Psychophysiology, 34 (1997), 383-398. Cambridge University Press. Printed in the USA.

PRESIDENTIAL ADDRESS, 1996

Views on how the electrical activity that the brain generates reflects the functions of different language structures

MARTA KUTAS

Departments of Cognitive Science and Neurosciences, University of California, San Diego, USA

Abstract

Human language is what it is because of its function and its implementation. We are far from understanding how language comprehension is carried out by the human brain. This task can be made easier by considering that evidence for the what and how of language comes from the study of linguistics, psychology, and neuroscience. The approach outlined herein describes how these different sources of evidence can be combined in studies of written and spoken sentence processing by using a measure of the brain's electrical activity. The outcome is a more temporally precise view of the analysis of language structures in our minds and brains.

Descriptors: Language, ERPs, Working memory, Word frequency, Sentence processing, Parsing, Relative clauses

We use language daily to leave each other messages via e-mail, voice mail, or, heaven forbid, in person. We use it to convey facts, rumors, thoughts, and wishes; we use it to discuss, to sway, to imagine, and sometimes to educate. If we do not suffer from some sort of brain damage, we use language quite well from the age of 2 years. How is this capability achieved? What computational, psychological, and physiological processes support language comprehension and production? Why have bioengineers not yet built a robot that can understand what anyone says and means and be able to get a word in edgewise-something comprehensible if not insightful? In my opinion, the difficulty of this enterprise calls for a cognitive neuroscience approach that spans theories and data from multiple disciplines including linguistics, psychology (including but not limited to psycholinguistics), and neuroscience. Thus, I believe that a true understanding of language processing requires knowledge of and an appreciation for (a) the structure of language, (b) the structure of different mental processes such as working memory (WM), (c) the organization of the brain, and (d) the strengths and limitations of the techniques and measures from which inferences are drawn. This article outlines such an approach, with an emphasis on how electrophysiological investigations have contributed to our understanding of language processing.

Over the past 5 years or so, event-related brain potentials (ERPs) have been used as a source of psycholinguistic evidence (although not always correctly). At about the same time, experimental designs in electrophysiological investigations of language have become increasingly elegant. As a consequence, electrophysiological data are now considered to be informative not only by those who use ERPs but also by mainstream psycholinguists, neurolinguists, and even theoretical linguists. In part, this development is because current-day linguistic theories are so subtle that primary linguistic data are often no longer sufficient to differentiate among them. By primary linguistic data I mean grammaticality judgments, that is, native speaker's judgments of linguistic well-formedness; for example, whether or not it is appropriate for a speaker to say "Charles doesn't shave because *him* tends to cut himself." Neither a trained linguist nor a woman on the street would consider this sentence grammatically correct for standard English. Coulson, King, and Kutas (in press) showed that the reader's brain reacts to such grammatical violations within a few hundred milliseconds of their occurrence; there is a late positivity between 500 and 800 ms in the ERP elicited by *him* relative to that elicited by *he* (see Figure 1). Osterhout and Holcomb (1992) likewise observed a late positivity in the ERP elicited by to in sentences such as "The broker persuaded *to* sell the stock. . .'and that elicited by was in "The broker hoped to sell the stock *was*. . .... Similar late positivities, among other effects, have been reported by several laboratories in re-

---

This study was supported in part by grants from NIMH (MH52893), NIA (AG08313), and NICHD (HD22614) to M. Kutas and by various predoctoral and postdoctoral fellowships to members of her laboratory from the McDonnell-Pew Center for Cognitive Neuroscience at San Diego, Center for Human Information Processing, and the Center for Research in Language.

Special thanks go to the members of my laboratory over the years, without whom much of this work might never have been carried out, in particular Seana Coulson, Jonathan King, Robert Kluender, Cyma Van Petten, and Jill Weckerly, who have taught me the complexities of language, to Paul Krewski and Ron Ohst who kept the various machines running, and to Steve Harris, Heather Mclsaac, and Mark Mitchiner, who provided technical assistance as needed. Thanks to Thomas Muente, Phil Holcomb, and Greg Miller for comments on a previous draft of this manuscript.

Address reprint requests to: Marta Kutas, Department of Cognitive Science, D-015, University of California, San Diego, La Jolla, CA 92093-0515, USA. E-mail: mkutas@ucsd.edu.

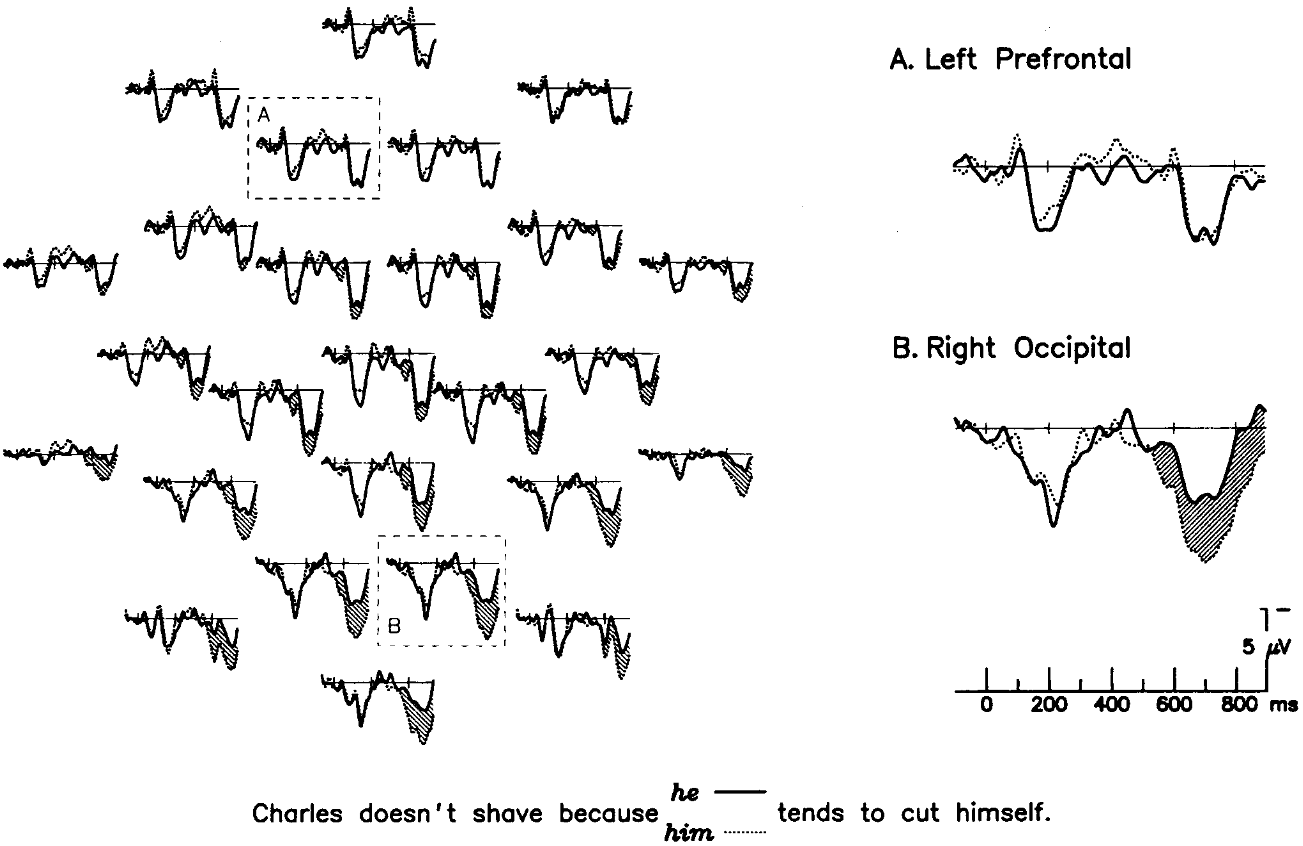

Figure 1. Comparison of the grand average ERPs to pronouns that are grammatically correct (solid line) versus incorrect (dotted line) in their grammatical case marking. The left side of the figure shows the data from all recording sites, with top being at the front of the head and the bottom at the back of the head. The right side of the figure shows data magnified from the two locations marked in the left side, over a left prefrontal (A) and a right occipital (B) site. Data are from Coulson, King, and Kutas (in press).

sponse to a host of grammatical violations including subject-verb agreement, pronoun case markers, subcategorization constraints, phrase structure, wh-movement (e.g., reviewed in Osterhout & Holcomb, 1995; Coulson et al., in press). Can this late positivity be considered a as a sign of the ungrammaticality? Yes, in the sense that the P600, or syntactic positive shift (SPS), reflects the fact that the listener or reader noticed that there was an unusual (in this case, ungrammatical) event. Certainly, grammatical errors of this type are not what one is wont to read under normal circumstances, but is it possible to infer the converse? That is, are we licensed to take the presence of a late positivity to mean that the person ran across an ungrammatical event? No, we are not! Moreover, as noted by Kluender (1991), the conceptual link between the P600 and ungrammaticality is even more tenuous when more complicated constructions such as wh-questions are investigated. For instance, he observed no late positivity to this in sentences such as "What did he wonder who he could coerce into this ...... although it is clearly ungrammatical (Kluender & Kutas, 1993b).

More than one linguistic theory accounts for the ungrammaticality of these and other such violations, which is the reason why some other source of evidence like ERPs is needed to help adjudicate among them. Using ERPs to this end requires delineating the antecedent conditions necessary for eliciting whatever ERP component is put into service for this purpose. This delineation involves determining which factor is responsible for eliciting a particular component and which factors modulate its amplitude or latency and identifying its functional significance. In the domain of language processing, there are a number of components for which the antecedent conditions have been sought or examined, including

1. the N400 to semantic violations as in "Bobcats hunt mice, squirrels, rabbits, laughs, and many other small animals" (e.g., Kutas & Hillyard, 1980, 1983, 1984);
2. the P600 or SPS to grammatical violations such as those in "Turtles will spit out things they *does* not like to eat" (e.g., Hagoort, Brown, & Groothusen, 1993; Kutas & Hillyard, 1983);
3. a left frontal negativity of 300-500 ms in response to syntactic violations such as those in "What did the scientist criticize Max's proof of" (e.g., Muente, Matzke, & Johannes, in press; Neville, Nicol, Barss, Forster, & Garrett, 1991);
4. the left anterior negativity (LAN effect) in a comparison of a wh-question ("Who has she ... ?") versus a simple yes-no question ("Has she ... ?") (Kluender & Kutas, 1993b), which some researchers equate to the frontal negativity just mentioned;
5. the N280 to closed-class items (Neville, Mills, & Lawson, 1992);
6. the N400-N700 to closed-class or function words within sentences (e.g., Neville et al., 1992; Van Petten & Kutas, 1991);

7. the clause-ending negativity in written and spoken sentences (e.g., Kutas & King, 1996; Mueller, King, & Kutas, 1997);
8. the ultraslow frontal positivity that develops across simple transitive clauses such as "The doctor answered the call" (e.g., Kutas & King, 1996).

We will return to several of these in the course of this article.

