

Reflections on the Neurobiology of Syntax

Peter Hagoort

Abstract

This contribution focuses on the neural infrastructure for parsing and syntactic encoding. From an anatomical point of view, it is argued that Broca's area is an ill-conceived notion. Functionally, Broca's area and adjacent cortex (together *Broca's complex*) are relevant for language, but not exclusively for this domain of cognition. Its role can be characterized as providing the necessary infrastructure for unification (syntactic and semantic). A general proposal, but with the required level of computational detail, is discussed to account for the distribution of labor between different components of the language network in the brain. Arguments are provided for the *immediacy principle*, which denies a privileged status for syntax in sentence processing. The temporal profile of event-related brain potential (ERP) is suggested to require predictive processing. Finally, since, next to speed, diversity is a hallmark of human languages, the language readiness of the brain might not depend on a universal, dedicated neural machinery for syntax, but rather on a shaping of the neural infrastructure of more general cognitive systems (e.g., memory, unification) in a direction that made it optimally suited for the purpose of communication through language.

Introduction

Recent years have seen a growing number of studies on the neural architecture of language, using both electromagnetic methods (EEG, MEG) and hemodynamic methods (PET, fMRI). These studies have added to, but also changed, previous views on the brain's infrastructure for language, which were based primarily on patient studies. Before discussing the relevant issues in detail, here is what I believe to be the major conclusions that can be drawn from the overall body of literature on the neurobiology of language:

1. The language network in the brain is more extensive than the classical language areas (Broca's area, Wernicke's area). It includes, next to

Broca's area, adjacent areas in the left inferior frontal cortex (LIFC), as well as substantial parts of superior and middle temporal cortex, inferior parietal cortex, and parts of the basal ganglia. In addition, homologue areas in the right hemisphere are often found to be activated to a lesser extent.

2. In contrast to classical textbook wisdom, the division of labor between Broca's area (frontal cortex) and Wernicke's area (temporal cortex) is not language production vs. language comprehension. LIFC is strongly involved in syntactic and semantic unification operations during comprehension. Wernicke's area is involved in language production, at least at the level of word form encoding (Indefrey and Levelt 2004).
3. None of the language-relevant areas and none of the language-relevant neurophysiological effects are language specific. All language-relevant ERP effects (e.g., N400, P600, (E)LAN) seem to be triggered by other than language input as well (e.g., music, pictures, gestures; see Kaan, this volume).
4. For language, as for most other cognitive functions, the notion of function-to-structure mapping as being one-area-one-function is almost certainly incorrect. More likely, any cortical area is a node that participates in the function of more than one network. Conceivably, top-down connections from supramodal areas could differentially recruit such a cortical node into the service of one network or another (Mesulam 1990, 1998).
5. Despite the syntacto-centrist perspective in most of modern linguistics, in terms of processing there is no evidence for a privileged status of syntactic information. Language comprehension, beyond the single word level, happens in accordance with the immediacy principle, which states that all relevant information types (e.g., syntactic, semantic, extra-linguistic information) are brought to bear on language interpretation as soon as they become available, without giving priority, on principled grounds, to the syntax-constrained combination of lexical semantic information. The immediacy principle does not apply to language production, which requires a conceptual specification that precedes syntactic and phonological encoding, at least to some extent.

In this chapter, I provide further background and additional reflections on the biological foundations of syntax. I begin with a closer look at Broca's area, from a neuroanatomical perspective. The focus on Broca's area will enable some general conclusions about the relevant features of language-relevant brain areas.

Deconstructing Broca's Area

Classically, Broca's area has been considered to be a key site for syntax. Despite some disagreement in the literature (Uylings et al. 1999), most authors agree that Broca's area comprises Brodmann areas 44 and 45 of the left hemisphere. In classical textbooks, these areas coincide at the macroscopic level with the pars opercularis (BA 44) and the pars triangularis (BA 45) of the third frontal convolution. However, given anatomical variability, in many brains these two parts are not easy to identify (Uylings et al. 1999), and clear micro-anatomical differences (see Amunts and Zilles 2006) have been missed when macro-anatomical landmarks are used (Tomaiuolo et al. 1999). Furthermore, cytoarchitectonic analysis (Amunts et al. 2003) shows that BA 44 and BA 45 do not neatly coincide with the sulci that form their boundaries in macro-anatomical terms. More fundamentally, one has to question the justification for subsuming these two cytoarchitectonic areas under the overarching heading of Broca, rather than, say, BA 45 and BA 47. Areas 44 and 45 show a number of clear cytoarchitectonic differences, one of which is that BA 45 has a granular layer IV, whereas BA 44 is dysgranular. In contrast, like BA 45, BA 47 is part of the heteromodal component of the frontal lobe, known as the granular cortex (see Figure 13.1; Mesulam 2002). In addition, BA 44 and BA 45 have clearly distinct postnatal developmental trajectories and show a difference in their patterns of lateral asymmetry (Uylings et al. 1999). Using an observer-independent method for delineating cortical areas, Amunts et al. (1999) analyzed histological sections of ten human brains. They found a significant left-over-right asymmetry in cell density for BA 44, whereas no significant left-right differences were observed for BA 45. However, BA 44 and BA 45




Figure 13.1 Lateral view of the frontal lobes. The numbers refer to Brodmann areas. Hashed markings: motor-premotor cortex; dotted markings: heteromodal association cortex. SF: Sylvian fissure; CS: central sulcus. After Mesulam (2002).

are cytoarchitecturally more similar to each other than BA 44 and BA 6, or BA 45 and BA 6 (Amunts and Zilles 2001).

