occur in parallel at all three levels (syntactic, semantic, and phonological). The unification framework is thus a type of constraint-based parsing architecture, in which syntactic and semantic information is assumed to be highly interactive from the earliest stages of the parsing process. In this way the mentalistic commitments of the unification framework (lexicalized syntactic frames, a single parsing process, highly interactive syntax and semantics) are very different from the mentalistic commitments of generative syntax (a distinction between lexical units and syntactic rules, which allows for the possibility of multiple parsing operations and/or a syntax-first parsing architecture), and begin to demonstrate the role that syntactic theory can play in a full cognitive neuroscience of sentences.

2.3 Other frameworks

Our focus on the generative framework and the unification framework in this chapter is not meant to suggest that these are the only two syntactic theories that are useful in interpreting electrophysiological and neuroimaging data. Our choices simply reflect the theme of this volume and our goal of concisely illustrating the important role that the mentalistic commitments of syntactic theories play. There are other frameworks that the interested reader may wish to investigate, such as the argument-dependency framework of Bornkessel and Scheslewsky (2006) based upon the theoretical framework of Van Valin and La Polla (1997) in which syntactic representations are built using a frame/unification approach, but argument linking is computed according to various supra-syntactic precedence hierarchies familiar from typological research (animacy, referentiality, case-marking). The argument-dependency framework will appear in the discussion of Broca's area in section 4; however, many other interesting frameworks had to be left out due to space limitations.

3 Electrophysiological responses

When it comes to non-invasively studying brain responses to cognitive stimuli, there really are only two widely available options: electrophysiological responses or hemodynamic responses. We will begin our discussion with electrophysiological responses. The most common method for measuring electrophysiological responses is electroencephalography (EEG), which is the placement of electrodes on the scalp to measure the underlying brain-based electrical activity. The advantage of all electrophysiological measures is that electricity travels very quickly, even in biological substrates, such that changes in the underlying brain state can be detected with millisecond level temporal accuracy. The advantage of EEG in particular is that the machines themselves are relatively affordable (on the order of tens of thousands of dollars), and have been around for decades (the first EEG recordings were done in the 1930s, the first language-related responses were recorded in the 1970s). The disadvantage of EEG is that it is very difficult to determine the spatial origin of electrical activity that is recorded on the scalp, because the distribution of electrical activity is distorted by passing through various layers of biological matter before being realized on the two-dimensional surface of the scalp. In other words, EEG sacrifices spatial resolution for temporal resolution and cost-efficiency (Nunez and Srinivasan 2006). More recently, magnetoencephalograms (MEG) have been developed that can measure the magnetic fields that are created by the electrical activity in the brain. MEG offers the same temporal resolution of EEG, and overcomes many of the spatial resolution problems of EEG as magnetic fields are not significantly affected by intervening biological matter. The two primary disadvantages of MEG are (i) the machines themselves are relatively rare because they are extremely expensive compared to EEG (on the order of millions of dollars), and (ii) MEG can only measure magnetic fields that leave the head, leaving some sources of neuronal activity undetectable. The majority of sentence-related electrophysiological studies have been conducted using EEG.

3.1 Event-related potentials

Because the brain simultaneously controls almost all human activity, from core bodily functions like breathing, to sensory functions like vision, to high-level cognitive functions like language, the electrical activity that can be measured on the scalp at any given moment is potentially composed of activity from hundreds of neuronal sources. This means that the mapping procedure from scalp-recorded activity to the cognitive process(es) of interest is extremely difficult. One highly successful technique for eliminating activity that is not related to the cognitive process(es) of interest is to look for activity that is both time-locked and phase-locked to a specific event, such as the onset of a critical word during sentence comprehension. The process for this is a simple averaging procedure: EEG activity is recorded for multiple trials of each condition, and then the multiple trials are averaged together using the event as the time reference point (time point 0). The use of the event as the reference point time-locks the resulting averaged activity, and the averaging procedure phase-locks the averaged activity according to the principles of constructive and destructive wave interference: for each frequency of activity present in the raw EEG recordings, waves that are in phase (peaks line up with peaks) survive, and waves that are out of phase (peaks line up with troughs) are eliminated. The idea behind this averaging procedure is that any non-event related electrical activity will be eliminated by destructive interference, leaving behind only the electrical activity that is related to the event of interest. The resulting averaged activity is called an Event-related potential, or ERP (see Luck 2005 for an introduction to the ERP technique).

Although time-locking the EEG to a stimulus is what makes it possible to estimate the timing of some processes down to a few milliseconds, this temporal precision paradoxically carries its own limitations. Operations that are tied very closely to a particular word of input, such as early visual or auditory processing, orthographic processing, and lexical access, are easier to associate with a particular time window in the ERP following the stimulus presentation. In contrast, one cannot always assume that the sentence-level and discourse-level processes that support syntactic and semantic combination will happen at a neat fixed time interval from the presentation of a particular word, as their timing is likely to be dependent on many more factors (e.g., how much structure has been built already, how much structure could be predicted before the word was presented, how much time was required for basic lexical access and selection). For this reason, much of the ERP literature on sentence processing has focused on the ERP responses to syntactic and semantic violations of different types, which allows investigation of processing at a given representational level in a fixed time-window by ensuring that the presentation of one particular word will disrupt processing at that representational level. Although this approach has proved very powerful, we must be careful to keep in mind that a syntactic violation does not only impact basic syntactic operations, but may also invoke reanalysis or rereading operations and other strategic mechanisms.

There is some terminological confusion in the ERP literature. This is because an ERP is a physiological response to a stimulus (e.g., the critical word in a sentence), but cognitive theories are generally interested in unobservable cognitive processes that can only be isolated by comparing *two* (or more) stimuli that are hypothesized to differ with respect to that process.

Therefore there are times when the same label is given to both a component of the ERP for an individual stimulus and the relative difference between the ERPs of two experimentally matched stimuli. For example, we will shortly review one prominent ERP known as the N400 (Kutas and Hillyard 1980). The N400 is a negative-going wave peaking around 400ms that is elicited by all meaningful stimuli. As such, every word is expected to have an N400 ERP. However, the label N400 is also used for the relative difference in the amplitude of the N400 between two experimentally matched words. Although such terminological ambiguity is potentially troublesome, context is usually sufficient to disambiguate the intended meaning. Nonetheless, it is important to be aware of these two senses, as well as the fact that the cognitive interpretation of ERPs depends entirely on the theoretical assumptions underlying the hypothetical difference between two stimuli. As we shall see, the mentalistic commitments of syntactic theories play a large role in those assumptions.

While many ERPs have been identified in the broader EEG literature, four ERPs have played a central role in the sentence processing literature: the early left anterior negativity (ELAN), the left anterior negativity (LAN), the N400, and the P600. Although a complete review of each of these ERPs would be beyond the scope of this article, in this section we will provide a basic review of each and discuss how these ERPs have been interpreted according to various syntactic theories.

The ELAN

As the name suggests, the ELAN (early left anterior negativity) is a negative-going deflection that peaks in a relatively early processing window (100-250ms post-stimulus onset) and is greatest over left anterior electrode sites. The ELAN was first reported by Neville et al 1991 to a specific phrase structure violation in which a preposition appears in an ungrammatical position (note that the critical position must contain either a noun or an adjective):

(2) a. The boys heard Joe's stories *about* Africa.
b. *The boys heard Joe's *about* stories Africa.

A similar effect was reported by Friederici et al. (1993) in German, in this case when a participle appears in a position that must contain a noun or adjective (**Das Baby wurde im gefürttert*/*The baby was in-the fed*). The ELAN has since been elicited to very similar phrase structure violations in Spanish (Hinojosa et al. 2003), French (Isel et al. 2007), and further replicated in English (Lau et al. 2006, Dikker et al. 2009) and German (e.g., Hahne and Friederici 1999, Hahne and Friederici 2002, Rossi et al. 2005). The ELAN is not affected by task (Hahne and Friederici 2002), by the probability of the violation in the experiment (Hahne and Friederici 1999), or by the frequency of a disambiguated structure (Ainsworth-Darnell, Shulman, and Boland 1998, Friederici et al. 1996). Taken as a whole, these results suggest that the ELAN is a very specific response to phrase structure violations, and not simply a response to difficult or unlikely structures.

Recent research on the ELAN has focused on the extremely early latency of the response. The 100-250ms post-stimulus window is remarkably early for syntactic analysis (and error diagnosis) given that estimates of lexical access often center around 200ms post-stimulus (Allopenna, Magnuson, and Tanenhaus 1998, van Petten, Coulson, Rubin, Plante, and Parks 1999). Three approaches have been offered to explain the early latency of the ELAN. Friederici

(1995) adopts a parsing model in which the earliest stage considers only word category information (e.g., Frazier 1978, 1987, 1990, Frazier and Rayner 1982), thus limiting the number of processes that need to be performed in the earliest time window. Lau et al. (2006) suggest that the early latency can be explained if the parser has predicted the properties of the critical word prior to encountering it, such that many of the syntactic features are in some sense "pre-parsed". Dikker et al. (2009) propose the "sensory ELAN hypothesis", in which the ELAN indexes a processing stage prior to lexical access that occurs in the sensory cortices (visual or auditory cortex). This pre-lexical processing is based purely on the form typicality of the words – i.e., the sensory cortices use the probability of certain phonetic forms to determine if the incoming string is most likely a noun, verb, etc. Though it is too early to declare a dominant view, it is clear that any interpretation of the ELAN must explain two facts: (i) that it only arises to very specific violations (phrase structure violations), and (ii) that it occurs in an extremely early time window.

The LAN

While the LAN and the ELAN share many properties (i.e., they are both negative-going deflections that occur primarily over left anterior electrode sites), they differ along two critical dimensions. First, the LAN occurs in a slightly later time window, usually 300-500ms post-stimulus onset, which eliminates many of the complex timing questions associated with the ELAN. Second, the LAN has been elicited by a broad array of (morpho-)syntactic violations, such as agreement violations (Coulson et al. 1998, Gunter et al. 1997, Münte et al. 1997, Kaan 2002, Osterhout and Mobley 1995), case violations (Münte and Heinze 1994), phrase structure violations, (Friederici, Hahne, and Mecklinger 1996, Hagoort, Wassenaar, and Brown 2003) island constraint violations (Kluender and Kutas 1993b), and even garden-path sentences (Kaan and Swab 2003). The LAN has also been elicited during the processing of long-distance dependencies such as wh-movement, at both the displaced wh-word and the unambiguous cue for the gap location (Kluender and Kutas 1993a, Phillips, Kazanina, and Abada 2005).

The N400

The N400 is a negative-going deflection that is generally largest over centro-parietal electrode sites, and tends to occur 300-500ms post-stimulus onset (with a peak amplitude occurring at 400ms). The N400 was first found by Kutas and Hillyard (1980) when they presented participants with sentences that ended with unexpected words. They compared baseline sentence with semantically congruent endings (a) to sentences with semantically incongruent endings (b) and sentences with endings that were incongruent due to the physical properties of the stimulus such as words written in all capital letters (c):

(3) a. I spread the warm bread with butter.
b. I spread the warm bread with socks.
c. I spread the warm bread with BUTTER.