At this point, it is sufficient to note that it has not been easy to specify the antecedent conditions for many of these components or effects. To determine what a component is sensitive to, one must first identify it and then home in on exactly what sensory processing, cognitive operations, or motor event led to its elicitation. Both requirements raise problems that have haunted cognitive electrophysiologists for years, given the difficulty of unequivocally defining a component (Donchin, Ritter, & McCallum, 1978). For as soon as the experimental conditions (or the task demands) of an experiment are changed "too" much, it is difficult, if not impossible, to determine whether a component in one condition is exactly the same as that in another condition. One cannot necessarily use its latency as a criterion for identity because systematic variation in latency is what makes the N200 (whose latency changes with the difficulty of discrimination), the P300 (whose latency changes with the difficulty of categorization), and the lateralized readiness potential (whose latency changes with motor preparation) all useful measures in mental chronometry (for review, see Coles & Rugg, 1995). Similarly, one cannot resort to differences in scalp distributions in any straightforward way because that too is allowed to change even as a component maintains its name and functional significance. I refer to the presumed equivalence of the visual and auditory N200s, Contingent Negative Variations, and the P300s (see Coles & Rugg, 1995).

As psychologists, cognitive electrophysiologists most often have opted for functional identity as a criterion, but its application is often not so obvious without further support from the component itself. In fact, we often resort to similarity in the way a component or an effect looks (e.g., waveshape, duration, scalp distribution) when faced with uncertainties about its functional identity. For instance, most researchers have equated the N400 effect (i.e., the difference between ERPs to semantically incongruous and congruous words) in sentences to the relatedness or priming effect in the ERP to the second of two words in a related versus unrelated word pair. This equation is not simply because both ERP effects display a sensitivity to semantic processing and context, but because these effects have about the same peak latency (around 400 ms), a similar morphology (monophasic negativity), and a similar albeit not-identical distribution across the scalp (posterior and slightly larger over the right than left hemisphere; for review, see Kutas & Van Petten, 1994; Osterhout & Holcomb, 1995). Researchers have not been especially disconcerted by the finding that N400s elicited by anomalous sentence endings presented at faster rates (10 words/s) are generally later and appear somewhat more frontal in their distribution than those presented at slower rates, presumably because we can come up with an account of how the processing for meaning may be held up at fast presentation rates (Kutas, 1987). Likewise, we have not been discomfited by the longer latency N400s in elderly individuals or patients with Alzheimer's disease (relative to younger adults) because this change fits with the general slowing typically observed with advancing age (e.g., Gunter, Jackson, & Mulder, 1996; Schwartz, Kutas, Salmon, Paulsen, & Butters, 1996). On the whole, we also have not been concerned with the fact that the auditory N400 begins earlier and appears to last longer than the visual one because we can attribute the differences to coarticulation, among other factors (e.g., Holcomb & Neville, 1991; McCallum, Farmer, & Pocock, 1984).

Are these negativities in the ERPs to semantic anomalies in written and spoken English (or German, French, Dutch, etc.) and in American Sign Language the same N400? In such comparisons, cognitive electrophysiologists tend to overlook the basic waveform differences due to individual variability, modality, stimulus characteristics, and so on by looking at effects. That is, we look at difference waves (ERPs to incongruent minus congruent ERPs), and this does tend to reduce the contribution of irrelevant ERP differences, but we are still faced with nonidentity between different ERPs that make it difficult to know the extent of the relation between two effects. For instance, how are we to interpret the finding by Ganis, Kutas, and Sereno (1996) that the congruity effect for words and for line drawings at the ends of sentences are at once similar and different? Both incongruous written words and pictures yield a monophasic negativity and thus are similar, but the picture N400 has a decidedly frontal maximum, and the word N400 has a decidedly posterior maximum, and thus the two are different. What is the correct answer then: Are they the same or are they different? My inclination is to think that we are seeing the same functions performed across different inputs, that is, the same computations but carried out by the different brain regions that receive these different inputs. Nevertheless, how does one become more objective about this sort of comparison? Even at a strictly functional level, it is not always easy to be certain that all that has changed from one condition to another is the factor that one believes it to be or the factor that was manipulated.

A prime example of this difficulty follows. One supposedly language-related component is the N280. What are the necessary and sufficient conditions for eliciting an N280 component? Neville et al. (1992) proposed that the N280 is a special marker of closed-class items (e.g., articles, conjunctions, prepositions, auxiliaries). The closed class refers to the small words in the language that serve primarily a grammatical function of relating open-class (or content) words to each other. Content words (e.g., adjectives, adverbs, nouns, verbs) are those that convey meanings. Neville et al. reported that the ERP to closed-class (or function) words is characterized by a negative peak that is most clearly seen at frontal sites over the left hemisphere (i.e., over Broca's area), whereas the ERP to open-class words had no such N280. By contrast, the ERP to open-class words is characterized by a large negative wave (N400) over posterior sites (i.e., Wernicke's homolog) over the right hemisphere. This double dissociation was offered as evidence for a fundamental difference in how open- and closed-class words are processed, accessed, and stored. Thus, it seems that all that is needed to elicit an N280 is a closed-class word and all that is needed to elicit an N400 is an open-class word. However, by focusing on antecedent conditions, we find that function words can elicit N400s and content words can elicit the N280s, albeit at a longer latency.

First, let us examine two examples of N400s in response to closed-class words. Kluender (1991) observed N400s to the closed-class words *that,* if, and *who* in yes-no questions such "Can't you remember that, if, who . . . ," with the largest amplitude to the interrogative pronoun *who,* which he explains in terms of referential specificity (see Kluender & Kutas, 1993b). Likewise, King and Kutas (1995a) observed N400s to the definite article *the* in the relative clause of sentences such as "The professor that the student regularly drove crazy committed ...... although only in poor comprehenders. Such findings clearly demonstrate that N400 amplitude changes with how expected a word is within a given semantic context (and thus how easy it is to integrate), regardless of the class of the eliciting word. Both open- and closed-class words can be highly predictable or unexpected and thus easy or difficult, respectively, to accommodate in the ongoing discourse; it is this semantic expectancy and its consequences for online processing and not lexical class that are reflected in fluctuations in N400 amplitude. Thus, the N400 cannot be a marker for open-class words. In fact, by the end of a sentence, as semantic context has built up, even the response to open-class words appears to have fairly small N400s (Van Petten & Kutas, 1990); a similar reduction in N400 is seen with the third repetition of a semantically anomalous sentence-terminal word (Besson, Kutas, & Van Petten, 1992). Closed-class words are far more predictable from sentential contexts than are open-class words, and we think that it is primarily this difference in predictability that accounts for the larger N400 generally elicited by open-class words.

Similarly, we find that the data do not support the proposal that the N280 is a unique response to closed-class words. Of relevance is the work of Jonathan King who sorted the ERPs to words in hundreds of sentences into 10 grammatical categories, 7 closed class and 3 open class, and then examined the ERP at the site where the N280 was maximal, namely a left anterior temporal site close to Broca's area. A comparison of the ERPs at this site to definite articles, adverbial prepositions, adjectives, verbs, and nouns, for example, reveals that all of these word types contain a negativity but at different peak latencies (Figure 2, left-hand column). Thus, what factor is the brain sensitive to in this case?

King hypothesized that if this potential were related to a word's processing, then it would change with one or more of the lexical factors that have been shown by reading time, reaction time, and eye movement gaze duration studies to be critical for determining how quickly a word is read. Two such lexical factors are the length of a word and its frequency of occurrence in daily usage. Typically, shorter words are processed more quickly than are longer words and common words more quickly than rare words (Just & Carpenter, 1980). King tested his hypothesis by regressing the peak latency of the negativity for each of 10 categories onto the sum of its mean length and scarcity (a transformation of frequency). The resulting regression accounted for 86% of the variance (in mean peak latency). Thus, the negativity at left frontal sites apparently is sensitive to some combination of the length and frequency or a

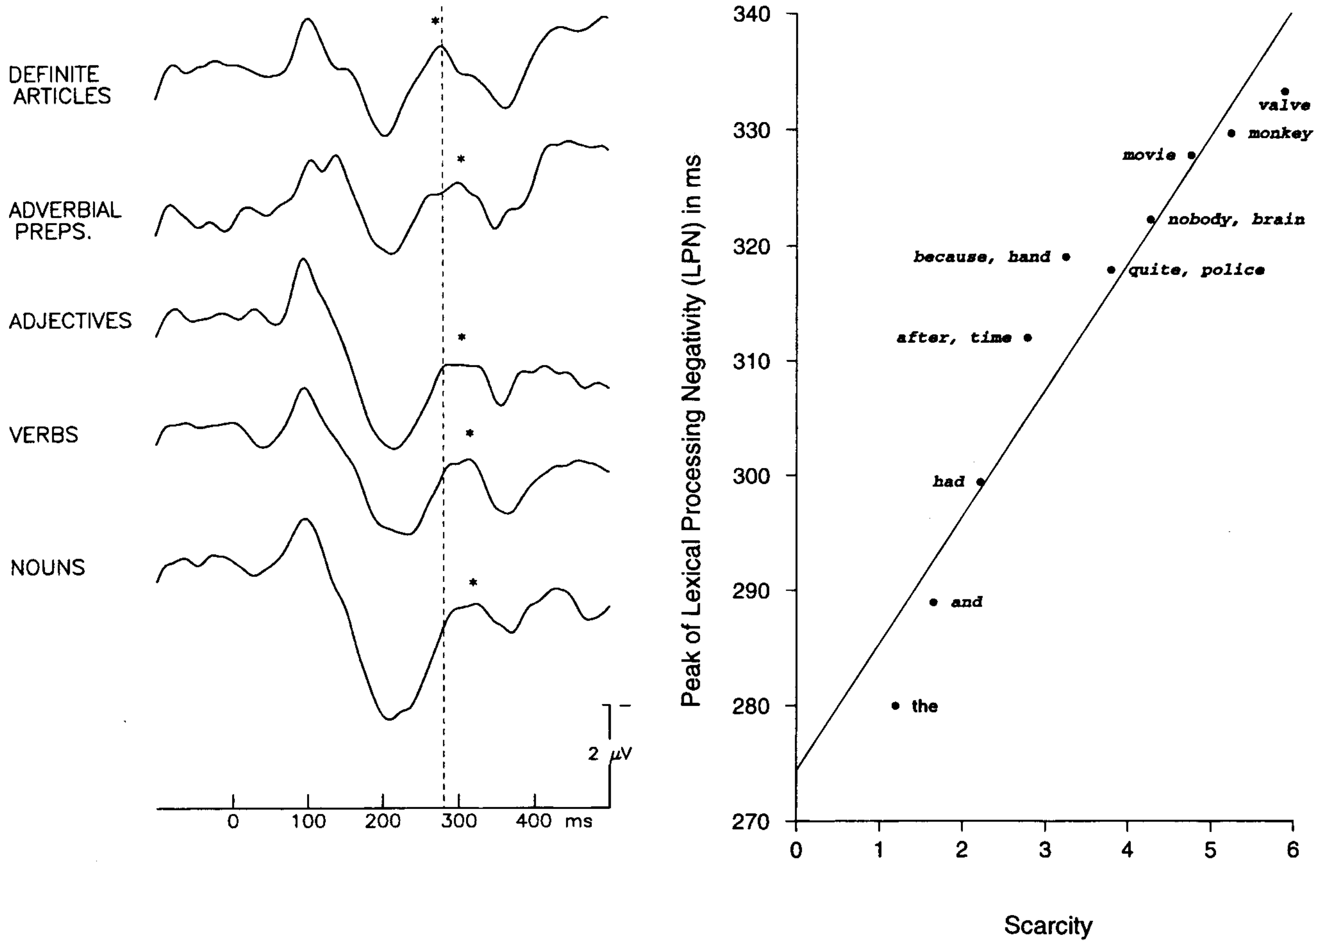

Figure 2. On the left-hand side are grand-average ERPs from a left frontal recording site for five different lexical classes. On the right-hand side is grand mean lexical processing negativity (LPN) plotted against scarcity (transformed log word frequency, i.e., 6-log *F*, where *F* is the frequency of occurrence in a million-word corpus). Each data point is the grand mean over subjects of peak latency measurements made on individual waveforms for that frequency bin. The minimum number of items per bin per subject = 198. Points are labeled with representative lexical items from each frequency class; classes containing open- and closed-class words are double labeled.