Studies on corresponding regions in the macaque brain (Petrides and Pandya 2002a) have shown that BA 44 receives projections primarily from somatosensory and motor-related regions, the rostral inferior parietal lobule, and supplementary and cingulate motor areas. There is input from portions of the ventral prefrontal cortex, but only sparse projections from inferior temporal cortex (Pandya et al. 1996). Conversely, BA 45 receives massive projections from most parts of prefrontal cortex, from auditory areas of the STG, and visually related areas in the posterior STS. In other words, the connectivity patterns of macaque BA 44 and 45 suggest clear functional differences between these areas. Differences in connectivity have also been found in human studies, using a technique called *diffusion tensor imaging* (Glasser and Rilling 2008). Friederici et al. (2006a) report syntax relevant differences in connectivity from Broca's area and from the frontal operculum to different parts of the temporal lobe.

Xiang, Norris, and Hagoort (submitted) performed a resting state functional connectivity study to investigate directly the functional correlations within the perisylvian language networks by seeding from three subregions of Broca's complex (*pars opercularis*, *pars triangularis*, and *pars orbitalis*) and their right hemisphere homologues. A clear topographical connectivity pattern in the left middle frontal, parietal, and temporal areas was revealed for the three left seeds in Broca's complex. These results demonstrate that a functional connectivity topology can be observed in the perisylvian language areas in the left hemisphere, in which different parts of Broca's area and adjacent cortex show a differential pattern of connectivity. This pattern is only seen in the left hemisphere and seems to be organized according to information type: semantic, syntactic, phonological (Figure 13.2).

Finally, studies on the receptor architecture of left inferior frontal areas indicate that functionally relevant subdivisions within BA 44 and BA 45 might be necessary (for more details, see Amunts and Zilles 2006). For instance, there is a difference within BA 44 of the receptor densities, for example of the 5HT₂ receptor for serotonin, with relatively low density in dorsal BA 44 and relatively high density in ventral BA 44.

In short, from a cytoarchitectonic and receptor architectonic point of view, Broca's area, comprising BA 44 and BA 45, is a heterogeneous patch of cortex; it is not a uniform cortical entity. The functional consequences of this heterogeneity are unclear, since the degree of anatomical uniformity required for an inference of functional unity is unknown. Here, two different views exist about the functional relevance of architectural differences in brain structure, which can be made clear in connection to the Brodmann map.

A prime example of the contribution of neuroanatomy is the famous map by Korbinian Brodmann (1869–1918). This map consists of 47 different areas, usually referred to by expressions such as BA 44 for Brodmann BA 44.

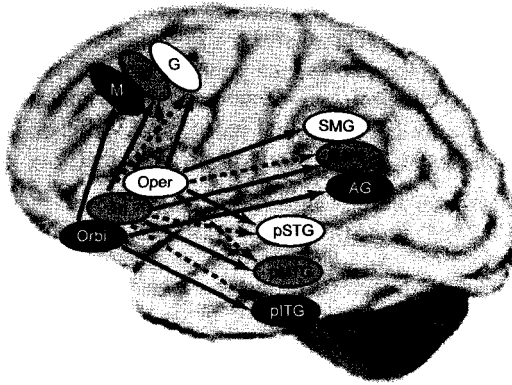


Figure 13.2 The topographical connectivity pattern in the perisylvian language networks. Connections to the left pars opercularis (oper), pars triangularis (tri) and pars orbitalis (orbi) are shown in white, gray, and black arrows respectively. The solid arrows represent the main (most significant) correlations and the dashed arrows represent the extending (overlapping) connections. Brain areas assumed to be mainly involved in phonological, syntactic and semantic processing are shown in white, gray, and black circles, respectively. MFG (middle frontal gyrus) is separated into three parts. SMG: Supramarginal gyrus; AG: Angular gyrus; S/IPL: the area between SMG and AG in the superior/inferior parietal lobule; pSTG: posterior superior temporal gyrus; pMTG: posterior middle temporal gyrus; pITG: posterior inferior temporal gyrus.