Kutas and Hillyard (1980) observed a larger N400 for (b) compared to (a), and a larger P300 (also known as a P3b) to (c) compared to (a). This qualitative difference in the responses to (b) versus (a) suggests that the N400 is specifically related to semantic processes rather than general error detection. In the decades since its discovery, the N400 has been elicited by a broad array of linguistic and non-linguistic stimuli, with the common pattern being that they are all meaningful in some way: spoken words, written, words, signed words, pseudowords, acronyms, environmental sounds, faces, and gestures (Kutas, Van Petten, and Kluender 2006).

Although the idea that the N400 is related to semantic processes is almost universally accepted, there has been quite a bit of debate about the exact nature of those processes in the literature due to the complex pattern of N400 results that have been reported. For example, although semantic incongruence often elicits a larger N400 than semantic congruence, this is not always the case. Congruent endings that are less predictable elicit a larger N400 than congruent endings that are more predictable: a sentence like *I like my coffee with cream and* ***honey*** produces a larger N400 than *I like my coffee with cream and* **sugar** because *honey* is less predictable than *sugar*, though both are semantically plausible endings. It is also the case that the N400 is affected by the degree of semantic relatedness between the realized ending and the unrealized predicted ending. For example, both *salt* and *socks* are incongruent endings to the "coffee" sentence above. However, *salt* produces a smaller N400 than *socks*, presumably because *salt* shares more semantic features with the predicted ending *sugar* than *socks* does (Federmeier and Kutas 1999, Kutas and Federmeier 2000).

There are two popular theories of the N400 effect. The first is that the N400 indexes processes related to the semantic integration of the critical word with the preceding semantic context. Under this view, increases in N400 amplitude reflect the increased difficulty of integrating incongruent, unexpected, or semantically unrelated words into the preceding context (Hagoort 2008, Osterhout and Holcomb 1992, Brown and Hagoort 1993). The second view is that the N400 indexes processes related to the activation of semantic features in the mental lexicon (long-term or semantic memory). Under this view, decreases in N400 amplitude reflect the ease of activation (or pre-activation) for congruent, predicted, and semantically related words (Federmeier and Kutas 1999, Kutas and Federmeier 2000, Lau et al. 2009).

Though this chapter is primarily focused on brain responses to syntactic processing, it would be a mistake to overlook the role of the N400 in sentence processing theories. The N400 is often interpreted as an index of "semantic" processing, especially when it comes to the relative timing of syntactic and semantic processes. However, it is important to regard macroscopic labels such as "semantics" and "syntax" with healthy skepticism: though no one would argue that the N400 is related to aspects of the semantic processing of words in sentences, it is unlikely that the N400 is a direct reflection of the compositional semantic processes that occupy much of semantic theory, especially within generative grammar. For example, the N400 is not sensitive to classic sentence-level compositional semantic effects such as negation and quantification. Sentences such as *A robin is a* ***tree*** elicit a larger N400 at the final word when compared to *A robin is a* ***bird***. However, adding negation to the sentence does not reverse the N400 pattern (*A robin is not a* ***tree/bird***); in fact, *tree* still elicits the larger N400, even though *tree* is the true continuation (Fischler et al. 1983). The same is true for other scope-taking elements (Kounios and Holcomb 1992). Furthermore, N400 effects are derived from non-sentential paradigms such as word-word priming paradigms just as often as they are derived from sentential paradigms. Although there is likely some form of composition between two words in a priming paradigm, it is unlikely to be identical to the compositional processes at the core of sentence level compositional semantics. N400 effects have also been modulated by discourse-level manipulations (van Berkum et al. 2003, Niewland and van Berkum 2006), which, taken together with the lexical-level effects from priming and the lack of negation/quantification effects, suggests that the mapping between the N400 and sentence-level compositional semantics is anything but straightforward.

The P600

The P600 (alternatively the "syntactic positive shift") is a positive-going deflection that is generally largest over centro-parietal electrode sites and tends to occur 500-800ms post-stimulus onset (although there is a good deal of variability in the latency in the ERP literature). Like the LAN, the P600 has been reported for a broad array of syntactic violations, in many cases co-occurring with a preceding LAN. For example, P600's have been elicited to phrase structure violations (Hagoort, Brown, and Groothusen 1993, Friederici et al. 1993, Hahne and Friederici 1999, Friederici and Frisch 2000, Osterhout and Holcomb 1992), agreement violations (Hagoort, Brown, and Groothusen 1993, Kaan 2002), syntactic garden-paths (Friederici et al. 1996, Kaan and Swaab 2003, Osterhout, Holcomb, and Swinney 1994), and island violations (McKinnon and Osterhout 1996). The sheer number of violation types that elicit a P600 has led some researchers to suggest that the P600 may be a (slightly delayed) version of the P300 (or P3b), which is a general response to unexpected stimuli (Coulson et al. 1998, see Osterhout and Hagoort 1999 for a response). P600's have also been elicited by the processing of grammatical sentences with particularly complex syntactic properties, such as ambiguous structures (Frisch, Schlesewsky, Saddy, and Alpermann 2002) and wh-movement (Fiebach, Schlesewsky, and Friederici 2002, Kaan, Harris, Gibson, and Holcomb 2000, Phillips, Kazanina, and Abada 2005). Current research on the P600 has focused on cases of unexpected theta-role assignment (Kim and Osterhout 2005, Kuperberg, Sitnikova, Caplan, and Holcomb 2003, van Herten, Kolk, and Chwilla 2005, Kuperberg 2007, Bornkessel-Schlesewsky and Schlesewsky 2008, Stroud and Phillips 2011), which we will discuss in more detail in the next section as these P600's have been interpreted as evidence of an independent semantic processing stream.

3.2 The role of syntactic theory in the interpretation of ERPs

As the brief review above makes clear, ERPs do not map cleanly to single parsing operations, but rather seem to track macroscopic classes of violations or processing difficulties. Consequently, there is quite a bit of variation in the literature when it comes to the interpretation of ERPs. It is not our intention to review the entire ERP literature in this chapter, but rather to illustrate how the mentalistic (architectural) commitments of syntactic theory play a key role in the interpretation of electrophysiological responses. Before continuing it should be noted that for convenience throughout the chapter we are going to use terminology such as the *generative framework* or the *unification framework* to refer to the mentalistic commitments of the respective syntactic theories. The fact that we describe certain interpretations of electrophysiological (and later hemodynamic) responses as using the *generative framework* does not necessarily mean that the cited authors explicitly endorse generative syntactic theories. Instead, we simply mean that the architectural assumptions that they assume for human sentence processing system are consistent with the mentalistic commitments of generative syntactic theories. Because syntactic theories are not themselves theories of sentence processing, it is not the case that a set of architectural assumptions is uniquely specified by a single syntactic theory. Therefore it is possible that some of the authors have other syntactic theories in mind that also lead to the same mentalistic commitments as generative theories.

ERPs and the Generative framework

Recall from section 2 that generative syntax is likely committed to a mental architecture in which (bundles of) features are stored in the lexicon and composed by a number of structure building operations that obey a set of complex syntactic rules. Generative syntax is also straightforwardly compatible with a parsing architecture in which the syntactic structure building precedes (compositional) semantics, as generative syntactic theories tend to be committed to an interpretive (compositional) semantics (Frazier 1978, 1987). Friederici (1995, 2002) has argued that the ERPs identified in the sentence processing literature are consistent with just such an architecture. Under this view, the ELAN reflects an initial stage in which syntactic structure is built according to syntactic rules using only the word category information of the incoming word (i.e, no semantic information is used to direct the structural parse at this stage). Friederici argues that three properties of the ELAN point to such a function: (i) it is extremely early in the parse, suggesting that only partial information could be available at that time, (ii) it is only elicited by phrase structure violations, suggesting that it is specific to syntactic structure building, and (iii) it is not affected by any task-level factors (such as the likelihood of encountering a violation in the experiment), suggesting that it is a relatively automatic process. Because the LAN is elicited by many different kinds of morphosyntactic violations, this theory interprets the LAN as an index of the second half of the initial syntactic stage where other morphosyntactic properties are established or checked, such as agreement and case marking. The N400 reflects a second stage where lexical semantic information is processed and argument relations are established. Finally the P600 represents a third stage of processing in which the syntactic and semantic information is integrated into a single representation, and any mistakes are rectified by reanalysis (which explains the fact that the P600 is elicited by both syntactic violations and garden-path sentences, and can be affected by the likelihood of encountering a violation in the experiment).

One of the defining characteristics of the three stage model proposed by Friederici (1995, 2002) is that it temporally orders syntactic processing (stage 1) and semantic processing (stage 2), before finally integrating both types of information in stage 3. This is often called a *syntax-first* parsing architecture (Frazier 1978, 1987, Frazier and Rayner 1982). The primary evidence for this temporal ordering comes from the relative ordering of the ERPs themselves: the ELAN occurs 150-250ms post-stimulus, the N400 occurs 300-500ms post-stimulus, and the P600 occurs 500-800ms post-stimulus. However, some studies have been designed to directly test the syntax-first hypothesis. For example, Friederici et al. 2004 used a sentence that violates both syntactic and semantic constraints as a way to test the syntax-first view:

(4) Das Buch wurde trotz *verpflanzt* von einem Verleger, den wenige empfahlen.
The book was despite replanted by a publisher, who(m) few recommended

The critical word *verpflanzt/replanted* violates both the phrase structure rules of German, and the selectional requirements of the subject *Buch/book*. Friederici et al. (2004) argued that a syntax-first architecture predicts an (E)LAN-P600 response to the syntactic violation caused by the verb, but no N400 because the syntactic integration failure would mean that the verb cannot be (subsequently) semantically integrated into the structure. On the other hand, the elicitation of an

N400 would suggest that semantic integration does not require successful syntactic integration. They found a LAN and a P600 but no N400.

The double violation paradigm was also used by Hagoort (2003) to determine the type of interaction between syntactic violations and semantic violations during processing. In this case, Dutch NPs containing an agreement error (between the determiner and the NP) and a selection error (between an adjective and the NP) were compared to sentences containing only one of the errors:

(5) Het / De zoute / bekwame *vaklieden* zien kwaliteit van het produkt.
The_{sg} / The_{pl} salty / skilled craftsmen_{pl} appreciate the quality of the product.

On the one hand, Hagoort found that the N400 to the selection violation was larger for the double violation than for the selection violation in isolation. This suggests that syntax and semantics interact in a way that boosts the N400 effect. On the other hand, Hagoort also found that the P600 to the agreement violation did not change between the double violation and the agreement violation in isolation, suggesting that the semantic violation has no impact on the P600 response to the syntactic violation. This asymmetry in the interaction (syntax affects semantics, but semantics does not affect syntax) is consistent with a syntax-first view in which syntax functionally precedes syntax.