word; therefore, we (re)named this negativity the lexical processing negativity (LPN) (King & Kutas, 1995b). However, in this data set, most of the long high-frequency words are open class, whereas all of the short low-frequency words are closed class with no words in between. Accordingly, one might contend that this really is mostly a lexical class difference, although it is unlikely given that we have observed variability in the latency of the so-called N280 even for different closed-class items. Nonetheless, to be on the safe side, we analyzed more data including more frequency classes from more electrodes from a different group of individuals. Also, in this case, we regressed the peak latency of the negativity onto word scarcity (cutting across word class). Whereas some frequency bins contain only open- or only closed-class items, four frequency classes contain words from both lexical classes (Figure 2, right-hand column). The average regression accounted for 90% of the variance in peak latency of the negativity at one left frontal site; none of the analyses within the same latency range at any other recording site (except the few nearby) or within any other latency window showed such a reliable relation with word frequency. Clearly, this negativity cannot be taken as proof that different brain systems processed the two classes because the words from both classes showed the same response but at different latencies; the LPN is simply shorter for the closed class because these words occur more frequently.

Further evidence in support of our contention that the LPN is the same component in the ERP to open- and closed-class words comes from the similarity of their scalp distributions. The original argument for the involvement of different brain systems for these two lexical classes was based on their different spatial distributions across the scalp; only ERPs to closed-class items were supposed to have an N280. However, if we take their different latencies into account, that is, 280 ms for the closed class and 330 ms for open class on average, then we find that their scalp distributions are remarkably similar. This similarity can be seen in the current source density (CSD) maps for the two latencies (CSD maps represent the second spatial derivative of the potential fields). In sum, the existence of the LPN for both word classes indicates that at some point the two are processed similarly, albeit at different speeds that are proportional to their different frequencies of occurrence. Naturally, such a difference in processing speed could have real functional effects on open- and closed-class words. Observed processing differences between the two classes thus can be attributed to this difference in processing speed rather than to any specific closed-class processor of the brain such as Broca's area.

The CSD maps of this sort reflect a relatively big change in ERP recording, namely the use of electrode caps or nets, which combined with the somewhat reduced cost of amplifiers has made it easier to record more channels. The availability of more scalp locations has had the positive consequence of providing a more precise description of the spatial distribution of a potential at any given moment (although only rarely are we looking at the activity of only one generator) and better localization (see Yvert, Bertrand, Echallier, & Pernier, 1996), but this precision comes at a significant cost, inasmuch as the techniques for analyzing, comparing, and contrasting distributions have not kept up with the increase in the number of recording sites. Spatial principal components analysis (PCA) (much like the temporal PCA of the past except that time points are replaced by electrodes) like other spatial filtering algorithms such as the generalized Eigen systems analysis (A. M. Dale, personal communication) and the independent components analysis (Bell & Sejnowski, 1995) may help reduce the dimensionality of ERP data from across the scalp in a comprehensible manner.

The difficulty of equating language-related ERP components across different experiments has been exacerbated further by the increased complexity of the experimental designs that are being utilized. These days, language processing refers to more than simply whether the stimulus materials are verbal or nonverbal. Moreover, more sophisticated questions are being asked than what is the difference between a word and a nonword or is there a brain or ERP correlate of the known left hemisphere advantage for language processing. Rather, many researchers are employing different ERP effects to address questions that are in the mainstream of psycholinguistics. Personally, I think some of the more controversial, supposedly fundamental, questions about the "innateness" or "modularity" of language are misguided because these questions can never be answered unequivocally. The brain is plastic, and no infant pops out of the womb understanding or speaking Swahili. Even so, under normal circumstances, all infants have the potential to speak and understand any language; whether or not they eventually do so depends on to which language, if any, they are exposed. The potential for learning is unquestionably innate, and for the moment there is no reason to believe that language differs in any substantive way from other higher cognitive skills in this regard. However, is there some part of the brain (e.g., a language acquisition device) that is present at birth that makes it most suited of all the areas in the brain for learning to comprehend and produce speech because of its inherent organization? In other words, is there an area specialized for language processing? Perhaps, but in my opinion, the relevant constraints on this area(s) are not in its innate internal organization but rather on where it is situated in relation to the inputs to it and the outputs from it, given the function(s) it must perform. In any case, I submit that if there is such a region, then it is quite large and is distributed both within and across the hemispheres and certainly not in exactly the same place in everyone.

Over the past 20 years or so, there have been some significant insights in our understanding of the involvement of the brain in language, and these are of relevance for electrophysiological investigations. Until quite recently, the predominant model of how language is represented in the brain was based on data from patients with brain damage. According to this model, in the great majority of right-handed people, language is subserved primarily by two interconnected areas in the left hemisphere: Broca's area in frontal cortex and Wernicke's area in the temporal lobe near its junction with the parietal lobe (e.g., Geschwind, 1965). We have yet to agree on the functional significance of these brain areas, despite more than 100 years of accumulated neuropsychological data. Among the different proposals are Broca's area for production and Wernicke's area for comprehension, Broca's area for syntax and Wernicke's area for semantics, Broca's area for grammar and Wernicke's area for the lexicon, Broca's area for regularity and Wernicke's area as a memory store for exceptions, and so on. Naturally, there is some evidence supporting each of these dichotomies, making it difficult to reach a consensus on exactly what Broca's area does. In fact, determining the functional significance of Broca's area is no less of a problem than defining the functional significance of an ERP component such as the N280 or the N400 for essentially the same reasons: (a) it is difficult to identify Broca's area, that is, to define its boundaries, and (b) it has proven difficult to link Broca's area to a particular function from patients' behavior when the area is compromised.

The structural images of the brain provided by computed axial tomography and magnetic resonance imaging (MRI) scans have made it easier to delimit exactly which areas are actually damaged by stroke. At the same time, however, these images also have made it more difficult to accept the view that damage to Broca's (and no other) area, for example, always leads to a definable set of symptoms (Broca's aphasia) in all people. The correlation between damage to Broca's area and aphasic symptoms is quite low (e.g., Willmes & Poeck, 1993).[1] In part, this low correlation may reflect the fact that functional damage as revealed by positron emission tomography (PET) scans in these patients often extends beyond the boundaries of the observable structural damage. However, it also seems likely that what has traditionally been called Broca's area is not a single area. Extrapolating from what is known about the numbers and sizes of visual, auditory, somatosensory, and motor areas in macaque monkeys, the average brain area in the human covers approximately 1,000 mm^2 of cortical surface. Thus, what is typically considered Broca's area probably includes four to five distinct brain areas (also see Deacon, 1992). This account also explains the observed variability in the symptoms associated with damage to Broca's area. An analogous argument can be made for Wernicke's area.

In summary, the number of anatomically and functionally distinct brain regions within the classical language areas seems to be larger than previously supposed. This notion may explain the large variability in the loci of activity observed across many PET and functional MRI (fMRI) studies, all of which presumably engaged "language" processing areas (e.g., Roland, 1993). Over the past decade, a third language area on the bottom of the temporal lobe, called the basal temporal language area, has been discovered (Luders et al., 1986).

What might the functions of these different "language areas" widely (although not randomly) distributed in the brain be? How likely is it that all these areas are used only in the service of language processing? We do not yet know. However, given the way the senses and the motor system are wired, it would follow that some areas would become specialized by experience to best deal with certain types of inputs and to emit certain classes of outputs. In this sense, with experience the brain becomes both highly specialized and modular. That said, I still find it more informative and potentially rewarding to ask not whether language is innate or modular but rather, What categories of inputs is the brain sensitive to? What are the natures of the computations that are performed by the neocortex, among other regions, and how are they put to use in the service of language? What external and internal factors affect processing, and what are the time courses of their influences? What can humans ignore and to what must they attend? What sorts of regularities in time and space does the human brain accommodate and assimilate? What aspects of our reality come to consciousness, and what aspects simply cannot? Although completely satisfying answers to any of these questions do not yet exist, there are ample data to show that ERPs are a good index of many of the factors that the brain is sensitive to, can ignore, must attend to, and so on. Moreover, there is precious little evidence that any of these factors is specific to language.

Thus far, I have provided a brief overview of the functional significance of several language-related ERP effects and have highlighted the increased ability of electrophysiologists to collect from more recording sites, the existence of different neuroimaging techniques to complement ERPs and constrain their interpretation, and a new emerging view of how language is organized in the brain. In addition, it is important to mention a new view of the ERP. On this contemporary view, the ERP is more than a multidimensional behavioral measure and more than a continuous reaction time. Twenty years ago, when the analytic techniques for localizing scalp ERP generators were extremely limited, it was argued that the ERP would be useful as a tool for tracking cognitive activity even if it were generated in the big toe. This is a defensible position and the argument still holds today. In point of fact, however, ERPs are not generated in the big toe but rather in the brain. The currency of the brain is electrochemical, and the electroencephalogram and magnetoencephalogram are the only neuroimaging measures that directly reflect neural activity in real time in an intact human. Both techniques provide a snapshot of a significant, albeit limited (to open-field activity), picture of the activity of the brain at any given moment.

Cognitive ERP researchers have tended to be shy in pushing ERPs as a neuroimaging technique because the inverse problem of determining the locations, orientations, and time courses of sources producing the electric recordings at the scalp is ill-posed. In other words, the inverse problem has no unique solution because there are infinitely many distributions of dipole moments inside the brain that are consistent with any set of electric scalp recordings (e.g., Nunez, 1981). Moreover, only open-field activity can be recorded at the scalp. The inverse problem does not disappear even with recordings taken from inside the brain, as in those suffering from medically intractable epilepsy. Nonetheless, such data have increased the stock of ERPs as a neuroimaging measure because they allow better localization of certain scalp ERP effects. For example, intracranial recordings led McCarthy and colleagues (McCarthy, Nobre, Bentin, & Spencer, 1995; Nobre, Allison, & McCarthy, 1994) to suggest that the N400 is generated in the parahippocampal anterior fusiform gyrus.

Precise localization notwithstanding, scalp-recorded ERPs provide a fairly accurate picture of the activity of neural ensembles in the neocortex because they meet all the necessary constraints for producing an externally observable electric field: (a) neurons are aligned in some systematic fashion, (b) sources and sinks have a nonradial symmetric distribution, and (c) neurons are activated in synchrony. Approximately 70% of the neurons in the neocortex are pyramidal cells that have apical dendrites extending from the soma toward the cortical surface; these dendrites give the cortex a columnar appearance. When activated, these dendrites create an approximate dipolar source-sink configuration oriented perpendicularly to the cortical sheet. Although any one neuron generates too small a signal to be seen at any considerable distance, hundreds of thousands of them can produce a potential field that is strong enough to be detected at the scalp. These fields are considered to be the primary source of scalp-recorded ERPs.