The numbers of the Brodmann areas were determined by the order in which Brodmann went through the brain, as he analyzed one area after the other. Brodmann's classification is based on the cytoarchitectonics of the brain, which refers to the structure, form, and position of the cells in the six layers of the cortex. Quantification was done by Brodmann on postmortem brains, which were sectioned into slices of 5–10 micron thickness that underwent Nissl-staining and were then inspected under the microscope. In this way, the distribution of different cell types across cortical layers and brain areas could be determined. Even today Brodmann's map (Brodmann 1909) is recognized as a hallmark in the history of neuroscience. Brodmann's work reveals that the composition of the six cortical layers, in terms of cell types, varies across the brain. Cell numbers can vary as well. The primary visual cortex, for instance, has about twice as many neurons per cortical column as other brain areas (Amaral 2000).

The classical view among neuroanatomists is that these architectural differences in brain structure are indicative of functional differences, and, conversely, that functional differences demand differences in architecture (Bartels and Zeki 2005; Brodmann 1905; Vogt and Vogt 1919; Von Economo and Koskinas 1925). Following the classical view, through different ways of characterizing brain structure (i.e., cyto- myelo- and receptor architectonics; Zilles and Palomero-Gallagher 2001), brain areas can be identified, for which differences in structural characteristics imply functional differences. Accordingly,

it follows that one should look for the structural features that determine why a particular brain area can support, for instance, the processing of a first or second language.

In contrast to the classical view in neuroanatomy, more recent accounts have argued that from a computational perspective, different brain areas are very similar. For instance, Douglas and Martin argue that:

The same basic laminar and tangential organization of the excitatory neurons of the neocortex, the spiny neurons, is evident wherever it has been sought. The inhibitory neurons similarly show a characteristic morphology and patterns of connections throughout the cortex....all things considered, many crucial aspects of morphology, laminar distribution, and synaptic targets are very well conserved between areas and between species (Douglas and Martin 2004, p. 439).

Functional differences between brain areas are, according to this perspective, due mainly to variability of the input signals in forming functional specializations. Domain specificity of a particular piece of cortex might thus not be determined so much by the heterogeneity of brain tissue, but rather by the way in which its functional characteristics are shaped through input. Moreover, findings of neuronal plasticity (e.g., the involvement of visual cortex in verbal memory in the congenitally blind; Amedi et al. 2003), suggest substantial plasticity in structure-to-function relations.

The above considerations result in a view of Broca's area that is different from the classical perspective. With respect to language areas in frontal cortex, it has become clear that, in addition to BA 44 and 45, at least BA 47 and the ventral part of BA 6 should be included in the left frontal language network. Recent neuroimaging studies indicate that the pars orbitalis of the third frontal convolution (roughly corresponding to BA 47) is involved in language processing (e.g., Devlin et al. 2003; Hagoort et al. 2004). From a functional anatomical perspective, it thus makes sense to use the term *Broca's complex* for this set of areas. Broca's complex is used here to distinguish it from Broca's area as classically defined, which is both too broad (since it comprises anatomically and functionally distinct areas, with differences in their connectivity patterns) and too narrow (since it leaves out adjacent areas that are shown to be crucial for language processing). Broca's complex, as defined here, is the set of anatomical areas in left inferior frontal cortex that are known to play a crucial, but by no means exclusive, role in language processing.

The Role of Broca's Complex

A particular cognitive function is most likely served by a distributed network of areas, rather than by one local area alone. In addition, a local area participates in more than one function. A one-to-one mapping between Broca's area and a specific functional component of the language system would thus be a highly unlikely outcome. Even for the visual system, it is claimed that the

representations of, for example, objects and faces in ventral temporal cortex are distributed and overlapping (Haxby et al. 2001). Moreover, Broca's area has been found activated in imaging studies on nonlanguage functions, such as action recognition (Decety et al. 1997; Hamzei et al. 2003) and movement preparation (Thoenissen et al. 2002). Of course, all this does not mean that cognitive functions are not localized and that the brain shows equipotentiality. It only means that the one-area-one-function principle is in many cases not an adequate account of how cognitive functions are neuronally instantiated.


If Broca's complex is not a domain-specific area, what properties does it have that makes it suitable for recruitment for unification operations in the language domain?

The answer that I propose is based on (a) an embedding of this complex in the overall functional architecture of prefrontal cortex, and (b) a general distinction between memory retrieval of linguistic information and combinatorial operations on information retrieved from the mental lexicon. These operations are referred to as *unification* or *binding*. The notion of binding is inspired by the visual neurosciences, where one of the fundamental questions concerns how we get from the processing of different visual features (color, form, motion) by neurons that are far apart in brain space to a unified visual percept. This is known as the *binding problem*. In the context of the language system, the binding problem refers to an analogous situation, but is now transferred to the time domain: How is information that is incrementally retrieved from the mental lexicon unified into a coherent overall interpretation of a multi-word utterance? Most likely, unification needs to take place at the conceptual, syntactic, and phonological level, as well as between these levels (see Figure 13.3; Jackendoff 2002). In this context, binding refers to a problem that the brain has to solve, not to a concept from a particular linguistic theory.


