Cahiers de Psychology Cognitive/
Current Psychology of Cognition
1998, 17 (45), 951-969

Current thinking on language structures

Marta Kutas

University of California, La Jolla, U.S.A.

Abstract

There are multiple language processing areas that are widely but not randomly distributed in the brain. To map these brain areas onto language functions in real time requires knowing not only how they are distributed in space but how they are coordinated in time. Any psychological theory of language processing must include a description of the mental representations involved, computations on them, the timecourse of their operations and interactions. The method of choice for precise temporal tracking of sentence comprehension is the event-related brain potential (ERP) technique. Its application to the analysis of written and spoken sentences of varying structural complexity reveals the critical roles of working memory operations and background knowledge in normal comprehension. Such data show that many of the processes of language include large parts of the brain, are neither modality nor language-specific, are subject to individual differences, and overlap each other in neural and mental space and time, such that analyses of both dimensions are needed to tease them apart so as to understand how we understand.

Key words: ERPs, sentence processing, working memory.

Correspondence should be sent to Marta Kutas, Department of Cognitive Science, UCSD, La Jolla, California 92093-0515, U.S.A. (e-mail: mkutas@ucsd.edu).

Language is not just sandwiched
between Broca's and Wernicke's.
Heard, spoken or seen
meaning emerges from activities
in multiple sites
each scheming by all rights
at much the same time
to yield both the literal
and the metaphorical,
the banal
and the sublime.

INTRODUCTION

What is current thinking on language structures? The answer to this question is not as obvious as it may seem, as certainly in isolation and sometimes even within context, language input is ambiguous at many levels. In fact, I will assert that it is pointless to view language simply as a code which combines context-invariant meanings of words via a set of rules (grammar); rather it makes more sense to view language as providing clues to the language user for flexibly addressing his/her knowledge base so as to make sense of a particular utterance within the current context (for elaboration see Coulson, 1997). Without some background knowledge, in other words, more context than is literally a part of this article's title, no one can know exactly what meaning its author intended. The title "current thinking on language structures" is ambiguous in meaning. Moreover, several of its individual words are lexically ambiguous. For example, is "current" serving as an adjective or a noun in this case? Which meaning was intended - the swiftest part of a stream or a flow of electric charge? And, what about "structures", does it refer to something constructed or to the organization of the parts as dominated by the general character of the whole? This is not to mention "on", which many people may not consider ambiguous until they look it up in the dictionary only to find over twenty different definitions.

How does the reader's brain choose and the writer's mind come to know which of these dictionary entries is the right one? In other words

- which "current", which "structures", and which "on" are to be accessed and recombined to yield the title's intended meaning? Perhaps it is those structures in the brain that turn changes in air pressure on the eardrums into a meaningful sound as well as those that allow us to know that "I am pleased as punch to be here" does not mean that I am thirsty (for punch) or that I want to be punched (i.e., hit).

Language is an exquisite medium for passing along propositions, for talking about the past and the future, and for getting people to do or to think what you want them to do or think, respectively. But decoding linguistic input, especially in real time, is a difficult analytic problem. Human brains are massively parallel in their processing but linguistic input (speech) enters as an essentially serial stream of acoustic inputs: words come in one at a time, perhaps with a little forewarning (e.g., co-articulation) but more often than not it is necessary to hold onto parts of the input, i.e., to wait for more than a word or two in order to figure out the structure and the meaning of an utterance or a line of text. Perhaps, this is the language structure to which my title refers.

So, what are some current thoughts on language structures? Some current thinking on language (brain) structures is that the traditional view of Wernicke's area in the temporal lobe near its junction with the parietal lobe to take in linguistic information and Broca's area in the frontal cortex of the left hemisphere to control its output needs to be modified (Damasio, 1997). However one defines language, the cognitive neuroscience literature attests to the involvement of many more than two areas of the brain in understanding a single sentence much less an entire discourse. Areas in the left and right hemispheres in both anterior and posterior regions of the brain have been implicated in some aspect of language processing (e.g., Binder, Frost, Hammeke, Cox, Rao, & Prieto 1997). Damage somewhere in the peri-sylvian region of the left hemisphere can be devastating to certain aspects of language, especially speech. But damage to the right hemisphere seriously compromises an individual's ability to appreciate the meaning of an indirect request, a metaphor, an idiom or a joke (Joanette & Brownell, 1990; Joanette, Goulet, & Hannequin, 1990). The traditional model of language representation in the brain is undergoing fine-tuning and expansion based on what is now known about the size of brain areas, about various language functions and brain activity and/or damage to the brain, and about the functions of language as well as about various linguistic phenomena.