ERPs and the unification framework

Recall that under the unification framework, words are stored in the lexicon as part of a structural frame that contains the syntactic environment(s) for that word. Parsing then consists of a single combinatorial operation called unification that joins two structural frames together (Hagoort 2003, 2005). In cases of ambiguity, the assumption is that two (or more) frames will be accessed from the lexicon, and a process of lateral inhibition based on various factors (plausibility, frequency, temporal decay, etc) will yield a single phrasal configuration. This model also assumes that syntactic and semantic processing is highly *interactive*: unification takes place at the syntactic, semantic, and phonological levels simultaneously, and both syntactic and semantic information will interact to produce the correct phrasal unification.[2] Given these properties, the unification architecture requires very different interpretations of the major ERP components. For example, the interpretation of the ELAN as an index of a purely syntactic structure building process, and the interpretation of the P600 as an index of structural reanalysis and syntactic-semantic integration, are both meaningless within a highly interactive architecture.

Within the unification framework, the three "syntactic" ERP components are not interpreted as indexing distinct processes or stages of processing, but rather are interpreted as indexing different aspects of the unification process itself. Under this view, the ELAN is a response to impossible unification such as when there are no nodes that can be combined between two structural frames, as is the case with the assumed structure of the phrase structure

[2] Highly interactive architectures are sometimes also known as *constraint-satisfaction* architectures to make reference to the fact that both syntactic constraints and semantic constraints are just constraints that need to be satisfied in whatever order possible (Altmann & Steedman 1988, Levy 2008, MacDonald, Pearlmutter, and Seidenberg 1994, Tanenhaus, Spivey-Knowlton, Eberhard, and Sedivy 1995).

violations that have historically elicited an ELAN. The LAN is also an index of failed unification, but whereas the ELAN is a response to impossible structural unification, the LAN is a response to morphosyntactic mismatches (e.g., agreement) that occur after two syntactic frames have been combined. The other major syntactic component, the P600, is not viewed as a response to impossible or invalid unification, but rather an index of the difficulty of the unification, which explains the fact that the P600 is elicited by various types of grammatical sentences that are either structurally ambiguous or syntactically complex. Finally, the N400 is interpreted in this framework as an index of the (lexical-)semantic unification that is assumed to occur in parallel with syntactic unification (Hagoort 2003).

As the interpretation of the ERPs above suggests, the interaction of syntactic and semantic processes plays a critical role in the architecture of the unification framework. Electrophysiological evidence for this type of interaction (and against a syntax-first architecture) was presented by van den Brink and Hagoort (2003) using the word-final morphology of Dutch verbs as a way to ensure that (lexical-)semantic information became available before word category information. For example, during auditory presentation of the Dutch word *kliederde* (*messed*), the information that this is the past tense verb 'to mess' as opposed to the noun *klierderboel*/*mess* is not available until the word-final passed tense morpheme *de* is encountered (approximately 300ms after the onset of the first phoneme). By embedding verbs with this property in a sentential context that creates both a syntactic and semantic violation, van den Brink and Hagoort were able to induce both an ELAN and an N400 to the two kinds of violations. However, in this case the ELAN actually occurred after the N400, as the syntactic violation could not be recognized until the word-final suffix *de* was encountered. Van den Brink and Hagoort argue that this demonstrates that syntactic integration does universally precede semantic integration, as the N400 (an index of difficult semantic integration) can precede the ELAN (an index of failed syntactic integration). Of course, this interpretation assumes that no word-category was assigned to the incomplete string *klierder* prior to the final morpheme *de*. If the category noun were assigned to the incomplete string prior to the final morpheme, perhaps due to the fact that the syntactic environment strongly predicts a noun, correct syntactic integration could occur prior to semantic integration.

Thematic P600s and the possibility of independent semantic processing

By focusing on the interpretation of ERPs within the generative and unification frameworks, we have seen both a syntax-first architecture in which syntactic processing is independent of semantic processing, and a highly interactive (or constraint-satisfaction) architecture in which syntactic and semantic processing are co-dependent. Recent work in the ERP literature has suggested that a third type of model may be possible in which semantic processing actually occurs independently of syntactic processing, leading to situations in which an ungrammatical structure associated with a plausible interpretation may be selected by the parser over a grammatical structure which is implausible. These models run contrary to the (likely) assumptions of many syntacticians, who assume that a syntactically licit structure will always be selected over an illicit structure, even if the licit structure leads to an implausible interpretation. These studies focus on "thematic P600" effects in which it appears that the comprehender has constructed a semantic interpretation that is not licensed by the syntactic structure of the sentence, thus eliciting a P600 (an index of syntactic violation, or of conflict between the syntactic and semantic representations) when the mismatch between the independent semantic interpretation and the syntactic structure is encountered. For example, Kim and Osterhout 2005 presented the following sentences to participants and recorded ERPs to the italicized verbs:

(6)		sentence	prediction	result
	a.	The hungry boy was *devouring* the cookies	control	
	b.	The dusty tabletop was *devouring* the kids	N400	N400
	c.	The hearty meal was *devouring* the kids	N400	P600

In the (a) sentence *the hungry boy* is an agentive subject of the active verb *devouring*. In the (b) sentence, *the dusty tabletop* is in subject position, but does not match the agent requirements of the active verb *devouring*. A familiar N400 effect results, apparently due to the mismatch between the lexical-semantic requirements of the verb and the subject. A similar scenario appears to be unfolding in the (c) sentence: *the hearty meal* is in the subject position, but it does not match the agent requirements of the active verb *devouring*. However, instead of eliciting an N400, the (c) sentence elicits a P600[3]. Similar "thematic P600s" have been found in English (Kim and Osterhout 2005, Kuperberg et al. 2003, Kuperberg, Caplan, Sitnikova, Eddy, and Holcomb 2006, Kuperberg, Kreher, Sitnikova, Caplan, and Holcomb 2007), Dutch (e.g., Kolk et al. 2003, van Herten, Chwilla, and Kolk, 2006, van Herten, Kolk, and Chwilla 2005), German (e.g., Friederici and Frisch 2000), and Spanish (Stroud and Phillips 2011).

Kim and Osterhout (2005) argue that the asymmetry between the ERP responses to (b) and (c) suggests that the semantic relationship (what they call semantic *attraction*) between *the hearty meal* and the verb *devour* is processed independently of the syntactic structure of the sentence. Specifically, they argue that at some point prior to the processing of the verbal morphology of *devouring*, the parser recognizes based on world knowledge that the NP *the hearty meal* is likely to be the theme of *devour* (but not the agent), and assigns that the thematic relationship. This sets up a syntactic prediction for passive verbal morphology on the verb *devour*. Therefore when the active verbal morphology is encountered, they argue that a syntactic violation occurs, eliciting a P600 (as opposed to a semantic violation which they assume would elicit an N400). Although data from languages with different word orders has led to several different accounts for these effects (see Bornkessel-Schlesewsky and Schlesewsky 2008, Kuperberg 2007, van Herten et al. 2005, 2006 for three other major approaches), they all appear to share the assumption that some form of semantic processing can occur independently of syntactic processing (but see Stroud and Phillips 2011 for a different view).

3.3 Synchronous neuronal oscillations

An underlying assumption of electrophysiological studies of neuronal oscillations is that when groups of neurons begin to work in concert, they will adopt a similar rate of oscillatory firing – something akin to a rhythm (Nunez and Srinivasan 2006). The various rates of oscillatory firing

[3] As discussed above, previous evidence suggests that the N400 may not be a reliable index of compositional semantic processes. Therefore, the N400 contrast between (a) and (b) may simply reflect the degree to which the theme primes/predicts the verb root (*tabletop – devour* vs. *meal - devour*), and thus may not be informative about whether or not the thematic mismatch is perceived in both cases. Therefore, it is not the absence of the N400 effect but the presence of the P600 effect that is hardest to account for under previous assumptions.

are conventionally divided into various frequency bands: activity between 0-4Hz is called *delta* activity, 4-7Hz is called *theta* activity, 8-12Hz is called *alpha* activity, 13-30Hz is called *beta* activity (and it sometimes subdivided into upper and lower), and 30Hz and above is called *gamma* activity. When a group of neurons begin to act in concert, the synchrony of their activity means that the EEG will show an increase in amplitude in the frequency band that the group of neurons has adopted. In this way, an increase in the amplitude of activity in a specific frequency band over a single scalp region suggests that a single group of neurons was recruited to perform a task related to the experimental event. Similarly, an increase in the oscillatory coherence at a specific frequency band between two or more *distinct* scalp regions suggests that multiple groups of neurons have been recruited by the experimental task. These two oscillatory responses – an increase in amplitude or an increase in coherence – provide a method for studying the local and long-distance synchrony of populations of neurons (Nunez and Srinivasan 2006).

To date, the majority of EEG studies of sentence processing has focused on ERPs, so it may be useful to briefly discuss the difference between ERPs and neuronal oscillations. As mentioned previously, ERPs are calculated by averaging the EEG activity of several trials using the event as time point zero. This averaging procedure means that only time-locked and phase-locked activity remains (which is called *evoked* activity). Neuronal oscillations are similarly time-locked to an event, but do not need to be phase-locked (which is called *induced* activity): two neurons can both adopt a synchronous rate of firing without the peaks and troughs of the waves necessarily lining up, just as two musicians can play a piece together at the same rhythm without necessarily creating notes at the same time. This means that oscillation activity can reveal patterns of activations that are very different from the ERP activity.

Although the vast majority of EEG studies of sentence processing have focused on ERPs, some recent research has started to look at neuronal oscillations to sentence processing events (Bastiaansen, van Berkum, and Hagoort 2002a, Bastiaansen, van Berkum, and Hagoort 2002b, Bastiaansen and Hagoort 2003, Bastiaansen et al. 2005, Hald et al. 2006). For example, Bastiaansen et al. (2002b) looked at two paradigms in Dutch that have been shown to elicit a P600 response: number agreement violations and gender agreement violations (van Berkum et al 2000).

(7) number agreement violation
Ik zag enkele donkere *wolk* aan de horizon
I saw several dark cloud on the horizon

(8) gender agreement violation
Ik zag een donker *wolk* aan de horizon.
I saw a dark$_{\text{NEUT}}$ cloud$_{\text{NON-NEUT}}$ on the horizon

They found local increases in amplitude in the theta band (4-7Hz) between 300-500ms post stimulus onset for both violation types; however, the scalp distribution of the maximum amplitude increase was different for each violation type: number agreement violations elicited a theta power increase over left anterior electrode sites, whereas gender agreement violations elicited a theta power increase over right anterior electrode sites. This result is interesting in a number of ways. First, it demonstrates a possible electrophysiological difference between number and gender agreement violations (although the precise functional interpretation of this difference remains unclear). Second, it demonstrates a difference between the latency of the ERP effect (500-800ms post stimulus onset) and the latency of the theta band increase (300-500ms). Third, it demonstrates a difference between the scalp distribution of the ERP (centro-parietal for both violation types) and the scalp distribution of the theta band increase (left and right anterior). And finally, it demonstrates an asymmetry in the types of differences between the two analyses: the ERP analyses revealed a quantitative difference between the two violations (the P600 to number violations was larger than the P600 to gender violations), whereas the oscillation analysis revealed a qualitative difference in scalp distribution. These results suggest that the two analyses yield different types of information about EEG activity during sentence processing (see also Hald et al. 2006 for similar results for the N400 to semantic violations).