Integration is an important part of the function of the prefrontal cortex. This holds especially for integration of information in the time domain (Fuster 1995). To fulfill this role, prefrontal cortex needs to be able to hold information online (Mesulam 2002) and to select among competing alternatives (Thompson-Schill et al. 1999). Electrophysiological recordings in the macaque monkey have shown that this area is important for sustaining information triggered by a transient event for many seconds (Miller 2000). This allows prefrontal cortex to establish unifications between pieces of information that are perceived or retrieved from memory at different moments in time (Fuster 1995).

Recent neuroimaging studies indicate that Broca's complex contributes to the unification operations required for binding single word information into larger structures. In psycholinguistics, integration and unification refer to what is usually called post-lexical processing. These are the operations on information that is retrieved from the mental lexicon. It seems that prefrontal cortex is especially well suited to contribute to post-lexical processing, since this includes selection among competing unification possibilities, so that one unified representation spanning the whole utterance remains.


Phonological structure



Syntactic structure



Semantic/conceptual structure



Spatial structure



Figure 13.3 The phonological, syntactic, and semantic/conceptual structures for the sentence *The little star's beside the big star* (Jackendoff 2002). The unification operations involved are suggested to require the contribution of Broca's complex.

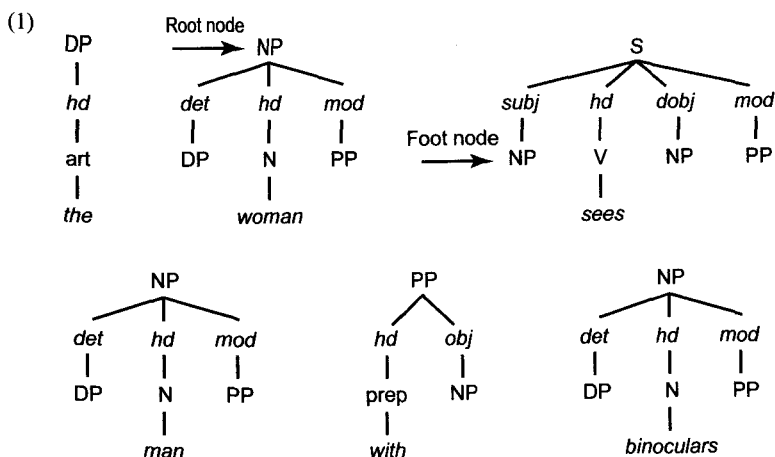
In short, the properties of neurons in the prefrontal cortex of macaques suggest that this part of the brain is suitable for integrating pieces of information that are made available sequentially, spread out over time, irrespective of the nature of the material to be handled (Owen et al. 1998). Clearly, there are interspecies differences in terms of the complexity of the information binding operations (Fitch and Hauser 2004), possibly supported by a corresponding increase in the amount of frontal neural tissue from monkey to man (Passingham 2002). With respect to language processing in humans, different

complex binding operations take place. Subregions in Broca's complex might contribute to different unification operations required for binding single word information into larger structures.

A Unification Model of Parsing

Progress in the neurobiology of language suffers from the lack of (or the unawareness about the availability of) detailed, preferably computationally explicit models of language processing. In many ways, this situation holds for most domains of language processing. In many ways, this situation holds for most domains of cognitive neuroscience research. If it comes to the neurobiology of syntax, the specification of a grammar does not suffice. If anything, all there is to be found in the brain is a system capable of parsing and syntactic encoding. Therefore, we need explicit models of parsing (and syntactic encoding) to guide our search and interpretation of results in ERP, MEG, or fMRI experiments. Here I offer a proposal for an explicit model of parsing. An interesting aspect of the model is that it also accounts for syntactic encoding, with the input–output relations reversed (from concept to phonology). This model is not to be taken as a final theoretical commitment. Instead, it is an illustration of the explicitness that is necessary to make progress in this domain of research. I offer an account based on the Unification Model for parsing (Vosse and Kempen 2000).

According to this model, each word form in the mental lexicon is associated with a structural frame. This structural frame consists of a three-tiered unordered tree, specifying the possible structural environment of the particular lexical item:

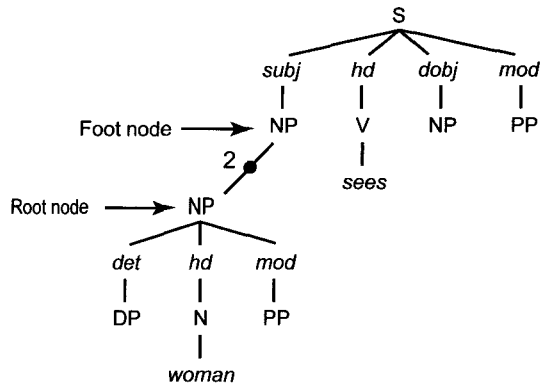


where DP: determiner phrase; NP: noun phrase; S: sentence; PP: prepositional phrase; art: article; hd: head; det: determiner; mod: modifier; subj: subject;