The neocortex may not be the entire brain, but it is the principal neural substrate of recognition, has a crucial role in motor execution and planning, and certainly underlies much of what is meant by "higher cognition" such as thinking, problem solving, and language. Furthermore, much is known about its circuitry and the types of computations that are performed in neocortical columns (e.g., White, 1989). All the neocortical areas are organized according to a common set of principles including a macrostructural

---

[1] This discussion does not question the data that damage in the perisylvian areas typically lead to some type of aphasia, whereas damage elsewhere in the brain almost never does (with the possible exception of the more transitory aphasia consequent to damage of the basal temporal language area). Rather, this discussion is intended to question the extent to which the brain honors the classical aphasic syndromes described by Wernicke and Lichtheim, which is still prevalent in current neuropsychological texts.

division into areas with inhibitory and excitatory lateral local connections and topographic long-distance connections, a mesostructural division into six layers each with its own inputs or outputs, and a microstructural division into a set of universal cell types. Because the brain electrical activity recorded at the scalp reflects primarily the activity of these neocortical areas at the scalp, its likely dipole generators can be localized via modeling techniques constrained by additional anatomical information provided by fMRI or PET (e.g., Dale & Sereno, 1993; personal communication, 1996) and the temporal information provided by the ERPs.

To understand language processing, it is as important to consider the structure and functions of language as it is to consider the structures and functions of the brain in language. Language comprehension includes a wide range of analytic and synthetic processes. Sentences are not just words lined up end to end, and yet much of the research on brain and language (especially with the newer neuroimaging techniques based on metabolic or hemodynamic activity) takes place at the level of letter strings. Language has structure such that different orders of the same lexical items (words) convey different meanings. It is the structure of language that makes the meaning of "in science the benefits of collaboration outweigh those of competition" different from that of "in science the benefits of competition outweigh those of collaboration" but similar to that of "in science the benefits of competition are outweighed by those of collaboration" and makes "are outweighed in by those science collaboration of benefits competition the" not a meaningful sentence at all. Structures have processing consequences. Some are easier than others. Structure is used to make sense of utterances quickly, and somehow this structure is appreciated by the brain, which encodes light and sound into linguistic categories and ultimately constructs meaning. For after all conveying meaning, in other words, passing along information that was not already known, is one of the main aims of many language exchanges. Writers and speakers are often in the business of using language to change a reader's or listener's state of knowledge, or construal of a given situation, and perhaps even their behavior.

Exactly how this conversion takes place at a psychological or a physiological level is both unknown and highly controversial, but one outcome is some sort of a (re)depiction of activities of different discourse entities (people, places, things) in the mind of the listener. Thus, discourse entities must be selected and their interactions encoded in a dynamic mental model or a frame that supports the kinds of inferences that are required to yield new information or new construals to the listener. For any listener to understand the sentence "Netanyahu whom Clinton halfheartedly supported soon planned to live up to his agreement with Arafat," that listener would have to keep some temporary representation of Netanyahu and Clinton until it became clear who did what to whom. Although sentences provide structure, for example, in the order that words come, there are still many moments of syntactic ambiguity when it is not clear who is doing and who is being done unto. Moreover, in addition to understanding the meaning of each word in the sentence, more information needs to be retrieved from long-term memory to understand the full meaning of an utterance (although I am of the opinion that even understanding a single word is a constructive process as opposed to a simple lookup). Thus, for example, knowing that Clinton is president of the United States and as such has taken on a mediator role in the Middle East and that Netanyahu is the recently elected hawkish prime minister of Israel who has not wanted to live up to his predecessor's agreement to return Hebron to the Palestinians led by Arafat frames this sentence very differently than simply knowing that all three are public figures whose activities are regularly featured on the evening news. To understand the bulk of language, one must be concerned with levels beyond words, which means dealing with some temporary representations and some more long-lived discourse representations. Temporary representations in turn require some form of WM.

Exactly what form such WM would take and how it is used remain areas of active research. It is generally agreed that WM is capacity limited. However, there is considerable disagreement as to whether WM is best viewed as discrete chunks, activations within a production system, activations across units of a neural network, or an articulatory loop, and so forth (e.g., Burgess & Hitch, 1992; Just, Carpenter, & Keller, 1996; Waters & Caplan, 1996). For our purposes, suffice it to say that WM is capacity limited and that its contents are temporary. Items are around only as long as they are needed for certain computations; thereafter, they are lost or actively suppressed. Moreover, it is equally important to realize that there is an order to the computations in WM and that their temporal sequencing can serve to enforce some degree of separation among different processes. Such separation is essential to avoid ambiguity and misunderstanding. In short, both the capacity limitations and the temporal characteristics of WM impact real-time language comprehension. How and when are questions that are more likely to be answered than is language innate or is there a language module. Moreover, in the long run, I expect this approach to language processing to reveal just how much of what are usually considered linguistic phenomena to be explained practically in terms of more general cognitive functions.

Our laboratory has studied sentence processing to search for ERP signs of hierarchical relations between words independent of their actual meanings. We began by investigating a simple transitive sentence such as "The secretary answered the phone," which consists of a subject (S), a verb (V), and an object (O). An examination of the ERPs to simple transitive clauses revealed that there are electrophysiological properties that emerge across sentences that are more than the responses to individual words lined up end to end, specifically, some very slow effects that cumulate and fluctuate across the course of the sentence (Kutas & King, 1996).

A revealing way of looking at these types of slow potential effects is to apply a low-pass digital filter to the cross-sentence ERPs, leaving only the slow activity. The consequence of applying such a low-pass filter to an oversentence waveform is shown in Figure 3; the tracing represented in the dotted line is the potential after filtering. Although this representation clearly simplifies the waveforms, a temporally and spatially rich and complex pattern of activity remains at the scalp. In other words, these slow potentials show systematic variation in time across the extent of the clause and in space across the scalp in both the anterior-posterior and lateral dimensions.

These particular data were recorded as individuals read sentences presented one word at a time once every 500 ms for comprehension, which was probed after half of them with a true/false question. Seventy-two of these sentences began with a simple transitive clause consisting of an article, a noun, a verb, a noun, and a causative conjunction, in that order (e.g., "The secretary answered the phone because . . ."). Structurally, sentences cannot be much simpler than this. Moreover, unlike sentences with violations of grammar, semantics, or pragmatics or grammatically correct sentences with syntactic ambiguities or complicated structures, these sentences do not draw attention to themselves. Nonetheless, their comprehension requires analyses at what would be considered the orthographic, lexical, semantic, and syntactic levels

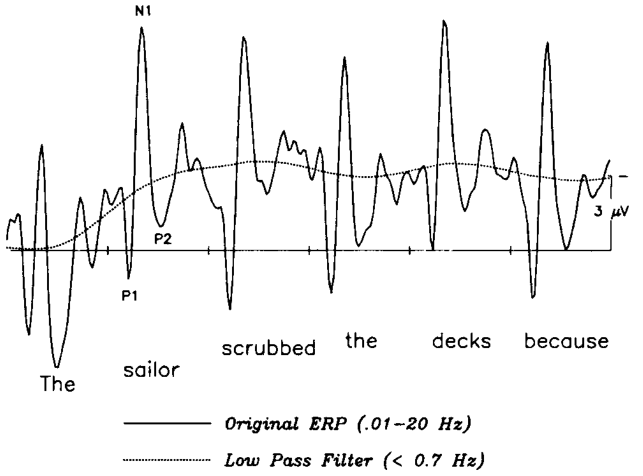

Figure 3. Grand-average ERP across the first six words of a transitive clause. The solid line represents the ERP averaged from EEG recordings taken over a left occipital site. The dotted line represents the same average after the application of a low-pass filter.

and must tap into something akin to a limited WM and a much larger capacity long-term memory.

Let us examine what, if any, aspects of sentence processing are evident in these cross-clause slow potentials across the scalp by focusing on analogous sites at the most lateral sites of the two hemispheres (Figure 4). The electrodes over the lateral occipitotemporal cortex cover brain regions that include areas crucial to the early processing of visual stimuli. Although it can be misleading to assume that electrical activity observed at a particular scalp location is generated directly beneath the recording electrode, evidence from other functional imaging (e.g., Petersen, Fox, Snyder, & Raichle, 1990; Tootell, Dale, Sereno, & Malach, 1996), monkey neurophysiology (e.g., Sereno, McDonald, & Allman, 1994), and human neuropsychology (e.g., Milner & Goodale, 1995) studies supports our working hypothesis that these potentials are related to the processing of visual features.

This sustained negativity over the occipital region is somewhat larger over the left than over the right hemisphere and appears to be insensitive to lexical features such as word class. It seems to reflect the continuous processing of the visual input. If this hypothesis is correct, then for spoken sentences we should see similar slow potentials with a more central distribution characteristic of early auditory processing (and potentials) generated in or near superior temporal gyrus (Näätänen, 1992; Woldorff et al., 1993).

At more anterior and temporal sites, both fast and slow activities are sensitive to lexical class (open, closed) and even higher level features of linguistic stimuli. In the low-pass filtered data, there is a phasic positivity beginning at Word 3 (i.e., the verb of the clause). Our current working hypothesis is that this phasic positivity reflects some aspect of thematic role assignment based on information contained in the lexical representation of the verb. Loosely, thematic role assignment refers to associating a verb with its different noun arguments (Dowry, 1986; Fillmore, 1968). If this

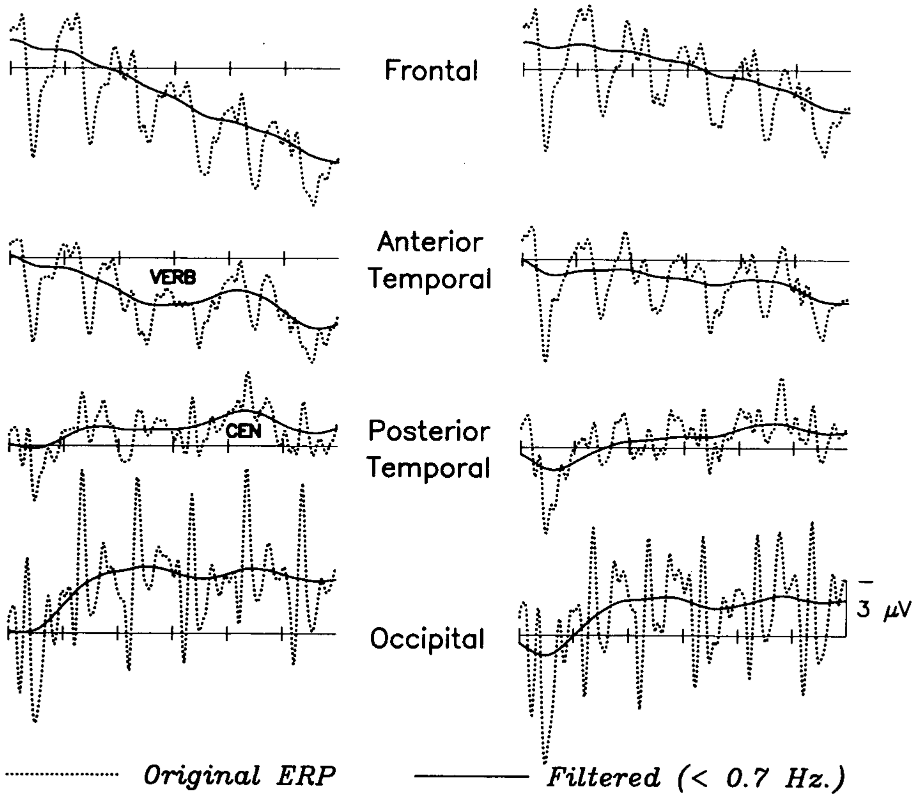

Figure 4. Superimposed are originally recorded cross-clause grand-average ERPs and the slow activity only (low pass filtered at 0.7 Hz) for four pairs of sites going from front to back of the head over left and right hemispheres, separately. The CEN labels the clause-ending negativity.

hypothesis is correct, then the amplitude of this potential should change systematically with the number of thematic roles verbs assigned. Specifically, we would expect to see variation in its amplitude when triggered by intransitive verbs, such as *sleep,* that assign one thematic role (the sleeper), versus transitive verbs, such as kill, that assign two thematic roles (a "killer" agent and an unfortunate individual who is killed), versus ditransitive verbs, such as *give,* that assign three thematic roles (an agent to give, a recipient, and something to be given).