4. Hemodynamic responses

As the previous section demonstrates, electrophysiological studies provide information about the types of processes that occur during syntactic processing, and about the timecourse of information flow (e.g., the relative ordering of syntactic and semantic processes). However, a complete cognitive neuroscience of syntactic processing also requires a theory of how those processes are implemented in the neural circuitry of the brain. The first step to a complete implementational theory is to identify the neural circuits involved in these processes. To be clear, this first "mapping" step is not itself an implementational theory; but it does constrain the search for how individual (or populations of) neurons implement individual processes. First, what we know about properties of particular brain regions may give us insight into how computations are implemented there. Although cellular organization is surprisingly similar throughout the brain (e.g. Mountcastle 1997), we know from many patient studies and neuroimaging studies that cognitive functions are often localized to particular regions, which may correspond to subtle cytoarchitectural differences (e.g. the Brodmann map – see Zilles and Amunts 2010 for recent review). While our understanding is still very limited, future research may show that these differences at the cellular level indeed constrain the kinds of computations that can be done in a given region. Second, knowledge about the anatomical connections that exist between different brain regions may help us to understand the larger circuits involved in sentence comprehension. We have rudimentary knowledge about the strongest of these connections from histology studies of gross anatomy, and new methods from functional neuroimaging such as diffusion imaging are continually improving our estimates of both structural and functional connectivity. Third, many brain regions seem to implement similar computations across domains; therefore if one localizes a syntactic computation to a region known to also be involved in another domain, it may provide clues to how the syntactic computation is implemented. Finally, in a practical sense, localization of a computation to a specific brain area can provide perhaps one of the most unambiguous ways of querying properties of that computation for future research, as compared with behavioral or scalp-recorded electrophysiological data that may reflect a mix of syntactic and more general processes.

Currently the most popular methodology for localizing cognitive function is measuring changes in blood flow related to neural activity with functional magnetic resonance imaging (fMRI). Neural firing depletes local energy stores, so in response, fresh blood carrying more glucose is delivered. The magnetic properties of fresh, oxygenated blood are different from old, deoxygenated blood, which is the source of the BOLD (Blood Oxygen-Level Dependent) signal that can be measured with MRI. If neural firing suddenly increases in a given region in response to a stimulus, the amount of oxygenated blood sent to this region should also increase (often referred to as the *hemodynamic response*), allowing the backwards inference that a region is involved in a computation if it demonstrates an increase in the BOLD signal when this computation is performed.

BOLD fMRI has a number of advantages as a localization methodology. It has much better spatial resolution than equally non-invasive techniques such as MEG and EEG, usually localizing the signal within 3-5 mm. The model for recovering the spatial coordinates of the MR signal is straightforward and universally agreed upon, in contrast to MEG or EEG source localization models that require many controversial assumptions. And although MRI scanning is expensive, MRI scanners are available all over the world, because of their clinical applications. However, BOLD fMRI also has well-known weaknesses for measuring cognitive processing. The most significant problem is the very poor temporal resolution of the technique, which is constrained more by the signal itself than by the measurement technology. As one might imagine, blood flows much more slowly than neurons fire; while electrophysiology suggests that basic processing stages take place on the scale of tens or hundreds of milliseconds, blood oxygenation changes on the scale of seconds. Therefore, if multiple subcomputations are sequentially engaged in the process of interest, a BOLD contrast isolating this process will show effects in all of the underlying regions in one image, as if they were all occurring simultaneously. It is left to the investigator to be sufficiently clever and lucky to be able to map the multiple regions that almost inevitably show up in a whole-brain contrast to the multiple hypothesized subcomponents of the computation.

All mainstream localization techniques (fMRI, MEG, PET) face difficulties in generalizing results across individuals, due to the large individual variation in brain size and morphology. Although algorithms for converting individual brains to a common space are constantly improving, these algorithms cannot solve the problem of variation in cytoarchitecture or 'functional' anatomy, e.g. if characteristic neuron types are not in exactly the same place with respect to larger brain structures. Regions of inferior frontal cortex historically implicated in language comprehension have been shown to have a particularly large amount of this form of variability (Amunts et al. 1999, Fedorenko and Kanwisher 2009).

Broca's Area

Our models of the cortical networks involved in the earliest stages of language processing such as speech or visual wordform recognition are still very crude. How can we even begin to approach questions about the cortical regions that process syntactic information? One simple approach is to manipulate the number or type of formal operations hypothesized by the syntactic theory in sentences presented to subjects during neuroimaging recordings. Even without committing to a specific parsing algorithm, we can start from the simple assumption that parsing representations that differ at the computational level will also differ in the operations required to parse them. Once we identify regions for which activity seems to correlate with our manipulation of structural properties of the sentence, we can then begin the more challenging task of determining what part of the parsing algorithm this region is implementing.

Broca's area is probably the most famous brain region to be correlated with structural properties of sentences. The term *Broca's area* usually refers to a portion of the *left inferior frontal gyrus* (LIFG) composed of the more anterior *pars triangularis* (Brodmann area 45) and the more posterior *pars opercularis* (Brodmann area 44). Paul Broca originally identified this area as central to speech processing based on the post-mortem inspection of the brains of two patients that exhibited severe aphasia: one patient could only produce the word "tan", the other only a handful of basic words. With the advent of non-invasive neuroimaging techniques such as PET and fMRI, Broca's area has taken center stage in the investigation of the neural substrates of syntactic processing. At least two thirds of the neuroimaging studies of the brain areas involved in sentence processing (in health) over the past 15 years reveal an increased activation in (at least part of) Broca's area for at least one of the reported contrasts, suggesting that this area indeed plays a significant role in some aspect of sentence processing.

Although there has been great debate about the key property that modulates activity in Broca's area in sentence processing, perhaps the most theory-neutral description of the central data is that Broca's area tends to respond more to sentences with non-canonical word order than sentences with canonical word order. For example, relative to controls with canonical word order, Broca's area shows increased BOLD signal for relative clauses (e.g., Just et al 1996, Ben-Shachar et al. 2003), wh-movement (e.g., Ben-Shachar 2004, Santi and Grodzinsky 2007), topicalization (e.g., Ben-Shachar 2004), clefting (e.g., Caplan et al. 1999), and scrambling (e.g., Friederici et al. 2006, Bornkessel et al. 2005, Bornkessel-Schlesewsky et al. 2009). The question then is which cognitive processes these comparisons have in common – a question that is inextricably tied to assumptions about syntactic theory. As the ongoing debate over the role of Broca's area illustrates, the interpretation of neuroimaging results at the implementational and algorithmic level is fundamentally shaped by the syntactic theory assumed at the computational level. Three prominent hypotheses regarding what structural property of sentences might account for variation in Broca's area activity are based in three different syntactic frameworks: movement (based on generative syntactic theory), unification (based on unification theory) and linearization (based on argument-dependency theory).

Movement Hypothesis

Grodzinsky and colleagues have argued that Broca's area seems to be more active for non-canonical word orders because Broca's area supports the syntactic mechanism of *movement* that is familiar from generative syntactic theory (Grodzinsky 1986, see Grodzinsky & Santi 2008 for a recent review). To demonstrate the specialization of Broca's area to the syntactic mechanism of movement, Grodzinsky and colleagues have conducted a series of experiments designed to tease apart movement from other possible functions of Broca's area. For example, Ben-Shachar et al. (2004) argue against the idea that Broca's area is sensitive to the syntactic complexity of non-canonical word orders in Hebrew by demonstrating increased activation for topicalization, which they assume is derived via A′-movement, compared to dative shift, which is a superficially similar displacement of an argument, but is either derived via A-movement, or simply base generated (see Ramchand, *this volume*):

(9) Dative shift:	John natan [la-professor me-oxford] ['et ha-sefer ha-'adom] __	
	John gave [to-the-professor from-oxford] [the-book the-red] __	
Topicalization:	[la-professor me-oxford] John natan ['et ha-sefer ha-'adom] __	
	[to-the-professor from-oxford] John gave [the-book the-red] __	
	'John gave the red book to the professor from oxford'	

Unification Hypothesis

Remember that the unification framework argues that the smallest unit of syntactic combination are words already associated with structural frames, and that syntactic combination proceeds by combining these small trees together, much as in tree-adjoining grammars. Assuming this view, Hagoort and colleagues have suggested that Broca's area is responsible for syntactic unification (Hagoort, 2005), with the prediction that factors that increase the difficulty of syntactic unification will also increase activation in Broca's area. For example, Snijders et al. (2009) compared sentences containing a word-class ambiguous word (either noun or verb) to sentences containing no ambiguous words, with the prediction that the word-class ambiguity will force the parser to consider to syntactic frames simultaneously, thus increasing the unification difficulty in the ambiguous condition. They found increased activation for the ambiguous condition in the posterior portion of LIFG. On this view, the reason for effects of non-canonical word order must be that such sentences in some way increase syntactic unification difficulty.

(10) Ambiguous: Zodra zij *bewijzen*$_{(N/V)}$ leveren kunnen we beginnen.
As-soon-as they evidence/prove provide can we start
'As soon as they provide evidence we can start'

Unambiguous: Zodra zij *kopij* leveren kunnen we beginnen.
As-soon-as they copy provide can we start
'As soon as they provide a copy we can start'

Linearization Hypothesis

As mentioned briefly in section 2, Bornkessel-Schlesewsky and colleagues have proposed the argument-dependency framework, which assumes a parsing stage in which the argument relations of the sentence are computed according to several prominence hierarchies that are familiar from typological research (e.g., the animacy hierarchy, the case hierarchy, the definiteness hierarchy) (Comrie 1989, Bornkessel and Schlesewsky 2006, Wolff et al. 2008). Based on this framework, several studies have argued that Broca's area in fact supports the "linearization" processes that map word order to argument structure according to these prominence hierarchies. For example, Chen et al 2006 manipulated the animacy of the relevant arguments in center-embedded object relative clauses:

(11) a. The golfer [that the lightning struck __] survived the accident
b. The wood [that the man chopped __] heated the cabin

They found increased activation in several areas of the LIFG (BA 47, 9, 6) for the (a) sentences in which the relativized object was more animate than the subject of the relative clause (see Grewe et al. 2006 for similar effects of animacy in German). Bornkessel-Schlesewsky et al. (2009) manipulated both the order of the NPs in German sentences (subject first (SO) or object first (OS)) and the referentiality of the NPs (proper names first (REF) versus common nouns first (NREF)) to assess whether Broca's area is sensitive to both word order manipulations and referentiality manipulations, and found a gradient activation response in BA 44 that seemed to follow the referentiality hierarchy: SO-REF < SO-NREF < OS-REF < OS-NREF (see also Bornkessel et al. 2005 and Grewe et al. 2005).