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

This parsing account is “lexicalist” in the sense that all syntactic nodes (e.g., S, NP, VP, N, V) are retrieved from the mental lexicon. In other words, chunks of syntactic structure are stored in memory. There are no syntactic rules that introduce additional nodes. In the online comprehension process, structural frames associated with the individual word forms incrementally enter the unification workspace. In this workspace constituent structures spanning the whole utterance are formed by a unification operation of two lexically specified syntactic frames:

(2)



Unification takes place by linking the root node NP to an available foot node of the same category. The number 2 indicates that this is the second link that is formed during online processing of the sentence *The woman sees the man with the binoculars*. This operation consists of linking up lexical frames with identical root and foot nodes, and checking agreement features (e.g., number, gender, person). It specifies what Jackendoff (2002) refers to as the only remaining “grammatical rule”: UNIFY PIECES.

The resulting unification links between lexical frames are formed dynamically, which implies that the strength of the unification links varies over time until a state of equilibrium is reached. Due to the inherent ambiguity in natural language, alternative binding candidates will usually be available at any point in the parsing process. That is, a particular root node (e.g., PP) often finds more than one matching foot node (i.e., PP) with which it can form a unification link (for examples, see Hagoort 2003).

Ultimately, this results in one phrasal configuration. This requires that among the alternative binding candidates, only one remains active. The required state of equilibrium is reached through a process of lateral inhibition

between two or more alternative unification links. In general, due to gradual decay of activation, more recent foot nodes will have a higher level of activation than ones which entered the unification space earlier. In addition, strength levels of the unification links can vary as a function of plausibility (semantic) effects. For instance, if instrumental modifiers under S-nodes have a slightly higher default activation than instrumental modifiers under an NP-node, lateral inhibition can result in overriding a recency effect.

The Unification Model accounts for sentence complexity effects known from behavioral measures, such as reading times. In general, sentences are harder to analyze syntactically when more potential unification links of similar strength enter into competition with each other. Sentences are easy when the number of U-links is small and of unequal strength. In addition, the model accounts for a number of other experimental findings in psycholinguistic research on sentence processing, including syntactic ambiguity (attachment preferences; frequency differences between attachment alternatives) and lexical ambiguity effects. Moreover, it accounts for breakdown patterns in agrammatic sentence analysis (for details, see Vosse and Kempen 2000).

The advantage of the Unification Model is that (a) it is computationally explicit, (b) it accounts for a large series of empirical findings in the parsing literature and in the neuropsychological literature on aphasia, and (c) it belongs to the class of lexicalist parsing models that have found increasing support in recent years (Bresnan 2001; Jackendoff 2002; Joshi and Schabes 1997; MacDonald et al. 1994).

Further support for a distinction between a memory component (i.e., the mental lexicon) and a unification component in syntactic processing comes from neuroimaging studies on syntactic processing. In a meta-analysis of 28 neuroimaging studies, Indefrey (2004) found two areas that were critical for syntactic processing, independent of the input modality (visual in reading, auditory in speech). These two supramodal areas for syntactic processing were the left posterior STG and the left prefrontal cortex. The left posterior temporal cortex is known to be involved in lexical processing (Indefrey and Cutler 2004). In connection with the Unification Model, this part of the brain might be important for the retrieval of the syntactic frames that are stored in the lexicon. The unification space, where individual frames are connected into a phrasal configuration for the whole utterance, might recruit the contribution of Broca's complex (LIFC). Empirical support for this distribution of labor between LIFC and temporal cortex was recently found in a study by Snijders et al. (2008). They did an fMRI study in which participants read sentences and word sequences containing word category (noun-verb) ambiguous words at critical positions. Regions contributing to the syntactic unification process should show enhanced activation for sentences compared with words, and only within sentences display a larger signal for ambiguous than unambiguous conditions. The posterior LIFG showed exactly this predicted pattern, confirming the hypothesis that LIFG contributes to syntactic unification. The left posterior

middle temporal gyrus was activated more for ambiguous than unambiguous conditions, as predicted for regions subserving the retrieval of lexical-syntactic information from memory.