For instance, it is quite unlikely that what has traditionally been called Broca's area refers to a single brain area. Based on the numbers and sizes of visual, auditory, somatosensory and motor areas in macaques, the average brain area in the human probably covers about 1000 mm^2 of cortex. If what is typically labeled Broca's area does in fact subsume 4 to 5 distinct brain areas, it is easier to explain the substantial variability in the behavioral symptoms reportedly observed after damage to this region. Note that a similar argument could be made for Wernicke's area. The newer brain imaging techniques such as computerized axial tomography (CAT) and magnetic resonance imaging (MRI) scans have made it possible to define the exact areas damaged by a stroke, and thus have forced neuropsychologists to reconsider the view that damage to Broca's (and no other) area always leads to Broca's aphasia. In fact, brain images such as by positron emission tomography (PET) have revealed that functional damage often extends beyond the boundaries of the observed structural damage. Some researchers argue that this is simply an inferential limit inherent in neuropsychological data which therefore can be overcome by measuring brain activity instead of brain damage or by more localizing the damaged areas with greater anatomical precision. But, in fact, we can never make the inferential part of the mapping problem vanish fully. FMRI and PET data map brain activity, not linguistic or psychological functions, and it is function that we hope to specify so as to account for the understanding and/or production of utterances or reading of texts.

In fact, many researchers do seem to have abandoned the simple one-to-one mapping between Broca's area and production and Wernicke's area and comprehension. But, this view has been replaced with the equally simplistic view that Broca's area is responsible for syntax and Wernicke's area is responsible for semantics, or Broca's area is mainly involved in rule-based processing (e.g., regular verbs) whereas Wernicke is a memory store for exceptions (e.g., irregular verbs). Naturally, there are some data consistent with each of these alternatives. Thus we come back to our original question - what functions does Broca's area perform? This is a difficult question to answer given our current state of knowledge, but I suggest that knowing that Broca's area probably comprises four or five different areas rather than one will help to constrain hypotheses about the functions of its different subareas (Deacon, 1996).

Another language-sensitive area identified recently is the Basal Temporal Language Area at the bottom of the temporal lobe (e.g., Luders, Lesser, Hahn, Dinner, Morris, Resor, & Harrison, 1986). It is in this general area (anterior fusiform gyrus) within the brains of individuals with epilepsy that Nobre, Allison, and McCarthy (1994) observed a large electrical response whose amplitude was sensitive to semantically anomalous words in sentence context. Nobre et al.'s data also revealed that cells in both hemispheres respond not just to faces but also to words both in and out of context, lending support to the hypothesis that the basal temporal regions are involved in language processing.

In summary, there appear to be multiple "language areas" distributed throughout the brain. We might ask what the functions of these areas are and whether these functions are specific to language processing. Clearly, in order to link these areas with language processing functions in real time it is necessary to determine not only where they are located (in space) but also their temporal characteristics (duration and order of involvement). Most psycholinguists agree that different types of linguistic and non-linguistic information can influence how sentences are understood, but they disagree on when and how different sorts of information are used. In other words, they disagree on the temporal structure of processing, and on which processes work on input from only one level of analysis and which are typically influenced by input from more than one level. To choose among the existing hypotheses, psycholinguists need to know more about the temporal coordination of all the (distributed) brain areas involved in comprehension and production. This is where event-related brain potential (ERP) methodology can be useful. If the average ERPs from two experimental conditions differ reliably at any given point in time, it can then be inferred that the associated brain and mental activity also differ at least by that point. In fact, the onset latency of the ERP difference can serve as an upper limit for the minimum time it takes the brain to differentiate two or more classes of stimuli.

ERPs provide an exquisite measure of the brain's activity at any moment across scales from milliseconds to tens of seconds. Its temporal resolution ranges from the crucial milliseconds that differentiate one phoneme from another to the more extended time course needed to determine that "pencilist" is not a word, or that "pencils" does not fit with the sense of the phrase "The pigs were herded into their pencils" as well as the seconds that may be required to determine who did what to whom after reading ... "The American that the Frenchman that the Russian telephoned perplexed contacted the Spanish electrophysiologist about running a study". ERP records can tell cognitive neuroscientists *when* certain events in the brain take place. And, as this last example implies, ERPs can provide such timing information beyond the level of the single word. With the exception of the magnetoencephalogram (MEG), no other currently-used neuroimaging technique has the temporal precision of the ERP technique. Neither PET nor FMRI methods will ever attain the resolution of electromagnetic recordings because they are based on hemodynamic changes, which occur on the order of a second or more. By contrast, both PET and FMRI have a higher spatial resolution than ERP/MEG although electromagnetic methods have gained considerable ground in this arena with the advent of high EEG resolution methods (see Kutas & Federmeier, 1998, for a review). Moreover it is important to note that measures dependent on hemodynamic or metabolic changes are indirect reflections of neural activity whereas the EEG/MEG are direct indices. Recently it has been suggested that optical imaging may someday rival both ERPs/MEG and PET/FMRI if it can be determined that the optical changes reflect changes in neural activity per se (Gratton, Fabiani, & Corballis, 1997).