The role of linking hypotheses

The movement, unification, and linearization hypotheses all demonstrate the role that syntactic theory can play in the interpretation of brain activation. However, it is important to note that this role is predicated upon an often unstated assumption about the linking hypothesis between a given syntactic theory (the computational level), a given parsing theory (the algorithmic level), and the correct level of description of brain activation (the implementational level). For example, the movement hypothesis of Grodzinsky and colleagues assumes a relatively direct linking hypothesis between syntactic theory and brain areas, such that parsing theories do not play much of a role in the interpretation of brain activation. Grodzinsky has called this assumption the syntactotopic conjecture (Grodzinsky and Friederici 2006): formal mechanisms of syntactic theory are themselves neurologically distinct elements of linguistic knowledge that can be localized to distinct brain areas. Superficially, the unification framework of Hagoort (2003, 2005) also appears to assume a relatively direct linking between syntactic theory and brain areas. However, this direct linking between syntax and brain activation is in some ways an illusion created by the fact that there is a relatively direct linking between all three levels: the syntactic theory is based upon the grammatical operation of unification, the parsing theory is based upon the parsing process of unification, and brain activation is assumed to reflect aspects of the neural computation of unification. Of course, other linking hypotheses are logically possible. In this section, we would like to review one approach that has enjoyed widespread support in recent years, in which the syntactic configurations licensed by syntactic theory place demands on non-structure-building parsing processes such as working memory processes, and it is these working memory processes that are assumed to be driving the activation in Broca's area.

The idea that Broca's area may be involved in the working memory requirements of sentence processing is consistent with the pre-theoretical observation that non-canonical word order leads to activation in Broca's area. Much research in the sentence processing literature has focused on the role of working memory during the processing of non-canonical word orders (e.g., King and Just 1991, Just and Carpenter 1992, MacDonald, Just, and Carpenter 1992, Gibson 1998, Caplan and Waters 1999, Vos et al. 2001, Fiebach et al. 2002, Roberts and Gibson 2002, Phillips et al. 2005, Fedorenko et al. 2006, 2007). However, to truly dissociate the working memory hypothesis from those that postulate a more direct link between structural properties and brain regions, it is necessary to find cases that share the same structural properties but differ in working memory demands and vice versa. A number of studies have done this over the last 15 years, with mixed results. For example, Fiebach et al. (2005) used embedded object wh-questions in German to manipulate dependency distance (short versus long) without the confound of embedding type, and found increased activation in Broca's area (BA 44 and 45) for the longer dependencies, which would be predicted on a working memory view, but not on a view such as Grodzinsky's in which activity should only be modulated by differences in the presence or number of movements rather than their length.

In defense of the movement hypothesis, Santi and Grodzinsky (2007) have presented opposing evidence in a study in which they manipulated the distance between a moved wh-word and its gap (0 intervening NPs, 1 intervening NP, 2 intervening NPs), and compared that to the distance between an anaphor and its antecedent (0, 1, and 2 intervening NPs). They argued that both wh-movement and anaphoric binding dependencies require working memory resources to resolve, therefore the working memory interpretation of Broca's area should predict that both manipulations will increase activation. In contrast, they find that Broca's area is selectively sensitive to the wh-movement manipulation but not the binding manipulation.

(12) Distance manipulation

Move0: The mailman and the mother of Jim love **the woman** who Kate burnt __.
Move1: The mother of Jim loves **the woman** who the mailman and Kate burnt __.
Move2: Kate loves **the woman** who the mailman and the mother of Jim burnt __.

Bind0: The sister of Kim assumes that Anne loves **the mailman** who burnt **himself**
Bind1: The sister of Kim assumes that **the mailman** who loves Anne burnt **himself**
Bind2: Anne assumes that **the mailman** who loves the sister of Kim burnt **himself**.

Of course, Santi and Grodzinsky (2007, 2010) are careful to acknowledge that it is untenable to hold that Broca's area as classically defined is *only* involved in a very specific linguistic computation such as syntactic movement, because numerous neuroimaging studies have localized activity here for working memory tasks that do not involve sentences, such as the n-back task (Braver et al 1997, Smith and Jonides 1998, 1999) and semantic priming tasks (Gold et al. 2006). They argue instead that Broca's area may serve multiple functions across the various domains of cognition, but that its role in sentence processing is specifically related to movement, and crucially not other non-movement constructions (see also Makuuchi et al. 2009 for a different approach to the tension between generality and specificity in Broca's area).

Another approach to dissociating the working memory hypothesis from the others that has been proposed recently is to tax processes involved in working memory with an external manipulation to see if this affects the contrasts observed during sentence processing of non-canonical word orders. The logic of this approach is straightforward: if the difference observed between two sentences is due to different working memory requirements, then increasing the working memory requirements for both sentences with a non-linguistic working memory task should eliminate the activation difference between the two sentences (as both will show maximum activation). Rogalsky et al. (2008) asked participants to listen to sentences containing center-embedded subject relatives (a) and center-embedded object relatives (b) – a contrast that reliable activates Broca's area – while performing one of three concurrent tasks: (i) no concurrent task, (ii) whisper "ba da ga da" repeatedly, (iii) tap their fingers in sequence from thumb to pinky repeatedly. They found that object-relatives led to an increased activation in both BA 44 and 45 when there was no concurrent task as reported by previous studies. However, the concurrent articulatory task (ba da ga da) eliminated the activation in BA 44, leaving only activation in BA 45. Interestingly, they found a complementary pattern for the finger-tapping task: object-relatives led to increased activation in BA 44 but not in BA 45. Rogalsky et al. interpret these results as evidence that BA 44 supports the articulatory rehearsal component of working memory, and that non-canonical word orders like object relative clauses generate activation in Broca's area because they recruit articulatory rehearsal based working memory (such as silently repeating sentences). They also suggest that BA 45 may support domain general sequencing operations that are deployed during both sentence comprehension and sequential tasks like finger tapping. Clearly, the evidence gathered to date is consistent with multiple theories of the role of Broca's area in syntactic processing.

The temporal lobe

Although Broca's area has been the focus of many neuroimaging studies of syntax, there is a growing literature implicating portions of the temporal lobe in syntactic processing. Notably, different formal syntactic theories make different predictions about which part of temporal cortex is likely to play the central role in basic syntactic combination. Theories like unification theory and tree-adjoining grammar argue that much of syntactic structure is stored with lexical entries, and therefore that most of the important syntactic work is done by lexical retrieval. As we will see, it has been suggested that lexical retrieval can be localized to areas of the posterior temporal cortex. In contrast, generative syntactic theory assumes that syntactic structure is built up by 'merge'-type operations that combine simple lexical terminals, and therefore predict that the areas involved in syntactic combination will show more activity than areas involved in simple lexical retrieval. As we will see, several researchers have suggested that parts of the anterior temporal cortex are involved in syntactic combination.

Anterior temporal cortex

One of the most robust neuroimaging findings about sentence-level processing is that lateral anterior portions of the superior and middle temporal cortex show greater activation bilaterally for reading or listening to sentences than word lists (Mazoyer 1993, Stowe et al. 1998, Friederici et al. 2000, Vandenberghe et al. 2002, Humphries et al 2005, 2006, Brennan and Pylkkänen *submitted*). This pattern holds in at least some regions of the anterior temporal lobe (ATL) even with Jabberwocky sentences in which the content words are replaced with nonsense words (Humphries et al. 2006). Furthermore, voxel-based lesion mapping has associated damage to left lateral ATL with comprehension impairment for most sentences more complex than simple declaratives (Dronkers et al. 2004, although cf. Kho et al. 2008). These findings suggest that ATL supports sentence-level computations that do not rely on lexical semantics, but this leaves open a number of possible candidate processes: syntactic processes, argument structure processes, discourse processes, and even prosodic processes (although Humphries et al. 2005 show that the sentence > word list effect holds in some areas of ATL even when sentences are pronounced with list prosody instead of normal prosody).

If there were a brain region dedicated to basic syntactic phrase structure computation in comprehension (as opposed to the dependency structure building suggested for Broca's area), one would expect it to show a profile exactly like ATL, showing more activity for processing word strings with syntactic structure than those without. However, demonstrating that this area is specifically involved in syntax as opposed to other phrase-level computations has proved challenging. An interesting recent attempt comes from a study by Brennan et al. (2010), who used a naturalistic comprehension paradigm in which they asked participants to listen to a portion of *Alice in Wonderland* while recording fMRI. Brennan et al. counted the number of syntactic nodes being constructed at each time point in the story and found that this syntactic node count significantly correlated with the BOLD signal in left anterior temporal cortex (notably, this was not the case for Broca's area). This finding is thus consistent with a syntactic interpretation of ATL function; however, as the authors themselves point out, syntactic node count is likely to be correlated with other factors such as the number of compositional semantic operations required and therefore this evidence is not decisive.

A few studies have shown increased activity in ATL for phrase structure violations (Meyer et al. 2000, Friederici et al. 2003). For example, Friederici et al. (2003) have demonstrated that ATL is sensitive to the phrase structure violations that give rise to the ELAN in ERP research:

(13)	correct:	Das Hemd wurde gebügelt the shirt was ironed
	syntactic violation:	*Die Bluse wurde am gebügelt The blouse was on-the ironed

Similarly, Meyer et al. (2000) demonstrated that ATL is sensitive to other basic syntactic violations such as gender disagreement, number disagreement, and case-marking violations. However, the interpretation of these results really depends on our hypothesis about the processes invoked when the parser encounter a syntactic violation. If syntactic violations induce the parser to try additional phrase structure combinations, one might indeed expect more activity in the region involved in building phrase structure. Yet syntactic violations might also lead the parser to initiate non-structure-building processes with the goal of reanalysis/repair, for example working memory processes to recall the previous material, or semantic interpretation processes to get clues towards how the structure is repaired. Under this latter view, activation to syntactic violations in ATL may not index syntactic structure-building at all. Until we have a good theory of the algorithm used to deal with syntactic violations in comprehension, these data do not provide unambiguous evidence about the neural implementation.

Finally, in a novel attempt to get around the syntax/semantics confound in sentence-level processing, Rogalsky and Hickok (2008) used a selective attention paradigm to attempt to tease apart syntactic and semantic activation in ATL by telling subjects that their task was to either identify syntactic violations (syntactic attention) or identify semantic violations (semantic attention). The analysis was then performed only on correct sentences to avoid contamination from the error detection. They found that large portions of ATL were activated equally for both syntactic and semantic attention, with a small region activated more by semantic attention. The authors interpret these results as evidence that ATL supports both syntactic and semantic processing. However, although no violations were included in the analysis, this study shares some of the problems of interpretation of the violation studies: there is not yet a well-worked out theory of how processing changes when subjects are told to direct their attention to syntactic or semantic violations. If typical sentence comprehension processes are not under attentional control, then it would not be surprising that ATL is similarly activated (relative to word lists) in both cases, nor would it be informative about whether ATL is involved in syntactic or semantic computations.

Although the majority of neuroimaging studies of syntactic processing have focused on Broca's area, it is becoming increasingly obvious that ATL plays an important role in some aspect(s) of sentence-level processing. However, progress in this area is going to require specific theories about the types of syntactic and semantic processes that are incrementally deployed during sentence processing, which we have already seen are highly dependent upon the assumptions of syntactic theory.