A Unification Model Account of Syntactic ERP Effects

Since the discovery of the N400 effect in the beginning of the 1980s (Kutas and Hillyard 1980), a whole series of language-relevant ERPs have been observed. ERPs reflect the sum of simultaneous postsynaptic activity of a large population of mostly pyramidal neurons recorded at the scalp as small voltage fluctuations in the EEG, time locked to sensory, motor or cognitive processes. In a particular patch of cortex, excitatory input to the apical dendrites of pyramidal neurons will result in a net negativity in the region of the apical dendrite and a positivity in the area of the cell body. This creates a tiny dipole for each pyramidal neuron, which will summate with other dipoles provided that there is simultaneous input to the apical dendrites of many neurons, and a similar orientation of these cells. The cortical pyramidal neurons are all aligned perpendicular to the surface of the cortex, and thus share their orientation. The summation of the many individual dipoles in a patch of cortex is equivalent to a single dipole calculated by averaging the orientations of the individual dipoles (Luck 2005). This equivalent current dipole is the neuronal generator (or source) of the ERP recorded at the scalp. In many cases, a particular ERP component has more than one generator and contains the contribution of multiple sources. Mainly due to the high resistance of the skull, ERPs tend to spread, blurring the voltage distribution at the scalp. An ERP generated locally in one part of the brain will therefore not only be recorded at a nearby part, but also at quite distant parts of the scalp.

In connection to syntactic processing, two classes of syntax-related ERP effects have been consistently reported for over a period of more than ten years. One type of ERP effect related to syntactic processing is the P600 (Hagoort et al. 1993; Osterhout and Holcomb 1992). The P600 is reported in relation to syntactic violations, syntactic ambiguities, and syntactic complexity. This effect occurs in a latency range between roughly 500–800 ms following a lexical item that embodies a violation or a difference in complexity. However, the latency can vary, and earlier P600 effects have also been observed (Hagoort 2003; Mecklinger et al. 1995). Another syntax-related ERP is a left anterior negativity (referred to as LAN or, if earlier in latency than 300–500 ms, as ELAN; Friederici et al. 1996). In contrast to the P600, the (E)LAN has thus far (almost) exclusively been observed to syntactic violations. LAN is usually observed within a latency range of 300–500 ms. ELAN is earlier and has an onset between 100 and 150 ms. The topographic distribution of ELAN and LAN is very similar. The most parsimonious explanation is, therefore, that the same

neuronal generators are responsible for LAN and ELAN, but the temporal profile of their recruitment varies.

How does the Unification Model account for these effects? In the Unification Model, binding (unification) is prevented in two cases: (a) when the root node of a syntactic building block (e.g., NP) does not find another syntactic building block with an identical foot node (i.e. NP) to bind to; (b) when the agreement check finds a serious mismatch in the grammatical feature specifications of the root and foot nodes. The claim is that (E)LAN results from a failure to bind, as a consequence of a negative outcome of the agreement check or a failure to find a matching category node. For instance, the sentence *The woman sees the man because with the binoculars* does not result in a completed parse, since the syntactic frame associated with *because* does not find unoccupied (embedded) S-root nodes with which it can bind. As a result, unification fails. However, this does not necessarily mean that no interpretation of the grammatically ill-formed input will result. There is good evidence that semantic unification and syntactic unification both occur in parallel and, to some degree, independently. Moreover, ERP recordings in aphasic patients have shown that agrammatic aphasics can reduce the consequences of their syntactic deficit by exploiting a semantic route in online utterance interpretation (Hagoort et al. 2003). They thus provide evidence for the compensation of a syntactic deficit by a stronger reliance on another route in mapping sound onto meaning (multiple-route plasticity).

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

This account is not without problems, in terms of the available data. For instance, a word category violation often triggers both (E)LAN and P600. The current version of the model does not specify a mechanistic account of simultaneous unification failure and ongoing U-link attempts. Thus, empirical data

provide input for adaptations and improved versions of the model. However, this does not invalidate the conclusion that what we need is models beyond verbal description, with a sufficient level of computational explicitness to be able to characterize the immense body of empirical data on the electrophysiology of language.

The Immediacy Principle

The immediacy principle states that language comprehension beyond the single word level happens incrementally, in close temporal contiguity with information provided by the input signal. There is no principled priority for certain information types. This is not to deny that as a default there is a certain order in the flow of information. For instance, in hearing a word, its phonological information will be retrieved before its semantics; in speaking, the reverse relation holds. However, in the process of composing an interpretation from the lexical building blocks that make up a multi-word utterance, syntactic, semantic, and extra-linguistic information conspire and constrain the interpretation space in parallel. This view stands in contrast to a class of processing models that claim a priority for syntactic information (Frazier 1987a). The strong version of these serial syntax-first models of sentence processing assumes that the computation of an initial syntactic structure precedes semantic unification operations, because structural information is necessary as input for thematic role assignment. In other words, if no syntactic structure can be built up, semantic unification will not be possible. Recent electrophysiological evidence has been taken as evidence for this syntax-first principle (Friederici 2002). Alternative models (Marslen-Wilson and Tyler 1980; MacDonald et al. 1994) claim that semantic and syntactic information are processed in parallel and immediately used once becoming available.