Following the response to the verb is an asymmetric negativity better articulated over the left hemisphere, which we have dubbed the clause-ending negativity (CEN). The sensitivity of the brain to clause boundaries is consonant with the heavy demands made on WM processes. Both button press and eye movement data demonstrate that reading times are significantly slowed at clause boundaries (e.g., Aaronson & Ferres, 1986; Just & Carpenter, 1980). If we are correct about these hypothesized links between phasic verb positivity and thematic role assignments and between the CEN and WM operations at clause boundaries, then both effects also should be observed during the processing of spoken sentences.

One of the most striking effects in these cross-clausal data is the presence of an extremely slow (<0.2 Hz), cumulative positivity over frontal recording sites; although present over both hemispheres, the drift is significantly more positive over the left than over the right hemisphere. Its frontal maximum and the proposed role of the frontal lobes in executive functions of WM led us to propose that the frontal maximum may reflect some aspect of the integration of items in WM with information from long-term memory (e.g., Goldman-Rakic, 1987). Integration of this sort seems to be a prerequisite to comprehending sentences and to laying down a memory trace for later use, recognition, or recall. Such integrative processes should be more prolonged for more complex syntactic structures (sentences with embeddings) than for a string of transitive clauses. The amount of attentional resources devoted to integration may well be influenced by how much WM capacity is free; complex structures whose analysis is not aided by semantic and pragmatic cues are especially likely to leave little capacity for such integrative processes. The time course of these integrative processes also should be affected by how much a comprehender knows about the topic at hand and how readily that person can access that information. Again, we expect this potential to be independent of the input modality, if it is indeed related to integration of this type or to language processing in general as opposed to reading in particular.

In brief, I have touched on four laterally asymmetric slow potential effects that we have observed in sentence-level electrical recordings from individuals reading a transitive clause one word at a time: (a) a sustained negative shift over occipital sites reflecting early visual processing; (b) a positivity over temporal sites coincident with the verb, which may reflect some aspect of thematic role assignment; (c) a negativity coincident, with the clause ending and the associated wrap-up processes; and (d) a very slow going frontal positivity that may reflect the use of long-term memory to build a mental model, schema, or frame of the incoming sentence in WM. The very existence of these effects points to the utility of slow potentials as a means of investigating sentential structure. Moreover, the different spatial distributions of these different effects also highlight the distributed nature of aspects of sentence processing. With this technique, it is possible to monitor and tease apart some of the overlapping but different processes that take in multiple brain regions at the same time, although with different time courses.

Another observation during the course of this study was the notable difference in the pattern of effects for "good" and "poorer" comprehenders, where comprehension was inferred from how well individuals answered the true or false probes that followed the sentences. In this particular case, good comprehension means that the individual was more than 90% correct, and poorer comprehension refers to approximately 80% correct. Despite this relatively small difference in comprehension scores, we observed some notable ERP differences. Specifically, at occipital recording sites, the poorer comprehenders showed larger and more asymmetric early visual components (e.g., P1-N1-P2) on top of larger and more asymmetric slow potential activity. By contrast, at the prefrontal sites, the good, not the poorer, comprehenders showed a significantly larger and more asymmetric slow positive shift.

We also have examined these sentence level effects in sentences with more complex syntactic structures that describe more complex discourse relations such as those with embeddings (i.e., sentences within sentences). A case in point are sentences with relative clauses (Figure 5).

These two sentences are similar in that they both include a relative clause (surrounded by the dashed line), but they differ in the structure of the relative clause and in the role the main clause subject plays in the relative clause. Basically, they differ in whether it is the subject or the object of the relative clause that is missing. In subject relative constructions, the subject of the relative clause is missing: Somebody speedily rescued the cop. Who? The answer is same "fireman" who is the subject of the main clause. In object relative constructions, the object of the relative clause is missing: The cop speedily rescued somebody. Who was rescued? The answer is the same "fireman" who is the subject of the main clause. Thus, in the subject relatives, the "fireman" is the subject of both the main and the relative clauses; in the object relatives, the "fireman" is the subject of the main clause but the object of the relative clause.

Neither sentence type is ambiguous; nonetheless, object relatives are typically more difficult than the subject relatives, where difficulty is inferred from slowed reading times or comprehension errors (e.g., King & Just, 1991). The difficulties readers and listeners have with object relatives may stem from their greater taxing of WM processes. For example, object relatives require the reader to maintain an unattached constituent in WM for a longer duration. For example, in this case, *fireman* must be maintained over four words before it can be assigned as the direct object of *rescued.* This greater memory load is assumed to lead to greater processing difficulty, but there is no evidence in reading time data for this greater memory load until the end of the relative clause *(rescued),* just about when the load goes away. In fact, the largest reading time effects are typically observed right after, at the main clause verb *(sued).* Another potentially difficult aspect of object relatives is that they require multiple shifts in attentional focus for the comprehender between the discourse participants.

King and Kutas (1995a) recorded word-by-word and cross-sentence ERPs from people reading subject (e.g., "The reporter that harshly attacked the senator admitted the . . .") and object (e.g., "The reporter that the senator harshly attacked admitted the . . .") relatives presented one word every 500 ms. Participants were required to answer comprehension questions after a subset of sentences. Let us take a look at the ERPs at the location where the largest reaction time effects are generally observed, namely at the main clause verbs (see Figure 6). Not surprisingly, the ERP also shows a reliable effect at this point, which is manifested as a much greater negativity over left anterior sites (a LAN effect) for the

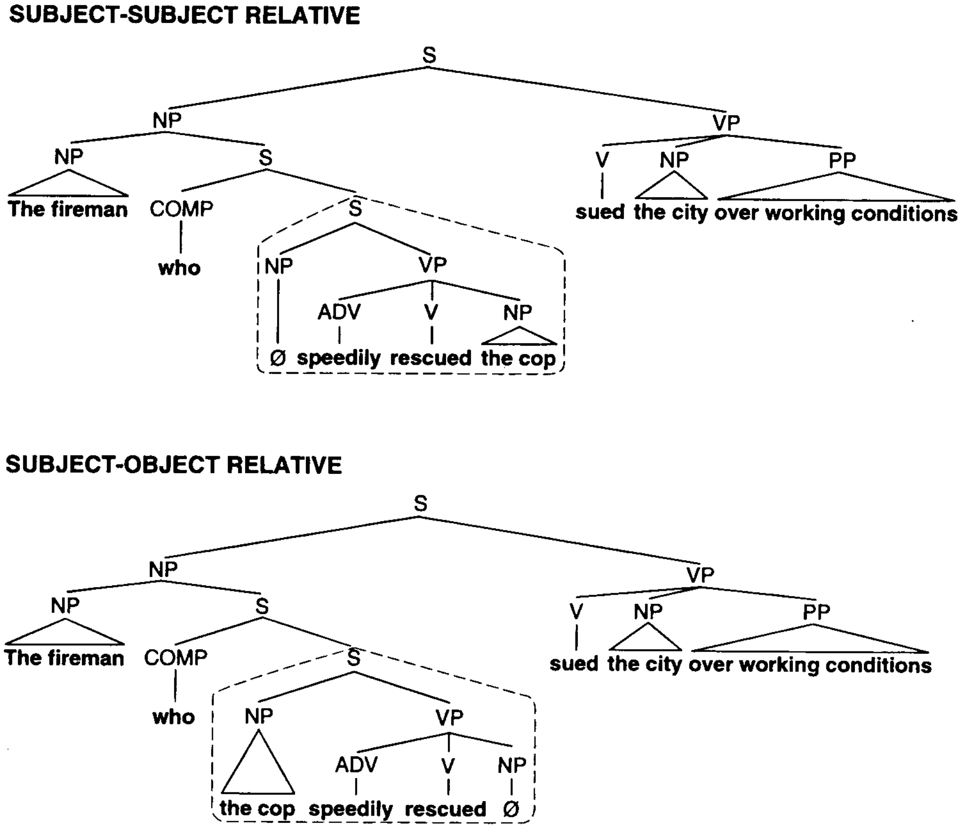

Figure 5. Tree diagrams showing the structure of sentences with an embedded relative clause. One example each is given for subject-subject relative (SS or SR) sentence and a subject-object relative (SO or OR) sentence.

object relatives. This type of effect was first described by Kluender (1991) for *wh* questions versus simple yes-no questions, where it takes a while before the reader knows who is the who in "Who do you think that . . ." (Kluender & Kutas, 1993a, 1993b). Kluender proposed that this LAN effect reflected holding unattached items (also known as fillers) such as the "who" in WM. In King's experiment, there is a LAN at the main clause verbs following both types of relative clauses when compared with the response to verbs in sentences that do not have any embedded clauses. Thus, the LAN effect appears to index some aspect of WM load. This particular finding meshes well with the reading time data and in some sense is no more informative.

But the comparison of the sentence-length ERPs shown in Figure 7 provides new information that reading time could not provide. In Figure 7, we see a processing difference between subject and object relatives much earlier in the sentence; specifically, this brain measure reflects a difference shortly after the reader encounters the relative clause. In other words, the sentence-level ERPs to the two types of relative clauses diverge as soon as there is a WM load difference, when the reader has to deal with the word cop without yet knowing what to do with the word fireman. The ERP difference is a sustained negativity over frontal and central sites for the object versus the subject relative constructions. There is a further ERP differentiation as the reader returns to the main clause, encounters the main clause verb *shed,* and can in fact finally decide who did what to whom.

A median split of good (87%) versus poorer (60%) comprehenders based on their comprehension scores also reveals strikingly different electrophysiological patterns. Among several differences, good comprehenders show a large difference between the subject and object relatives, whereas poorer comprehenders show almost none. Moreover, poorer comprehenders apparently find even the subject relative sentences demanding on their WM; in other words, their brains respond to subject and object relative sentences with prolonged negativities over frontal and central sites as if their WMs were loaded down even by subject relatives. The sentence-level ERPs of poorer comprehenders also do not exhibit the ultraslow frontal positivity characteristic of the response to the subject relatives in good comprehenders.

With these observations in hand, one of our main objectives was to investigate why some object relative sentences are easier to understand than others. However, critics have argued that reading sentences one word at a time every 500 ms is abnormally slow given that adult readers normally average about three to four words per second. The implication is that some, if not all, of the ERP differences observed between the subject and object relatives and between sentences with and without embedded clauses may be an artifact of this slow and abnormal presentation rate.