Posterior temporal cortex

Within the unification framework there are two primary parsing operations: lexical retrieval of lexically-bound syntactic frames, and unification of these frames into sentence structures (Hagoort 2005). Unification is argued to occur in left IFG, which is thought to account for the effects of syntactic structure observed in Broca's area, as we discuss above. However, unlike in traditional generative theories, much of syntactic structure is thought to be retrieved with the lexical entry rather than compiled online, and therefore, the other critical operation for syntactic processing in the unification framework is memory retrieval of stored lexical representations. A large body of functional neuroimaging evidence suggests that retrieval of lexical representations across both comprehension and production is supported by posterior temporal cortex, in particular the region encompassing mid-posterior middle temporal gyrus (MTG) and parts of the neighboring superior temporal sulcus (STS) and inferior temporal cortex (IT) (see Indefrey and Levelt 2004, Hickok and Poeppel 2004, Martin 2007, Lau et al. 2008 for review). In addition, neuropsychological studies show that lesions to posterior temporal cortex are associated with difficulty in comprehension of single words (Hart and Gordon 1990) and even the simplest sentences (Dronkers et al. 2004).[4]

Since the unification framework posits that parsing involves retrieving the syntactic frame stored with each lexical representation, it predicts that manipulations that affect this syntactic frame retrieval should change activity in posterior temporal cortex. In one recent study, Snijders et al. (2009) used syntactic category-ambiguous words presented in word lists and sentences (e.g. *duck*, which has both a noun and a verb meaning) to operationalize lexical retrieval, as category-ambiguous words require (at least) two frames to be retrieved from the lexicon, whereas unambiguous words only require one. They predicted that lexical retrieval would be more taxed for ambiguous than unambiguous words in both word lists and sentences (because two syntactic frames must be retrieved instead of one) but that because word lists do not require syntactic unification, only ambiguous words in sentences would be associated with increased syntactic unification difficulty. Consistent with this prediction, Snijders et al. found a main effect of ambiguity for left posterior middle temporal gyrus (LpMTG) in both word lists and sentence contexts, while a left inferior frontal gyrus (LIFG) region of interest only showed an ambiguity effect for sentence contexts. However, alternative interpretations of this pattern are possible; for example the increase in LpMTG may reflect the cost of retrieving multiple conceptual representations rather than multiple syntactic frames, and the increase in LIFG for sentences may reflect the resolution of lexical competition made possible by the sentence context (e.g. Bedny et al. 2008, Grindrod et al. 2008) rather than syntactic unification. To more

[4] Note that an alternative view has suggested that lexical-semantic information is localized in an anterior temporal 'semantic hub' region, because semantic dementia is characterized by severe anterior temporal atrophy (Patterson et al. 2007). However, there is evidence that the atrophy extends to posterior parts of the temporal lobe as well (e.g. Mummery et al. 1999), and recent studies support the idea that this posterior temporal damage is the cause of the general lexical semantic deficits, while the anterior 'semantic hub' region may have a more specific function in representing living things or people (Noppeney et al. 2007, Simmons et al. 2010). Also see Tyler and Marslen-Wilson (2008) for arguments that posterior temporal cortex contains distinct syntactic and semantic subregions.

conclusively show that posterior temporal cortex is involved in the retrieval of syntactic frames, more work will be needed; one possibility might be to contrast the processing of category-ambiguous words such as *duck* with words whose two meanings share the same syntactic frame such as *bank*, or with category-ambiguous words whose meanings are semantically similar such as *vote* (Lee and Federmeier 2006).

5. Conclusion: Toward a cognitive neuroscience of syntax

The preceding discussion is far from exhaustive, but we hope that it provides a coherent introduction to the major strands of research on syntactic processing in the human brain. We also hope that this chapter has also made it clear that syntactic theory, and the linking hypotheses between syntactic theory and sentence processing, plays a pivotal (but often unstated) role in the interpretation of electrophysiological and hemodynamic neuroimaging studies. This suggests that syntacticians may be well positioned to become central players in the progress toward a cognitive neuroscience of syntax, and help resolve some of the longstanding puzzles discussed in this chapter.

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

Allopenna, PD, JS Magnuson, and MK Tanenhaus. 1998. Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of memory and language* 38:419–439.

Altmann, G, and M Steedman. 1988. Interaction with context during human sentence processing* 1. *Cognition* 30:191–238.

Amunts, K., A Schleicher, U Burgel, H Mohlberg, H.B.M Uylings, and K Zilles. 1999. Broca's region revisited: Cytoarchitecture and intersubject variability. *Journal of Computational Neurology* 412:319–341.

Bastiaansen, M, JJA van Berkum, and P Hagoort. 2002a. Event-related theta power increases in the human eeg during online sentence processing. *Neuroscience Letters* 323:13–16.

Bastiaansen, M, JJA van Berkum, and P Hagoort. 2002b. Syntactic processing modulates the [theta] rhythm of the human eeg. *Neuroimage* 17:1479–1492.

Bastiaansen, M, and P Hagoort. 2003. Special issue event-induced theta responses as a window on the dynamics of memory. *Cortex* 39:967–992.

Bastiaansen, M, M Linden, M Keurs, T Dijkstra, and P Hagoort. 2005. Theta responses are involved in lexical-semantic retrieval during language processing. *Journal of cognitive neuroscience* 17:530–541. Closed class versus open class words.

Bedny, Marina, M Mcgill, and S. L Thompson-Schill. 2008. Semantic adaptation and competition during word comprehension. *Cerebral Cortex* 18:2574–2585.

Ben-Shachar, M, T Hendler, I Kahn, D Ben-Bashat, and Y Grodzinsky. 2003. The neural reality of syntactic transformations: Evidence from functional magnetic resonance imaging. *Psychological Science* 14:433–440. Object relatives versus embedded sentences.

Ben-Shachar, M, D Palti, and Y Grodzinsky. 2004. Neural correlates of syntactic movement: converging evidence from two fmri experiments. *Neuroimage* 21:1320–1336.

van Berkum, JJA, P Zwitserlood, P Hagoort, and CM Brown. 2003. When and how do listeners relate a sentence to the wider discourse? evidence from the n400 effect. *Cognitive Brain Research* 17:701–718.

Bornkessel, I, and M Schlesewsky. 2006. The extended argument dependency model: A neurocognitive approach to sentence comprehension across languages. *Psychological Review* 113:787–821.

Bornkessel, I, S Zysset, AD Friederici, DY Von Cramon, and M Schlesewsky. 2005. Who did what to whom? the neural basis of argument hierarchies during language comprehension. *Neuroimage* 26:221–233.

Bornkessel-Schlesewsky, Ina, and Matthias Schlesewsky. 2008. An alternative perspective on “semantic p600” effects in language comprehension. *Brain Research Reviews* 59:55–73. URL `http://dx.doi.org/10.1016/j.brainresrev.2008.05.003`.

Bornkessel-Schlesewsky, Ina, Matthias Schlesewsky, and D Yves von Cramon. 2009. Word order and broca’s region: Evidence for a supra-syntactic perspective. *Brain and Language* 111:125–139. URL `http://dx.doi.org/10.1016/j.bandl.2009.09.004`.

Brennan, Jonathan, Yuval Nir, Uri Hasson, Rafael Malach, David J Heeger, and Liina Pylkkänen. 2010. Syntactic structure building in the anterior temporal lobe during natural story listening. *Brain and Language* 1–11. URL `http://dx.doi.org/10.1016/j.bandl.2010.04.002`.

Brennan, Jonathan, and Liina Pylkkanen. submitted. The time-course and spatial distribution of brain activity associated with sentence processing 1–39.

Bresnan, J. 2001. *Lexical functional syntax*. Blackwell.

Brink, D, and P Hagoort. 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.

Brown, C, and P Hagoort. 1993. The processing nature of the n400: Evidence fiom masked priming. *Journal of cognitive neuroscience* 5:34–44.

Caplan, D, N Alpert, and G Waters. 1998. Effects of syntactic structure and propositional number on patterns of regional cerebral blood flow. *Journal of cognitive neuroscience* 10:541–552.

Caplan, D, N Alpert, and G Waters. 1999. Pet studies of syntactic processing with auditory sentence presentation* 1. *Neuroimage* 9:343–351.

Caplan, D, N Alpert, G Waters, and A Olivieri. 2000. Activation of broca’s area by syntactic processing under conditions of concurrent articulation. *Human Brain Mapping* 9:65–71.

Caplan, D, E Chen, and G Waters. 2008a. Task-dependent and task-independent neurovascular responses to syntactic processing. *Cortex* 44:257–275.

Caplan, D, L Stanczak, and G Waters. 2008b. Syntactic and thematic constraint effects on blood oxygenation level dependent signal correlates of comprehension of relative clauses. *Journal of cognitive neuroscience* 20:643–656. Subject relatives versus object relatives.

Caplan, D, S Vijayan, G Kuperberg, C West, G Waters, D Greve, and AM Dale. 2002. Vascular responses to syntactic processing: Event-related fmri study of relative clauses. *Human Brain Mapping* 15:26–38. Subject versus object relatives.

Chen, E, WC West, G Waters, and D Caplan. 2006. Determinants of bold signal correlates of processing object-extracted relative clauses. *Cortex* 42:591–604.

Chomsky, N. 1995. *The minimalist program*. MIT Press.

Chomsky, N, and G Miller. 1963. *Handbook of mathematical psychology*, chapter Introduction to the formal analysis of natural languages, 269–321. Wiley.

Cooke, A, EB Zurif, C DeVita, D Alsop, P Koenig, J Detre, J Gee, M Pinango, J Balogh, and M Grossman. 2002. Neural basis for sentence comprehension: Grammatical and short-term memory components. *Human Brain Mapping* 15:80–94. Subject versus object relatives and distance.

Coulson, Seana, Jonathan King, and SCJWKM Kutas. 1998. Expect the unexpected: Event-related brain response to morphosyntactic violations. *Language and cognitive processes* 13:21–58.

Crain, Stephen, and Janet Fodor. 1985. *Natural language parsing: psycholinguistic, computational, and theoretical perspectives*, chapter How can grammars help parsers?, 94–128. Cambridge University Press.

Croft, W, and D A Cruse. 2004. *Cognitive linguistics*. Cambridge University Press.

Dapretto, M, and SY Bookheimer. 1999. Form and content:: Dissociating syntax and semantics in sentence comprehension. *Neuron* 24:427–432.

Dikker, Suzanne, Hugh Rabagliati, and Liina Pylkkänen. 2009. Sensitivity to syntax in visual cortex. *Cognition* 110:293–321. URL `http://dx.doi.org/10.1016/j.cognition.2008.09.008`.

Dronkers, Nina F, David P Wilkins, Robert D Van Valin, Brenda B Redfern, and Jeri J Jaeger. 2004. Lesion analysis of the brain areas involved in language comprehension. *Cognition* 92:145–77.

Federmeier, KD, and M Kutas. 1999. A rose by any other name: Long-term memory structure and sentence processing. *Journal of Memory and Language* 41:469–495.

Fedorenko, E, and N Kanwisher. 2009. *Language and Linguistics Compass* 3:839–865.

Felser, C, H Clahsen, and TF Münte. 2003. Storage and integration in the processing of filler-gap dependencies: An erp study of topicalization and wh-movement in german. *Brain and Language* 87:345–354.

Fiebach, CJ, M Schlesewsky, and AD Friederici. 2002. Separating syntactic memory costs and syntactic integration costs during parsing: the processing of german wh-questions. *Journal of Memory and Language* 47:250–272.

Fiebach, CJ, M Schlesewsky, G Lohmann, DY Von Cramon, and AD Friederici. 2005. Revisiting the role of broca's area in sentence processing: Syntactic integration versus syntactic working memory. *Human Brain Mapping* 24:79–91.