Relevant empirical evidence for a syntactic head start in sentence comprehension is provided in a series of studies in which Friederici and colleagues found an ELAN to auditorily presented words, whose prefix is indicative of a word category violation. For instance, Hahne and Jescheniak (2001) and Friederici et al. (1993) had their subjects listen to such sentences as:

- (3) (a) Die Birne wurde im *gepflückt*.
 "The pear was being in-the *plucked*."
 (b) Der Freund wurde im *besucht*.
 "The friend was being in-the *visited*."

where the prefixes *ge-* and *be-* in combination with the preceding auxiliary *wurde* indicate a past participle and where the preposition *im* requires a noun. In this case a very early (between 100 and 300 ms) LAN is observed that precedes the N400 effect.

Although this evidence is compatible with a syntax-first model, it is not necessarily incompatible with a parallel, interactive model of sentence processing. As long as word category information can be derived earlier from the acoustic input than semantic information, as was the case in the above-mentioned studies, the immediacy principle predicts that it will be used as it comes in. The syntax-first model, however, predicts that even in cases where word category information comes in later than semantic information, this syntactic information will nevertheless be used earlier than semantic information in sentence processing. Van den Brink and Hagoort (2004) designed a strong test of the syntax-first model, in which semantic information precedes word category information. In many languages, information about the word category is often encapsulated in the suffix rather than the prefix of a word. In contrast to parallel models, a syntax-first model would, in such a case, predict that semantic processing (more in particular, semantic binding) is postponed until after the information about the word category has become available.

Van den Brink and Hagoort (2004) compared correct Dutch sentences (4a) with their anomalous counterparts (4b) in which the critical word (italicized in 4a/b) was both a semantic violation in the context and had the incorrect word category. However, in this case word category information was encoded in the suffix “-de.”

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

Figure 13.4 shows the waveform of the spoken verb form *kliederde* (messed). This verb form has a duration of approximately 450 ms. The stem already contains part of the semantic information. However, the onset of the suffix *-de* is at about 300 ms into the word. Only at this point will it be clear that the word category is a verb and not a noun as required by the context. We define this moment of deviation from the correct word category as the *category violation point* (CVP), because only at this time is information provided such that it can be recognized as a verb, which is the incorrect word category in the syntactic



Figure 13.4 A waveform of an acoustic token of the Dutch verb form *kliederde* (messed). The suffix *-de* indicates past tense. The total duration of the acoustic token is approximately 450 ms. The onset of the suffix *-de* is at approximately 300 ms. Only after 300 ms of signal, can the acoustic token be classified as a verb. Thus, for a context that does not allow a verb in that position, the category violation point is at 300 ms into the verb (see text).

context. Although in this case semantic information can be extracted from the spoken signal before word category information, the syntax-first model predicts that this semantic information cannot be used for semantic unification until after the assignment of word category.

Figure 13.5 shows the averaged waveforms that are time-locked to the CVP for two frontal sites where usually the ELAN is observed, and two posterior sites that are representative of N400 effects. As can be seen, the N400 effect clearly precedes the ELAN in time. Whereas the ELAN started at approximately 100 ms after the CVP, the N400 effect was already significant before the CVP. This finding provides clear evidence that semantic binding/unification (as reflected in the N400) can start before word category information is provided. This is strong support for the immediacy principle: information available in the signal is immediately used for further processing. In contrast to what a strong version of the syntax-first model predicts, semantic binding/unification does not need to wait until an initial structure is built on the basis of word category information.

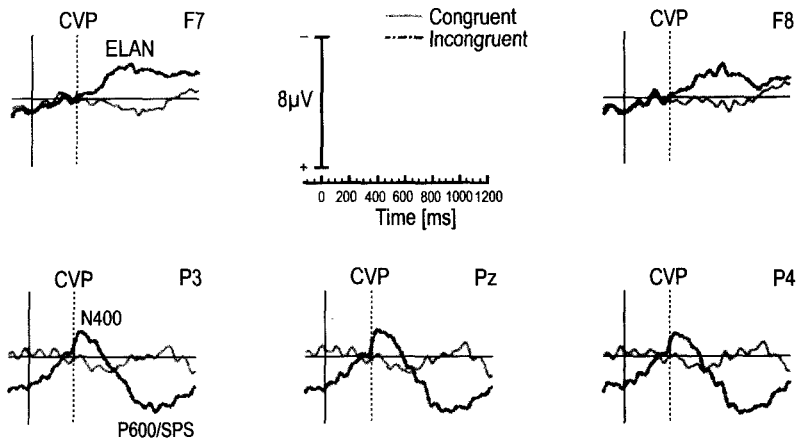


Figure 13.5 Connected speech. Grand average ERPs from two frontal electrode sites (F7, F8) and three posterior electrode sites (Pz, P3, P4) to critical words that were semantically and syntactically congruent with the sentence context (congruent: solid line), or semantically and syntactically incongruent (incongruent: alternating dashed/dotted line). Grand average waveforms were computed after time locking on a trial-by-trial basis to the moment of word category violation (CVP: category violation point). The baseline was determined by averaging in the 180–330 ms interval, corresponding to a 150 ms interval preceding the CVP in the incongruent condition. The time axis is in milliseconds, negativity is up. The ELAN is visible over the two frontal sites: the N400 and the P600 over the three posterior sites. The onset of the ELAN is at 100 ms after the CVP; the onset of the N400 effect precedes the CVP by approximately 10 ms (Van den Brink and Hagoort 2004).