Moreover, even at a normal presentation rate, one could argue that reading one word at a time is not natural and this is simply reading, whereas real language, that for which the brain has been adapted, is speech. Unlike in reading, during speech, one can neither control the rate of input nor literally go backward in the speech stream (similar to a regressive eye movement in natural reading) to check an interpretation, for example. However, speech segmentation raises substantial theoretical and practical difficulties. Because ERP components elicited by natural or even semi-connected speech tend not to be very well articulated, ERP

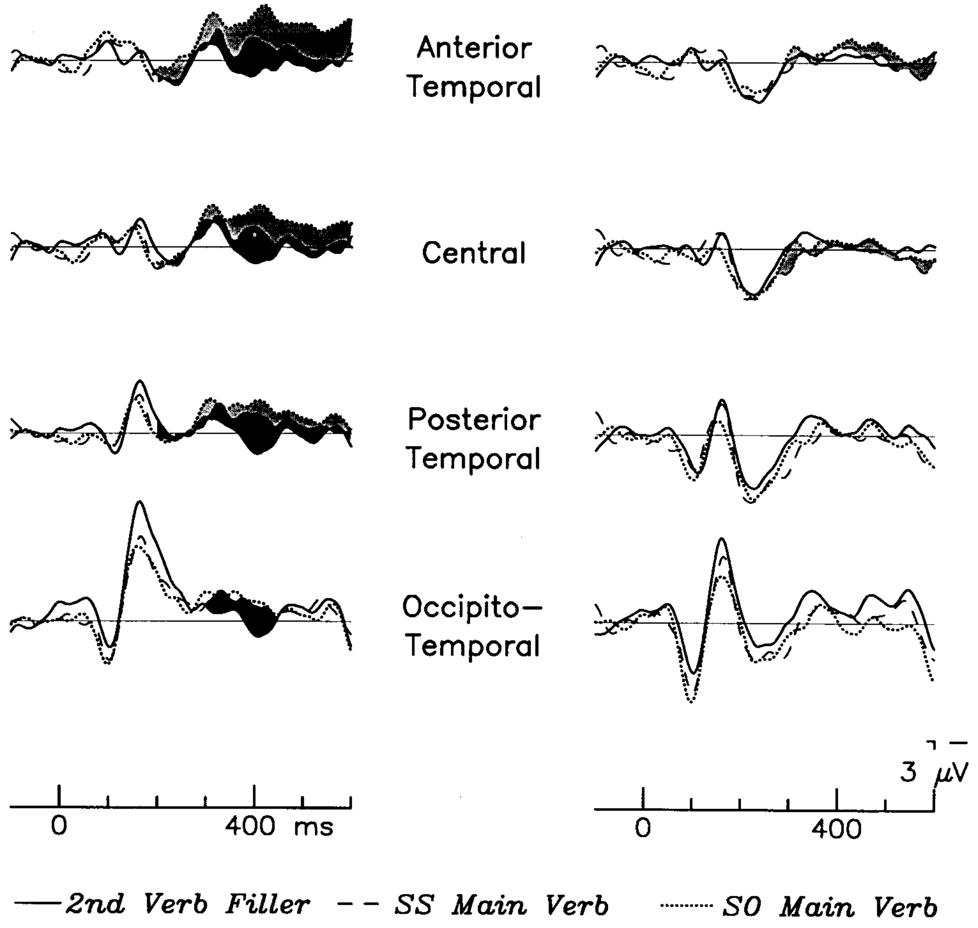

Figure 6. Grand-average ERPs to main clause verbs from subject relative (SS) and object relative (SO) sentences contrasted with the ERPs to second verbs in multiclausal sentences without embeddings ("The psychologist completed the experiment because she really wanted to know the answer"). The shaded differences in relative negativity over the more anterior sites is known as the LAN effect. The electrode pairs go from the front to the back of the head. Data are from King and Kutas (1995b).

researchers have shied away from their use until quite recently. However, because slow potentials were found to carry much of the information about sentence level processes, Mueller et al., *1997)* decided to replicate and extend their work with these subject and object relative sentences recorded as natural speech.

As expected, large ERP differences were observed between the subject and object relative sentences in the cross-sentence averages (Mueller et al., *1997).* The nature and timing of these differences were remarkably similar to the visual effects, although somewhat different in their distributions across the scalp. Specifically, the auditory effects were more widespread than the visual ones. In addition, the auditory effects exhibited a different lateral asymmetry, being larger over the right than over the left hemisphere. Over frontal sites, the ERPs revealed a very similar course of processing

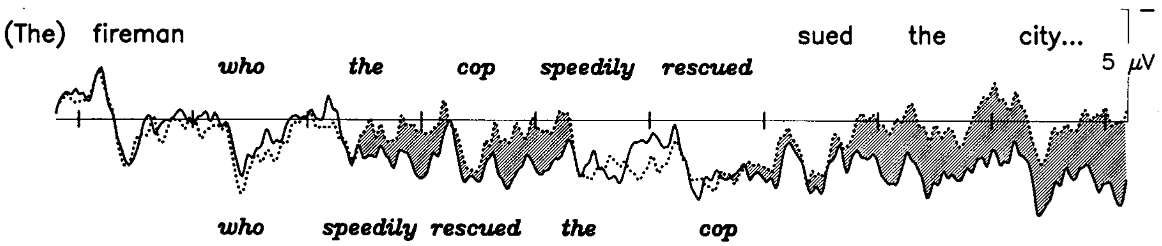

Figure 7. Comparison of the grand-average cross-sentence ERPs elicited by subject relative (solid line) and object relative (dotted line) sentences recorded over a left anterior location. The two sentence types are equivalent both before and after the relative clause. The relative clause above the baseline is an example of an object relative and that below the baseline is an example of a subject relative. Words were presented one at a time every 0.5 s for a duration of 200 ms. The shading represents the area of where object relative sentences are reliably more negative than the subject relative sentences.

of the relative clauses in the two modalities, suggesting that the effects reflect general language processes and not just purely reading or listening phenomena (see Figure 8).

Naturally, there are also some modality-specific effects. Thus, although slow potentials show that visual and auditory sentences generate dynamically similar effects, they are topographically distinct. Specifically, in the visual modality, there is a large resting negativity over occipitotemporal sites that is absent at central sites. By contrast, in the auditory modality, the opposite pattern obtains. Previously, we argued that the visual negativity is related to sustained word processing (Kutas & King, 1996), and we make a comparable argument for the auditory negativity in the present article. The pattern of ERP effects associated with comprehension skill in these data also parallel those observed for the word-by-word reading data (words presented once every 0.5 s; see Figure 9).

With these findings in hand, we have begun to delineate exactly what types of information (lexical, semantic, discourse, pragmatic) the brain is sensitive to while processing object relative constructions. In so doing, we aim to understand why some embedded sentences (e.g., "The student that the committee that the results surprised summoned listened closely") seem to be much harder to understand than others ("The painting that the artist that everyone loves painted hung in the living room"), despite having exactly the same syntactic structure and the same number of embeddings. Clearly, the difference must rest in the content or meaning of the words in the two sentences. However, theories of parsing, that is, theories that say who did what to whom (i.e., assign grammatical roles) differ considerably on when meaning is allowed to play any role and exactly what role it is allowed to play.

Weckerly and Elman (1992) pursued this question through different techniques. They first built a simple recurrent connectionist network in which the task was to predict the next word when presented with simple sentences and to predict the next word when presented with embedded clauses. Briefly, when this network was presented with sentences in which any noun could be linked to any verb, its predictions for words were based purely on structure (word order), but when the network was presented with sentences in which some nouns were more likely to be subjects than objects or vice versa and verbs were restricted as to which nouns could serve as their subjects or objects, then the network picked up on this nonstructural ("semantic") information to make its predictions and performed better (Weckerly & Elman, 1992).

These results suggested that the animacy of a noun may be one of the semantic factors that influences the ease of processing object relatives. This notion was tested empirically in three different ways. Some normative data were collected; undergraduates were asked to write simple transitive sentences given a pair of nouns, with no restrictions on which noun was to play which role. The results clearly showed that when only one of the one nouns was animate, that noun was chosen as a subject about 80% of the time. In contrast, when both nouns were animate, either was equally likely to be chosen as a subject (Weckerly, 1995). Thus, animate nouns seem to be favored as subjects of sentences, and we should not be surprised to see that the brain is sensitive to this correlation and uses it during sentence processing in real time.

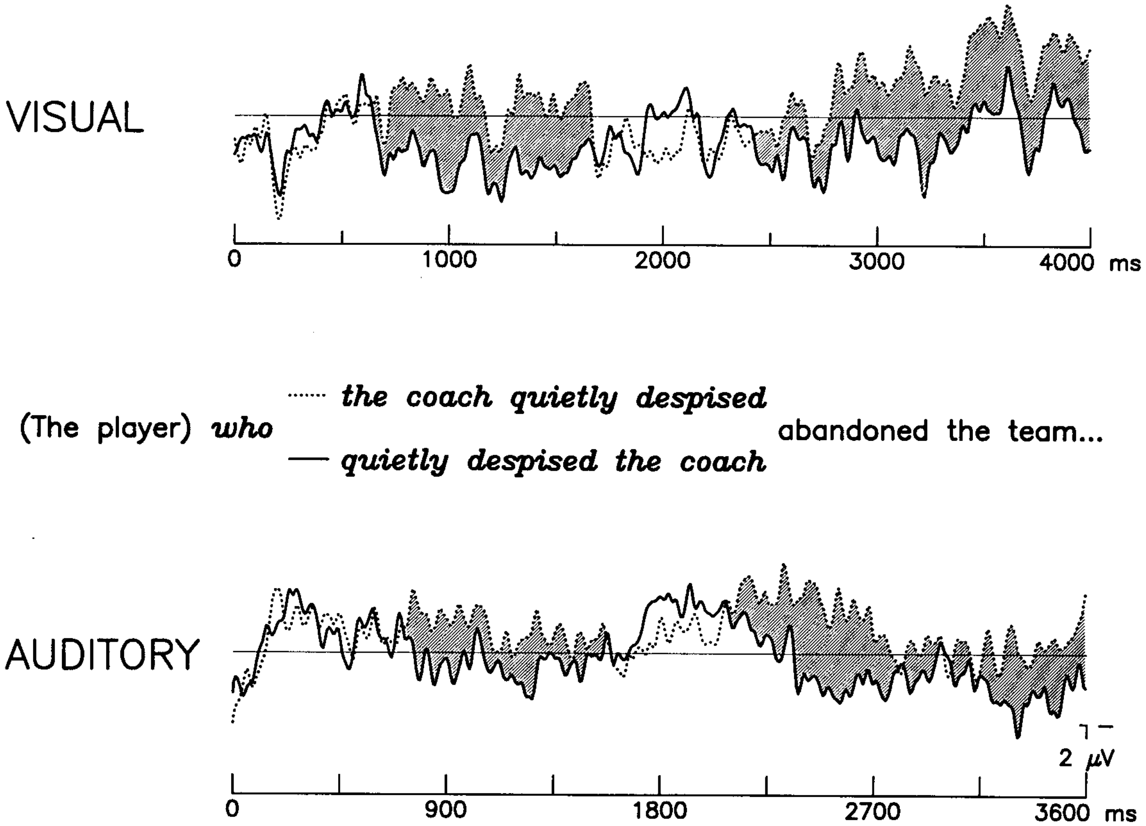

Figure 8. Grand-average ERPs recorded from a left anterior site in response to subject and object relative sentences during reading for comprehension (one word every 500 ms) or listening for comprehension (natural speech). Visual data are from King and Kutas (1995), and the auditory data are from Mueller, King, and Kutas (1997).