Fischler, I, PA Bloom, DG Childers, SE Roucos, and NW Perry Jr. 1983. Brain potentials related to stages of sentence verification. *Psychophysiology* 20:400–409.

Fodor, JD, and L Frazier. 1908. Is the human sentence pasing mechanism an atn? *Cognition* 8:417–459.

Frazier, L. 1978. On comprehending sentences: syntactic parsing strategies. Doctoral Dissertation, University of Connecticut.

Frazier, L. 1987. *Attention and performance xii: The psychology of reading*, chapter Sentence processing: A tutorial review. Lawrence Erlbaum.

Frazier, L. 1990. *Cognitive models of speech processing*, chapter Exploring the architecture of the language-processing system. MIT Press.

Frazier, L, and JD Fodor. 1978. The sausage machine: A new two-stage parsing model. *Cognition* 6:291–325.

Frazier, L, and K Rayner. 1982. Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive psychology* 14:178–210.

Friederici, A, M Meyer, and D Cramon. 2000. Auditory language comprehension: An event-related fmri study on the .... *Brain and Language* URL `http://linkinghub.elsevier.com/retrieve/pii/S0093934X00924386`.

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

Friederici, AD. 2009. Pathways to language: fiber tracts in the human brain. *Trends in Cognitive Sciences* 13:175–181.

Friederici, AD, J Bahlmann, S Heim, RI Schubotz, and A Anwander. 2006a. The brain differentiates human and non-human grammars: functional localization and structural connectivity. *Proceedings of the National Academy of Sciences of the United States of America* 103:2458.

Friederici, AD, CJ Fiebach, M Schlesewsky, ID Bornkessel, and DY Von Cramon. 2006b. Processing linguistic complexity and grammaticality in the left frontal cortex. *Cerebral Cortex* 16:1709.

Friederici, AD, and S Frisch. 2000. Verb argument structure processing: The role of verb-specific and argument-specific information* 1,* 2. *Journal of Memory and Language* 43:476–507.

Friederici, AD, TC Gunter, A Hahne, and K Mauth. 2004. The relative timing of syntactic and semantic processes in sentence comprehension. *NeuroReport* 15:165.

Friederici, AD, A Hahne, and Axel Mecklinger. 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, AD, and SA Kotz. 2003. The brain basis of syntactic processes: functional imaging and lesion studies. *Neuroimage* 20:S8–S17.

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

Friederici, AD, SA Ruschemeyer, A Hahne, and CJ Fiebach. 2003. The role of left inferior frontal and superior temporal cortex in sentence comprehension: localizing syntactic and semantic processes. *Cerebral Cortex* 13:170.

Frisch, S, M Schlesewsky, D Saddy, and A Alpermann. 2002. The p600 as an indicator of syntactic ambiguity. *Cognition* 85:B83–B92.

Garnsey, SM, MK Tanenhaus, and RM Chapman. 1989. Evoked potentials and the study of sentence comprehension. *Journal of Psycholinguistic Research* 18:51–60.

Gold, Brian T, David A Balota, Sara J Jones, David K Powell, Charles D Smith, and Anders H Andersen. 2006. Dissociation of automatic and strategic lexical-semantics: functional magnetic resonance imaging evidence for differing roles of multiple frontotemporal regions. *Journal of Neuroscience* 26:6523–32.

Goldberg, A. 2006. *Constructions at work: the nature of generalization in language*. Oxford University Press.

Goldberg, Adele. 1995. *Constructions: a construction grammar approach to argument structure*. University of Chicago Press.

Grewe, T, I Bornkessel, S Zysset, R Wiese, D. Y von Cramon, and M Schlesewsky. 2005. *Human Brain Mapping* 26.

Grindrod, CM, NY Bilenko, EB Myers, and Sheila E Blumstein. 2008. The role of the left inferior frontal gyrus in implicit semantic competition and selection: An event-related fmri study. *Brain Research* 1229:167–178.

Grodzinsky, Y. 1986. Language deficits and the theory of syntax. *Brain and Language* 27:135–159.

Grodzinsky, Y. 2010. The picture of the linguistic brain: How sharp can it be? reply to fedorenko & kanwisher. *Language and Linguistics Compass* 1–18.

Grodzinsky, Y, and AD Friederici. 2006. Neuroimaging of syntax and syntactic processing. *Current Opinion in Neurobiology* 16:240–246.

Grodzinsky, Y, and A Santi. 2008. The battle for broca's region. *Trends in Cognitive Sciences* 12:474–480.

Gunter, T, LA Stowe, and G MULDER. 1997. When syntax meets semantics. *Psychophysiology* 34:660–676.

Hagoort, P. 2003a. How the brain solves the binding problem for language: a neurocomputational model of syntactic processing. *Neuroimage* 20:S18–S29.

Hagoort, P. 2003b. 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. 2005. On broca, brain, and binding: a new framework. *Trends in Cognitive Sciences* 9:416–423.

Hagoort, P. 2008. The fractionation of spoken language understanding by measuring electrical and magnetic brain signals. *Philosophical Transactions of the Royal Society B: Biological Sciences* 363:1055.

Hagoort, P, C Brown, and J Groothusen. 1993. The syntactic positive shift (sps) as an erp measure of syntactic processing. *Language and cognitive processes* 8:439–483.

Hagoort, P, M Wassenaar, and CM Brown. 2003. Syntax-related erp-effects in dutch. *Cognitive Brain Research* 16:38–50.

Hahne, A, and AD Friederici. 1999. Electrophysiological evidence for two steps in syntactic analysis: Early automatic and late controlled processes. *Journal of cognitive neuroscience* 11:194–205.

Hahne, A, and AD Friederici. 2002. Differential task effects on semantic and syntactic processes as revealed by erps. *Cognitive Brain Research* 13:339–356.

Hald, LA, M Bastiaansen, and P Hagoort. 2006. Eeg theta and gamma responses to semantic violations in online sentence processing. *Brain and Language* 96:90–105. Good description of time-frequency methods.

Halle, M, and A Marantz. 1993. *The view from building 20*, chapter Distributed Morphology and the Pieces of Inflection, 111–176. MIT Press.

Heim, S, SB Eickhoff, and K Amunts. 2008. Specialisation in broca's region for semantic, phonological, and syntactic fluency? *Neuroimage* 40:1362–1368. Spontaneous generation tasks???

Herten, M Van, DJ Chwilla, and HHJ Kolk. 2006. When heuristics clash with parsing routines: Erp evidence for conflict monitoring in sentence perception. *Journal of cognitive neuroscience* 18:1181–1197.

Herten, M Van, HHJ Kolk, and DJ Chwilla. 2005. An erp study of p600 effects elicited by semantic anomalies. *Cognitive Brain Research* 22:241–255.

Hickok, G, and David Poeppel. 2004. Dorsal and ventral streams: a framework for understanding aspects of the functional anatomy of language. *Cognition* 92:67–99.

Hinojosa, J, M Martín-Loeches, P Casado, F Munoz, and F Rubia. 2003. Similarities and differences between phrase structure and morphosyntactic violations in spanish: An event-related potentials study. *Language and Cognitive Processes* 18:113–142.

Humphries, C, JR Binder, DA Medler, and E Liebenthal. 2006. Syntactic and semantic modulation of neural activity during auditory sentence comprehension. *Journal of cognitive neuroscience* 18:665–679.

Humphries, C, T Love, D Swinney, and G Hickok. 2005. Response of anterior temporal cortex to syntactic and prosodic manipulations during sentence processing. *Human Brain Mapping* 26:128–138.

Humphries, C, K Willard, B Buchsbaum, and G Hickok. 2001. Role of anterior temporal cortex in auditory sentence comprehension: an fmri study. *NeuroReport* 12:1749. Compared environmental sounds to sentences.

Indefrey, P, and WJM Levelt. 2004. The spatial and temporal signatures of word production components. *Cognition* 92:101–144.

Isel, F, A Hahne, B Maess, and A Friederici. 2007. Neurodynamics of sentence interpretation: Erp evidence from french. *Biological Psychology* 74:337–346.

Jackendoff, R. 1999. *The neurocognition of language*, chapter The representational structures of the language faculty and their interactions. Oxford University Press.

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

Joshi, Aravind, and Yves Schabes. 1997. *Handbook of formal language*, chapter Tree-adjoining grammars, 69–123. Springer.

Jr., John Hart, and Barry Gordon. 1990. Delineation of single-word semantic comprehension deficits in aphasia, with anatomical correlation. *Annals of Neurology* 1–7.

Just, MA, PA Carpenter, TA Keller, WF Eddy, and KR Thulborn. 1996. Brain activation modulated by sentence comprehension. *Science* 274:114. Subject versus object relatives.

Kaan, E. 2002. Investigating the effects of distance and number interference in processing subject-verb dependencies: An erp study. *Journal of Psycholinguistic Research* 31:165–193.

Kaan, E, A Harris, E Gibson, and P Holcomb. 2000. The p600 as an index of syntactic integration difficulty. *Language and cognitive processes* 15:159–201.

Kaan, E, and TY Swaab. 2002. The brain circuitry of syntactic comprehension. *Trends in Cognitive Sciences* 6:350–356.

Kaan, E, and TY Swaab. 2003. Electrophysiological evidence for serial sentence processing: a comparison between non-preferred and ungrammatical continuations* 1. *Cognitive Brain Research* 17:621–635.

Kho, KH, P Indefrey, P Hagoort, CWM van Veelen, PC van Rijen, and NF Ramsey. 2008. Unimpaired sentence comprehension after anterior temporal cortex resection. *Neuropsychologia* 46:1170–1178.

Kim, A, and L Osterhout. 2005. The independence of combinatory semantic processing: Evidence from event-related potentials. *Journal of Memory and Language* 52:205–225.

Kluender, R, and M Kutas. 1993a. Bridging the gap: Evidence from erps on the processing of unbounded dependencies. *Journal of cognitive neuroscience* 5:196–214.

Kluender, R, and M Kutas. 1993b. Subjacency as a processing phenomenon. *Language and cognitive processes* 8:573–633.

Kolk, HHJ, DJ Chwilla, M Van Herten, and PJW Oor. 2003. Structure and limited capacity in verbal working memory: A study with event-related potentials. *Brain and Language* 85:1–36.

Kounios, John, and PJ Holcomb. 1992. Structure and process in semantic memory: evidence from event-related brain potentials and reaction times. *Journal of Experimental Psychology: General* 121:459–479.

Kuperberg, G, D Caplan, T Sitnikova, M Eddy, and P Holcomb. 2006. Neural correlates of processing syntactic, semantic, and thematic relationships in sentences. *Language and cognitive processes* 21:489–530.

Kuperberg, GR. 2007. Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research* 1146:23–49.

Kuperberg, GR, DA Kreher, T Sitnikova, DN Caplan, and PJ Holcomb. 2007. The role of animacy and thematic relationships in processing active english sentences: Evidence from event-related potentials. *Brain and Language* 100:223–237.

Kuperberg, GR, T Sitnikova, D Caplan, and PJ Holcomb. 2003. Electrophysiological distinctions in processing conceptual relationships within simple sentences. *Cognitive Brain Research* 17:117–129.