Predictive Processing

One of the most remarkable characteristics of speaking and listening is the speed at which it occurs. Speakers easily produce 3–4 words per second; information that has to be decoded by the listener within roughly the same time frame. Considering that the acoustic duration of many words is on the order of a few hundred milliseconds, the immediacy of the ERP effects (discussed above) is remarkable. The ELAN has an onset on the order 100–150 ms, the onset of the N400 and the LAN is approximately at 250 ms, and the P600 usually starts at about 500 ms. Thus, the majority of these effects happen well before the end of a spoken word. Classifying visual input (e.g., a picture) as coming from an animate or inanimate entity takes the brain approximately 150 ms (Thorpe et al. 1996). Roughly the same amount of time is needed to classify orthographic input as a letter (Grainger et al. 2008). If we take this as our reference time, the earliness of an ELAN to a spoken word is remarkable, to say the least. In physiological terms, it might be just too fast for long-range recurrent feedback to have its effect on parts of primary and secondary auditory cortex involved in first-pass acoustic and phonological analysis. Recent modeling work on the *mismatch negativity* suggests that early ERP effects are best explained by a model with forward connections only. Backward connections become essential only after 220 ms (Garrido et al. 2007). The effects of backward connections are, therefore, not manifest in the latency range of at least the ELAN, since not enough time has passed for return activity from higher levels. In addition, LAN and N400 are following the word recognition points closely in time in the case of speech. This suggests that what transpires in online language comprehension is presumably based, to a substantial degree, on predictive processing. Under most circumstances, there is just not enough time for top-down feedback to exert control over a preceding bottom-up analysis. Very likely, lexical, semantic, and syntactic cues conspire to predict characteristics of the next upcoming word, including its syntactic and semantic makeup. A mismatch between contextual prediction and the output of bottom-up analysis results in an immediate brain response recruiting additional processing resources for the sake of salvaging the online interpretation process. As presented above, the Unification Model for parsing has prediction built into its architecture. Syntactic frames activated on the basis of an activated word form specify the local syntactic environment options and carry as such a prediction about the syntactic status of the next upcoming word; see Example (1). Recent ERP studies have provided evidence that context can indeed result in predictions about a next word's syntactic features (e.g., gender; Van Berkum et al. 2005) and word form (DeLong et al. 2005). Lau et al. (2006) have shown that the (E)LAN elicited by a word category violation was modulated by the strength of the expectation for a particular word category in the relevant syntactic slot. The role of structural predictions has, thus, found support on the basis of recent empirical findings. One possible source for these predictions is the language production system.

It might well be that the interconnectedness of the cognitive and neural architectures for language comprehension and production (Hagoort et al. 1999b) enables the production system to generate internal predictions while in the business of comprehending linguistic input. This “prediction-is-production” account has, however, not yet been tested empirically.

Afterthoughts on the Neurobiology of Syntax

Next to speed of language processing, another decisive characteristic of human language is its diversity. Decades of cross-linguistic work by typologists and descriptive linguists have shown “just how few and unprofound the universal characteristics of language are, once we honestly confront the diversity offered to us by the world’s 6–8000 languages” (Evans and Levinson 2009). Languages vary substantially in sound, lexical meaning, and syntactic organization. Nevertheless, children acquire language faster, almost universally, and seemingly more automatically than, for instance, musical skills. This indicates that evolution has provided humans with a brain that is characterized by a certain language-readiness. This might have happened by optimizing the neural infrastructure for the ensemble of cognitive systems (e.g., systems for memory, unification, executive control) that collectively provided language-readiness to the brain, instead of forming a universal neural machinery dedicated to syntax. According to the former view, there might also be overlap and common recruitment of certain brain regions in the service of different cognitive functions with similar requirements. For instance, domain-specific memories for syntax, music, and action schemata might all recruit the domain-general unification capacities of Broca’s complex to produce or decode the combinatorial aspects of language, music, and action (Patel 2003). In support of this account, a recent study with Broca’s aphasics characterized by a syntactic deficit showed that these same patients also were impaired in processing musical syntactic (harmonic) relations in chord sequences (Patel et al. 2008). How we continue to interrogate the brain will, to a large extent, depend on the perspective taken on this issue.

Acknowledgment

I am grateful to Karl Magnus Petersson for discussions about and comments on a previous version of this contribution.

Overleaf (left to right, top to bottom):
Tatjana Nazir, group discussion, Balázs Gulyás
Csaba Pléh, Anna Fedor, Peter Hagoort
Angela Friederici, Group discussion
David Caplan, Jens Brauer, Wolf Singer