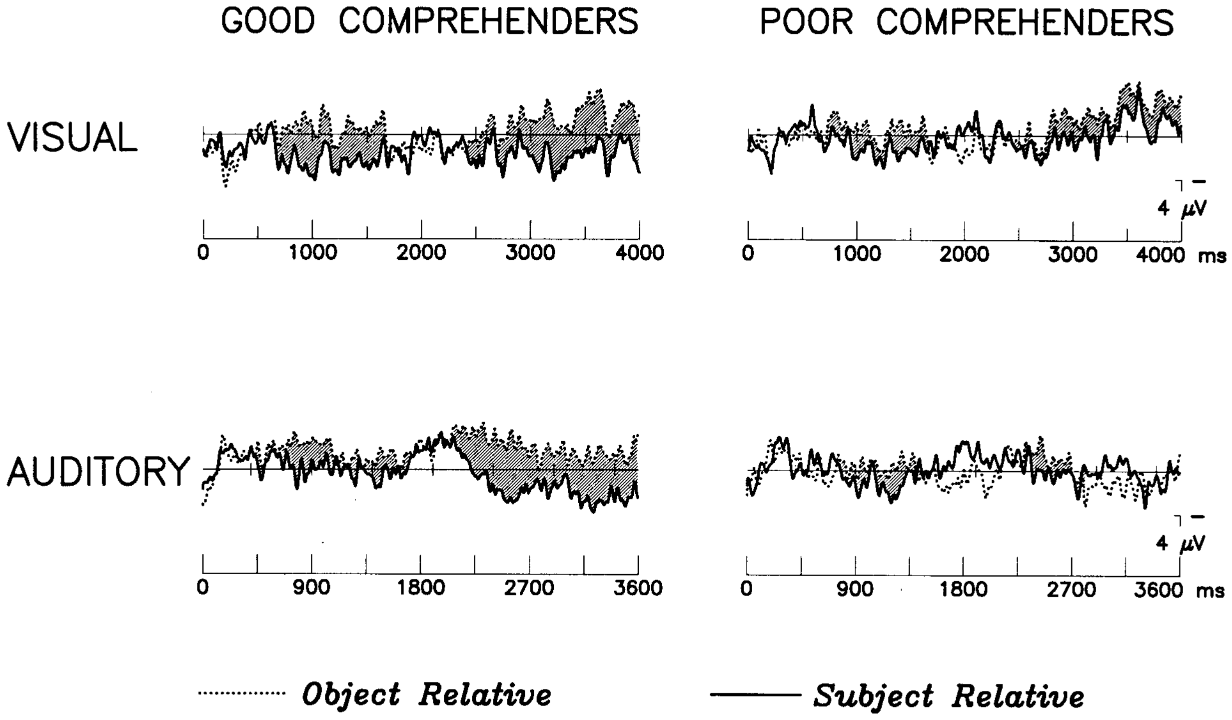

Figure 9. Grand-average ERPs to subject and object relative sentences from a left anterior site in good and poorer comprehenders. The visual sentences were presented one word every 500 ms, and the auditory sentences were presented as natural speech. The visual data are from King and Kutas (1995), and the auditory data are from Mueller, King, and Kutas (1997).

This finding was used to construct some critical sentences with two clauses each. The two types of clauses were combined to construct object relative sentences. One clause had an animate subject and an inanimate object, and the other had an inanimate subject and an animate object. These two clause types were combined into two sentence types, which were called I(A) and A(I), in which the first letter refers to the animacy of the noun in the main clause and the second letter (in parentheses) refers to the animacy of the noun in the relative clause.

I(A): The poetry that the editor recognized depressed the publisher of the struggling magazine.

A(I): The editor that the poetry depressed recognized the publisher of the struggling magazine.

These sentences have exactly the same words but in a different order.

In a word-by-word reading time study with these materials, the reading comprehension was about 10% better for sentences without embeddings (79%) than for either of the object relative sentences (63-68%). In addition, although the only difference between the critical sentences was the order of the noun phrases, there were some reliable reading time differences. Starting at the relative clause subject, I(A) sentences, that is, those with an animate noun as the relative clause subject, were read faster than A(I) sentences (Weckerly, 1995). This reading time advantage lasted across five words (through the direct object of the main clause). Thus, people's brains seem to be sensitive to noun animacy and may use its correlation with subjecthood to aid in comprehension. The brain may use the animacy configuration of the noun phrase to begin to assign grammatical roles such as subjects and objects, that is, to interpret syntactic structure even before the verb is encountered. Thus, these results are at odds with traditional theories of parsing wherein role assignments are contingent on a the presence of a verb. Nor do they fit neatly with any theory that maintains that initial parsing decisions proceed completely blind to the semantic content of the sentence elements.

Clearly, the reading time data have brought us far, but one reaction time difference is no different than another without an underlying theoretical account. By looking at the brain activity, we can determine whether there is any need to invoke different mental operations to account for the reading time differences observed at the different sentence locations. Therefore, these same materials were presented to a new set of people one at a time at a rate of one word every 500 ms.

Before perusing the data, let us consider what pattern of results could be expected based on what is known about language-related ERP components and what they are sensitive to. There should be some effect of animacy, although there is a priori prediction as to what form this might take in the ERP (Li & Thompson, 1976).

If readers use animacy information to bias role assignment, then they should be less disturbed by an animate relative clause subject in I(A) sentences in which the inanimacy of the main clause noun followed by the word *that* have suggested an object relative construction than by the inanimate subject of the relative clause in A(I) sentences.[2] Accordingly, given the sensitivity of the N400 to semantic expectancy and the difficulty of integrating a

[2] This is a simplified version of the different predictions. More detail can be found in the thesis by Weckerly (1995) and the study by Weckerly and Kutas (1997).

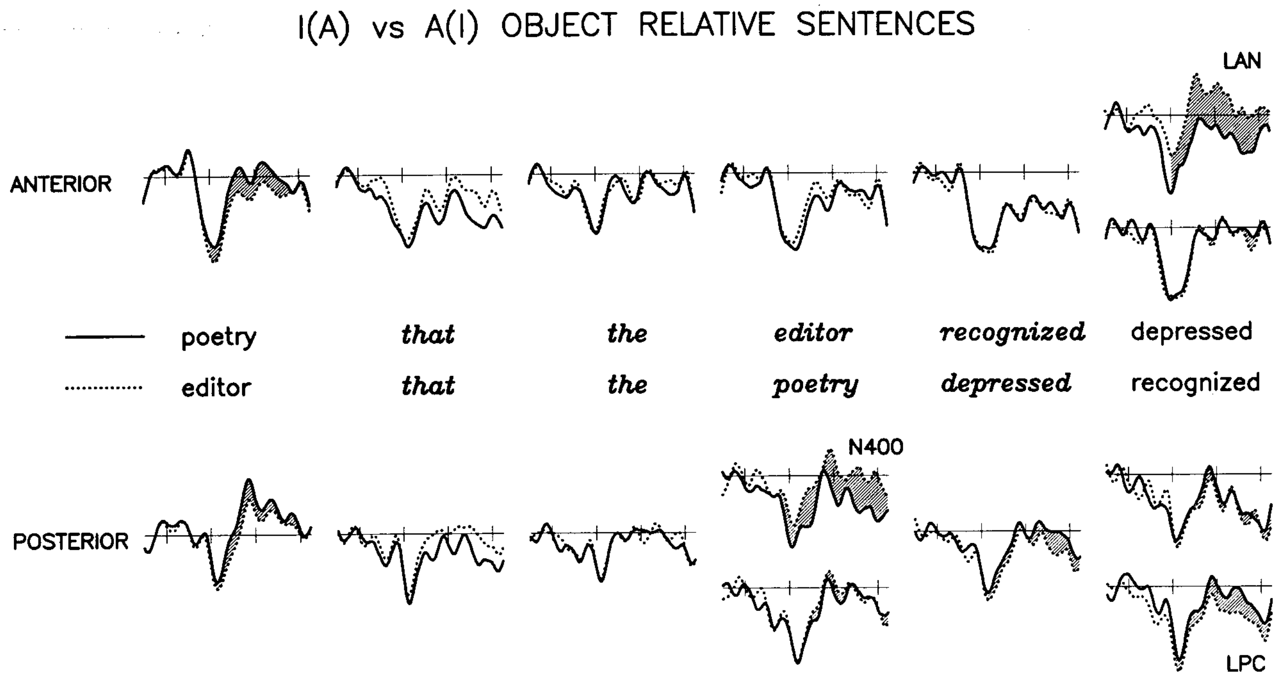

Figure 10. Comparison of grand-average ERPs to six words from the A(l) and [(A) sentences at one anterior and one posterior electrode site. From left to right, the responses are for the main clause subject, the complementizer (that), the article to the relative clause subject, the relative clause subject, and the relative and main clause verbs, respectively. The response to the relative clause subject at the posterior site and to the main clause verb at anterior and posterior sites are shown separately for the good (top) and poor (bottom) comprehenders, respectively. LAN = left anterior negativity, LPC = late positive component.

word into a sentential context, there should be a larger N400 for the A(I) sentences at this point, namely the relative clause subject.

The work of Kluender (1991) and King and Kutas (1995b) shows that WM loads are associated with the LAN effect, which is larger for greater loads. Such WM loads should be especially high at the main clause verb because by this point both nouns and both verbs are available and grammatical role assignments can be completed.

There is much evidence in the ERP literature that loci of grammatical difficulty are often, although not always, associated with a late positivity (P600). Thus, some enhanced late positivity should be expected at a number of locations in the sentence but certainly at or near both verbs.

A cursory look at the ERP data reveals that the different sentence-level effects are replicated, including the ultraslow positivity over frontal sites, the phasic positivity associated with verbs, the clause ending negativity, and LAN effects. More importantly, the data show that the brain is sensitive to noun animacy (see Figure 10). Within 200 ms of the presentation of the word, there is a greater negativity for inanimate than for animate nouns. Clearly, the brain has registered noun animacy at some level;[3] thus, the information is there to be used. This is true for both good and poorer comprehenders.

The next big effect, however, is only present in good comprehenders. These individuals seem to be surprised by the inanimate subject of the relative clause in A(I) sentences; they show a larger N400 to it than to the animate subject of the relative clause in I(A) sentences. This result is consistent with our hypothesis that animacy (along with other linguistic information) is used to generate expectations about the upcoming structure and the lexical items that might fill it in general terms.

What about at the verbs? It was here that we had seen the largest reading time differences. In addition, there are both greater negativity over left frontal sites (LAN effects) and greater late positivity over posterior sites (P600 effects), albeit to differing degrees depending on comprehension ability. Good comprehenders showed a large LAN to the main clause verb in the more difficult sentence type, the A(I)s. There also was a greater late positivity to both A(I) verbs versus I(A) verbs, which may reflect the ongoing difficulties of processing A(I) relative clauses; this effect is larger in the poorer comprehenders.[4]

In summary, the nature of the ERP effects across object relative sentences do change, thus confirming the suspicion that there are different causes for the slower reading times at different positions in A(I) versus I(A) sentences. Specifically, (a) the animacy of a noun is registered within 200 ms of its availability, which is three word positions earlier than any detectable difference in reading time measures; (b) the N400 to the relative clause subject supports

---

[3] The sensitivity of the brain may be to animacy per se or may be a response to the fulfillment or violation of the expectancy for the first noun in an English-language sentence to be animate or some combination thereof, which is manifest to different degrees in different individuals. In our data, good comprehenders seem to respond to inanimate initial nouns with an N400 response, whereas poorer comprehenders elicit an overall smaller negativity that is somewhat more frontally distributed.