Kutas, M, and KD Federmeier. 2000. Electrophysiology reveals semantic memory use in language comprehension. *Trends in Cognitive Sciences* 4:463–470.

Kutas, M, and S Hillyard. 1980. Event-related brain potentials to semantically inappropriate and surprisingly large words. *Biological Psychology* 11:99–116.

Kutas, Marta, C Van Petten, and R Kluender. 2006. Psycholinguistics electrified ii: 1994–2005. *Handbook of psycholinguistics* 659–724.

Lau, E, C Stroud, S Plesch, and C Phillips. 2006. The role of structural prediction in rapid syntactic analysis. *Brain and Language* 98:74–88.

Lau, EF, C Phillips, and D Poeppel. 2008. A cortical network for semantics:(de) constructing the n400. *Nature Reviews Neuroscience* 9:920–933.

Lau, Ellen, Diogo Almeida, Paul C Hines, and David Poeppel. 2009. A lexical basis for n400 context effects: Evidence from meg. *Brain and Language* 111:161–172. URL `http://dx.doi.org/10.1016/j.bandl.2009.08.007`.

Lee, C, and Kara D Federmeier. 2006. To mind the mind: An event-related potential study of word class and semantic ambiguity. *Brain Research* 1081:191–202.

Levy, Roger. 2008. Expectation-based syntactic comprehension. *Cognition* 106:1126–177.

Luck, Steven J. 2005. *An introduction to the event-related potential technique*. MIT Press.

MacDonald, Maryellen, Neal Pearlmutter, and Mark Seidenberg. 1994. Lexical nature of syntactic ambiguity resolution. *Psychological Review* 191:676–703.

Makuuchi, M, J Bahlmann, A Anwander, and AD Friederici. 2009. Segregating the core computational faculty of human language from working memory. *Proceedings of the National Academy of Sciences* 106:8362.

Marr, David. 1982. *Vision: A computational investigation into the human representation and processing of visual information*. MIT Press.

Martin, A. 2006. The representation of object concepts in the brain .

Mazoyer, BM, N Tzourio, V Frak, A Syrota, N Murayama, O Levrier, G Salamon, S Dehaene, Laurent Cohen, and J Mehler. 1993. The cortical representation of speech. *Journal of cognitive neuroscience* 5:467–479.

McKinnon, R. 1996. Constraints on movement phenomena in sentence processing: Evidence from event-related brain potentials. *Language and cognitive processes* 11:495–524.

McKoon, Gail, Roger Ratcliff, and Gregory Ward. 1994. Testing theories of language processing: An empirical investigation of the on-line lexical decision task. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:1219–1228.

Meyer, M, AD Friederici, and DY Von Cramon. 2000. Neurocognition of auditory sentence comprehension: event related fmri reveals sensitivity to syntactic violations and task demands. *Cognitive Brain Research* 9:19–33.

Mountcastle, Vernon B. 1997. The columnar organization of the neocortex. *Brain* 120:701–722.

Müller, O, and P Hagoort. 2006. Access to lexical information in language comprehension: Semantics before syntax. *Journal of cognitive neuroscience* 18:84–96.

Mummery, CJ, Karalyn Patterson, RJS Wise, R Vandenbergh, Cathy J Price, and JR Hodges. 1999. Disrupted temporal lobe connections in semantic dementia. *Brain* 122:61.

Münte, T F, and H J Heinze. 1994. *Cognitive electrophysiology*, chapter ERP negativities during syntactic processing of written words. Birkhauser.

Münte, TF, M Matzke, and S Johannes. 1997. Brain activity associated with syntactic incongruencies in words and pseudo-words. *Journal of cognitive neuroscience* 9:318–329.

Neville, H, JL Nicol, A Barss, KI Forster, and MF Garrett. 1991. Syntactically based sentence processing classes: Evidence from event-related brain potentials. *Journal of cognitive neuroscience* 3:151–165.

Nicol, J, and D Swinney. 1989. The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research* 18:5–19.

Nicol, JL, JD Fodor, and D Swinney. 1994. Using cross-modal lexical decision tasks to investigate sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:1–10.

Nieuwland, MS, and JJA Van Berkum. 2006. When peanuts fall in love: N400 evidence for the power of discourse. *Journal of cognitive neuroscience* 18:1098–1111.

Noppeney, U, Karalyn Patterson, LK Tyler, Helen E Moss, EA Stamatakis, P Bright, C Mummery, and Cathy J Price. 2007. Temporal lobe lesions and semantic impairment: a comparison of herpes simplex virus encephalitis and semantic dementia. *Brain* 130:1138.

Nunez, PL, and R Srinivasan. 2006. *Electric fields of the brain: The neurophysics of eeg*. Oxford University Press.

Osterhout, L. 1999. A superficial resemblance does not necessarily mean you are part of the family: Counterarguments to coulson, king and kutas (1998) in the p600/sps-p300 debate. *Language and Cognitive Processes* 14:1–14.

Osterhout, L, and PJ Holcomb. 1992. Event-related brain potentials elicited by syntactic anomaly* 1. *Journal of memory and language* 31:785–806.

Osterhout, L, PJ Holcomb, and D Swinney. 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 LA Mobley. 1995. Event-related brain potentials elicited by failure to agree. *Journal of Memory and Language* 34:739–773.

Osterhout, L, and J Nicol. 1999. On the distinctiveness, independence, and time course of the brain responses to syntactic and semantic anomalies. *Language and cognitive processes* 14:283–317.

Patterson, Karalyn, PJ Nestor, and TT Rogers. 2007. Where do you know what you know? the representation of semantic knowledge in the human brain. *Nat Rev Neurosci* 8:976–987.

van Petten, Cyma, Seana Coulson, Susan Rubin, Elena Plante, and Marjorie Parks. 1999. Time course of word identification and semantic integration in spoken language. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25:394–417.

Phillips, C. 2006. The real-time status of island phenomena. *Language* 82:795.

Phillips, C, N Kazanina, and SH Abada. 2005. Erp effects of the processing of syntactic long-distance dependencies. *Cognitive Brain Research* 22:407–428.

Phillips, C, and Shevaun Lewis. 2009. *Directions in derivations*, chapter Derivational order in syntax: Evidence and architectural consequences. Elsevier.

Pollard, C, and I Sag. 1994. *Head-driven phrase structure grammar*. University of Chicago Press.

Röder, B, O Stock, H Neville, S Bien, and F Rösler. 2002. Brain activation modulated by the comprehension of normal and pseudo-word sentences of different processing demands: a functional magnetic resonance imaging study. *Neuroimage* 15:1003–1014.

Rogalsky, C, and G Hickok. 2008. Selective attention to semantic and syntactic features modulates sentence processing networks in anterior temporal cortex. *Cerebral Cortex* 19:786–796.

Rogalsky, C, W Matchin, and G Hickok. 2008. Broca's area, sentence comprehension, and working memory: an fmri study. *Frontiers in Human Neuroscience* 2.

Rossi, S, MF Gugler, A Hahne, and AD Friederici. 2005. When word category information encounters morphosyntax: An erp study. *Neuroscience Letters* 384:228–233.

Santi, A, and Y Grodzinsky. 2007a. Taxing working memory with syntax: Bihemispheric modulations. *Human Brain Mapping* 28:1089–1097. More binding versus movment.

Santi, A, and Y Grodzinsky. 2007b. Working memory and syntax interact in broca's area. *Neuroimage* 37:8–17. Compares wh-movement and binding.

Santi, Andrea, and Yosef Grodzinsky. 2010. fmri adaptation dissociates syntactic complexity dimensions. *Neuroimage* 51:1285–1293. URL `http://dx.doi.org/10.1016/j.neuroimage.2010.03.034`.

Shetreet, Einat, Naama Friedmann, and Uri Hadar. 2009. An fmri study of syntactic layers: Sentential and lexical aspects of embedding. *Neuroimage* 48:707–716. URL `http://dx.doi.org/10.1016/j.neuroimage.2009.07.001`.

Simmons, W. K, M Reddish, P. S. F Bellgowan, and A Martin. 2010. The selectivity and functional connectivity of the anterior temporal lobes. *Cerebral Cortex* 20:813–825.

Snijders, T. M, T Vosse, G Kempen, J. J.A Van Berkum, K. M Petersson, and P Hagoort. 2009. Retrieval and unification of syntactic structure in sentence comprehension: an fmri study using word-category ambiguity. *Cerebral Cortex* 19:1493–1503.

Stowe, LA. 1986. Parsing wh-constructions: Evidence for on-line gap location. *Language and cognitive processes* 1:227–245.

Stowe, LA, CAJ Broere, AMJ Paans, AA Wijers, G Mulder, W Vaalburg, and F Zwarts. 1998. Localizing components of a complex task: sentence processing and working memory. *NeuroReport* 9:2995.

Stromswold, K, D Caplan, N Alpert, and S Rauch. 1996. Localization of syntactic comprehension by positron emission tomography. *Brain and Language* 52:452–473.

Stroud, C, and C Phillips. 2011. Examining the evidence for an independent semantic analyzer: An erp study in spanish. *Brain and Language* 118.

Sussman, RS, and J Sedivy. 2003. The time-course of processing syntactic dependencies: Evidence from eye movements. *Language and Cognitive Processes* 18:143–163.

Tanenhaus, MK, MJ Spivey-Knowlton, KM Eberhard, and JC Sedivy. 1995. Integration of visual and linguistic information in spoken language comprehension. *Science* 268:1632.

Townsend, D, and T Bever. 2001. *Sentence comprehension: The integration of habits and rules*. MIT Press.

Traxler, MJ, and MJ Pickering. 1996. Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language* 35:454–475.

Van Berkum, J. J. A., P. Zwitserlood, C. M. Brown, and P Hagoort. 2000. Processing gender and number agreement in parsing: An erp-based comparison. In *6th Conference on Architectures and Mechanisms of Language Processing*.

Van Valin, R, and R. La Polla. 1997. *Syntax: Structure, maening and function*. Cambridge University Press.

Vandenberghe, R, AC Nobre, and CJ Price. 2002. The response of left temporal cortex to sentences. *Journal of cognitive neuroscience* 14:550–560.

Vigneau, M, V Beaucousin, PY Herve, H Duffau, F Crivello, O Houde, B Mazoyer, and N Tzourio-Mazoyer. 2006. Meta-analyzing left hemisphere language areas: phonology, semantics, and sentence processing. *Neuroimage* 30:1414–1432.

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

Wagers, M, and C Phillips. 2009. Multiple dependencies and the role of the grammar in real-time comprehension. *Journal of Linguistics* 45:395–433.

Willems, RM, and P Hagoort. 2009. Broca's region: battles are not won by ignoring half of the facts. *Trends in Cognitive Sciences* 13:101. Just an angry letter.

Wolff, Schlesewsky M. Hirotani M. & Bornkessel-Schlesewsky I, S. 2008. The neural mechanisms of word order processing revisited: Electrophysiological evidence from japanese. *Brain and Language* 107.

Zilles, K, and K Amunts. 2010. Centenary of brodmann's map: conception and fate. *Nat Rev Neurosci* 11:139–145.