[4] Comprehension skill was defined by how well the participants answered different queries following the object relative sentences. Good comprehenders were more than 75% accurate, whereas poorer comprehenders fell below this cutoff.

our hypothesis that readers had expectations for an animate noun; (c) all of our data are consistent with animacy being used in making role assignments that begin before the verb and extend past the verb that relates the two nouns; (d) the electrophysiological data also show that multiple processes can occur simultaneously in different brain regions as reflected in the left anterior negativity and P600 to verbs; this would be difficult to tease out of reaction time measures alone. All in all, the results of these experiments indicate that role assignments are neither punctate nor time locked to the verb and that nonsyntactic information is used in the initial syntactic analysis as immediately as possible.

In conclusion, a cognitive neuroscience approach to language can be revealing insofar as it includes constraints from linguistics, psychology (including but not solely psycholinguistics), and brain anatomy and physiology on experimental design, data interpretation, and theory. More specifically, I have tried to demonstrate by example how it is possible to track the processing of written and spoken language and to delineate the factors that influence comprehension with scalp-recorded ERPs. Many more than two, putatively undifferentiated, areas in the left hemisphere are involved in language processing. Many aspects of language processing do not appear to be language specific, are not invariant across people (individual differences abound), and do overlap in space and time, such that analysis of both dimensions is needed to tease them apart. Language processing is not only fast but also includes appreciation of structure at multiple levels; in fact, ERPs can help validate the existence of these linguistic levels of analysis and indicate their time course of involvement and disclose the factors in language and the world that determine how an individual comes to know that not only this sentence but also this article has come to an — .

REFERENCES

Aaronson, D., & Ferres, S. (1986). Reading strategies for children and adults: A quantitative model. *Psychological Review,* 93, 89-112.

Bell, A. J., & Sejnowski, T. l. (1995). An information maximization approach to blind separation and blind deconvolution. *Neural Computation,* 7, 1129-1159.

Besson, M., Kutas, M., & Van Petten, C. (1992). & semantic, An event related potential (ERP) analysis of congruity and repetition effects in sentences. *Journal of Cognitive Neuroscience, 4(2),* 132-149.

Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language,* 31, 429-460.

Butler, S. R., & Glass, A. (1976). EEG correlates of cerebral dominance. In A. H. Riesen & R. F. Thompson (Eds.), *Advances in psychobiology* (pp. 219-272). New York: John Wiley & Sons.

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

Coulson, S., King, J. W., & Kutas, M. (in press). Expect the unexpected: Event-related brain responses to morphosyntactic violations. *Language and Cognitive Processes.*

Dale, A. M., & Sereno, M. I. (1993). Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience, 5(2),* 162-176.

Deacon, T. W. (1992). Cortical connections of the inferior arcuate sulcus cortex in the macaque brain. *Brain Research,* 573, 8-26.

Donchin, E., Ritter, W., & McCallum, W. C. (1978). Cognitive psychophysiology: The endogenous components of the ERP. In E. Callaway, P. Tueting, & S. H. Koslow (Eds.), *Event-related brain potentials in man (pp.* 349-441). New York: Academic Press.

Dowty, D. (1986). Thematic roles and semantics. *Berkeley Linguistics Society,* 12, 340-354.

Fillmore, C. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory (pp.* 1-88). New York: Holt, Rinehart and Winston.

Ganis, G., Kutas, M., & Sereno, M. I. (1996). The search for "common sense": An electrophysiological study of the comprehension of words and pictures in reading. *Journal of Cognitive Neuroscience,* 8, 89-106.

Geschwind, N. (1965). The organization of language and the brain. *Science,* 170, 140-144.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum & V. Mountcastle (Eds.), *Handbook of physiology: The nervous system (pp.* 373-417). Bethesda, MD: American Physiological Society.

Gunter, T. C., Jackson, J. L., & Mulder, G. (1996). Language, memory, and aging-An electrophysiological exploration of the N400 during reading of memory-demanding sentences. *Psychophysiology,* 32, 215-229.

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

Holcomb, P J., & Neville, H. J. (1991). Natural speech processing: An analysis using event-related brain potentials. *Psychobiology,* 19, 286-300.

Just, M. A., & Carpenter, P A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review,* 87, 329-354.

Just, M. A., Carpenter, P. A., & Keller, T. A. (1996). The capacity theory of comprehension-New frontiers of evidence and arguments. *Psychological Review,* 103, 773-780.

King, J., & Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language,* 30, 580-602.

King, J. W., & Kutas, M. (1995a). A brain potential whose latency indexes the length and frequency of words. *Newsletter of the Center for Research in Language,* 10(2), 3-9.

King, J. W., & Kutas, M. (1995b). Who did what and when? Using word- and clause-related ERPs to monitor working memory usage in reading. *Journal of Cognitive Neuroscience,* 7, 378-397.

Kluender, R. E. (1991). *Cognitive constraints on variables in syntax.* Ph.D. dissertation, University of California, San Diego, La Jolla, CA.

Kluender, R., & Kutas, M. (1993a). Bridging the gap: Evidence from ERPs on the processing of unbounded dependencies. *Journal of Cognitive Neuroscience,* 5, 196-214.

Kluender, R., & Kutas, M. (1993b). Subjacency as a processing phenomenon. *Language and Cognitive Processes, 8(4),* 573-633.

Kutas, M. (1987). Event-related brain potentials (ERPs) elicited during rapid serial visual presentation of congruous and incongruous sentences. *Electroencephalography and Clinical Neurophysiology, 40* (Suppl.), 406-411.

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

Kutas, M., & Hillyard, S. A. (1983). Event-related brain potentials to grammatical errors and semantic anomalies. *Memory and Cognition,* 11(5), 539-550.

Kutas, M., & Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature,* 307, 161-163.

Kutas, M., & King, J. W. (1996). The potentials for basic sentence processing: Differentiating integrative processes. In I. Ikeda & J. L. McClelland (Eds.), *Attention and performance (Vol.* 16, pp. 501-546). Cambridge, MA: MIT Press.

Kutas, M., & Van Petten, C. (1994). Psycholinguistics electrified. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics (pp.* 83-143). San Diego: Academic Press.

Li, C. N., & Thompson, S. A. (1976). Subject and topic: A new topology of language. In N. Li (Ed.), *Subject and topic.* New York: Academic Press.

Lichtheim, L. (1985). On aphasia. *Brain,* 7, 433-484.

Luders, H., Lesser, R. P., Hahn, J., Dinner, D. S., Moms, H., Resor, S., & Harrison, M. (1986). Basal temporal language area demonstrated by electrical stimulation. *Neurology,* 36, 505-510.

McCallum, W. C., Farmer, S. F., & Pocock, P. V. (1984). The effects of physical and semantic incongruities on auditory event-related potentials. *Electroencephalography and Clinical Neurophysiology,* 59, 477-488.

McCarthy, G., Nobre, A. C., Bentin, S., & Spencer, D. D. (1995). Language-related field potentials in the anterior-medial temporal lobe: I. Intracranial distribution and neural generators. *Journal of Neuroscience,* 15, 1080-1089.

Milner, A. D., & Goodale, M.A. (1995). *The visual brain in action.* New York: Oxford University Press.

Mueller, H. M., King, J. W., & Kutas, M. *(1997).* Event-related potentials to relative clause processing in spoken sentences. *Cognitive Brain Research, 5(3), 193-203.*

Muente, T F., Matzke, M., & Johannes, S. *(1997).* Brain activity associated with syntactic incongruencies in words and pseudo-words. *Journal of Cognitive Neuroscience, 9, 300-311.*

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

Neville, H. J., Mills, D. L., & Lawson, D. S. *(1992).* Fractionating language: Different neural subsystems with different sensitive periods. *Cerebral Cortex, 2(3), 244-58.*

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

Nobre, A. C., Allison, T., & McCarthy, G. *(1994).* Word recognition in the human inferior temporal lobe. *Nature, 372, 260-263.*

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

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

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

Petersen, S. E., Fox, P. T., Snyder, A. Z., & Raichle, M. E. *(1990).* Activation of extrastriate and frontal cortical areas by visual words and word-like stimuli. *Science, 249, 1041-1043.*

Roland, P. E. *(1993). Brain activation.* New York: Wiley-Liss.

Schwartz, T. J., Kutas, M., Salmon, D. P., Paulsen, J. S., & Butters, N. *(1996).* Electrophysiological insights into the nature of the semantic deficit in Alzheimer's Disease. *Neuropsychologia, 34, 827-841.*

Sereno, M. I., McDonald, C. T, & Allman, J. M. *(1994).* Analysis of retinotopic maps in extrastriate cortex. *Cerebral Cortex, 4, 601-620.*

Tootell, R. B. H., Dale, A. M., Sereno, M. I., & Malach, R. *(1996).* New images from human visual cortex. *Trends in Neurosciences, 19,481-489.*

Van Petten, C., & Kutas, M. *(1990).* Interactions between sentence context and word frequency in event-related brain potentials. *Memory and Cognition, 18, 380-393.*

Van Petten, C., & Kutas, M. *(1991).* Influences of semantic and syntactic context on open and closed class words. *Memory and Cognition, 19, 95-112.*

Waters, G. S., & Caplan, D. *(1996).* The capacity theory of sentence comprehension-Critique of Just and Carpenter *(1992). Psychological Review, 103, 761-772.*

Weckerly, J. *(1995). Object relatives viewed through behavioral, electrophysiological, and modeling techniques.* Ph.D. dissertation, University of California, San Diego, La Jolla, CA.

Weckerly, J., & Elman, J. *(1992).* A PDP approach to processing center-embedded sentences. In Hillsdale (Ed.), *Proceedings of the 14th annual conference of the Cognitive Neuroscience Society*. Hillsdale, NJ: Erlbaum.

Weckerly, J., & Kutas, M. *(1997). The interaction of grammatical role and noun animacy in object relatives.* Manuscript in preparation.

Wernicke, C. *(1874). Der aphasische symptoencomplex*. Breslau: Max Cohn & Weigert.

White, E. I. *(1989). Cortical circuits: Synaptic organization of the cerebral cortex-Structure, function and theory*. Boston: Birkhauser.

Willmes, K., & Poeck, K. *(1993). To* what extent can aphasic syndromes be localized? *Brain, 116, 1527-1540.*

Woldorff, M. G., Gallen, C. C., Hampson, S. A., Hillyard, S. A., Pantev, C., Sobel, D., & Bloom, F. E. *(1993). Modulation of early sensory pro-*cessing in human auditory cortex during auditory selective attention. *Proceedings of the National Academy of Sciences of the United States of America, 90(18), 8722-8726.*

Yvert, B., Bertrand, O., Echallier, J. F., & Pernier, J. *(1996).* Improved dipole localization using local mesh refinement of realistic head geometries-An EEG simulation study. *Electroencephalography and Clinical Neurophysiology, 99, 79-89.*

(ACCEPTED March *7, 1997;* ACCEPTED May *23, 1997)*