It is unlikely that you would tell someone who couldn't make it to a talk what it was all about via a string of isolated words. Language has structure - the very structure that Chomsky (1966) said could never be accounted for strictly in terms of chains of stimulus-response pairings in his devastating review of Skinner's *Verbal Behavior.* The exact same words occurring in different orders mean something different. Thus, it is the structure of language that makes "Kutas understood Jean" different from "Jean understood Kutas" but similar to "Kutas was understood by Jean" and "understood Jean was Kutas" not a meaningful sentence at all, and "Kutas who you understand understands that you understand her understanding" not so easy to understand. Structures have processing consequences which make some sentences easier to deal with than others.

The brain is sensitive to the structure(s) of language, as it encodes light and sound to yield meaning, a primary aim of many language exchanges. The details of the psychological or physiological processes that extract intended meaning from sensory input are unknown and highly controversial. But it is generally agreed that one outcome is some type of (re)depiction of activities of various people, places, and things in the comprehender's mind. The so-called discourse entities must be selected. Further, their interactions need to be encoded in a dynamic mental model that supports the kinds of inferences that are required to generate new information or new construals of old information by the comprehender. To understand "Jacques who Marta commented on presupposed the existence of special purpose innate processor for language" requires keeping some temporary representation of Jacques and Marta until it becomes clear who did what to whom. Sentential structure notwithstanding, there are still many moments of syntactic ambiguity during a sentence's processing when it isn't clear who is doing and who is being done unto as in "Jacques who Marta...". Moreover, as has become all too clear from the many failed attempts at machine translation from one language to another, to understand the meaning of such a sentence, information needs to be retrieved from long term memory; language comprehension and knowledge in long term memory go hand in hand. For example, knowing that Jacques' title included the phrase "language acquisition device" frames the sentence under discussion differently than simply knowing that Jacques and Marta both gave a presentation. The point, for present purposes, is that understanding the bulk of language requires processing at levels beyond words. Processing of units larger than single words, in turn, implies temporary representations in some form of working memory.

While the nature of the units of working memory (e.g., discrete chunks, activations within a production system, across units of a neural network, an articulatory loop) is highly contentious, it is commonly agreed that working memory is capacity limited and that its "contents" are temporary (for reviews see Richardson, Engle, Hasher, Logie, Stoltzfus, & Zachs, 1996). Furthermore, there is a growing appreciation for the importance of temporal sequencing of working memory operations in separating processes and thereby on our ability to comprehend and produce language. The capacity limitations and the temporal characteristics of working memory impact language comprehension in important ways, which as yet are not fully understood. Thus, I believe we need to track both working and long term memory use during online language comprehension so as to get a better picture of how different regions in neural space are coordinated in time to yield sense from a serial linguistic stream.

In our laboratory, we have been using ERPs recorded from the scalp as people read clauses or sentences as a way to look at word processing, language structure, and the relation between the two. Most recently, we have used ERPs to examine the hierarchical relations between words independent of their actual meanings, with the eventual aim of investigating how meaning influences the processing of sentences with certain structures. We started by recording ERPs across the human scalp to simple transitive sentences like "The doctor examined the child" - where there is a subject (S), a verb (V), and an object (O). The data revealed electrophysiological properties that emerge across sentences that are more than the responses to individual words lined up end to end: namely, very slow potentials that cumulate and fluctuate across the course of the sentence.

An informative way of looking at such slow potential effects is to apply a low pass digital filter to the cross-sentence ERPs such that only the slow activity remains. The filtering simplifies the representation of the data but still leaves a temporally and spatially rich and complex pattern of activity over the head. The slow potentials vary in time across the clause as sentence constituents are processed, and in space from the front to the back of the head as well as between the hemispheres. An important characteristic of these types of simple transitive sentences is that they do not draw any special attention to themselves but do nonetheless engage orthographic, lexical, semantic, and syntactic information and processes as well as working and long term memory systems.

Let us briefly examine the pattern of slow potentials (see Figure 1). Electrodes over the lateral occipitotemporal cortex are over brain areas that are critical for the early processing of visual stimuli including words. We cannot assume that electrical activity at the scalp is generated directly beneath the electrode. But as a first working hypothesis, we suggest that the sustained occipital potential reflects neural activity involved in processing of visual features. This negativity may reflect the continuous processing of the visual input in the ventro-lateral occipital areas proposed by Petersen and Fiez (1993) to be involved in processing word forms on the basis of their PET data. If our hypothesis is correct, then we would expect to record similar slow potentials during spoken sentences; however, these potentials would have a more central distribution characteristic of early auditory processing generated in or near the superior temporal gyrus. This prediction was borne out in subsequent work, as seen in Figure 3 (Mueller, King, & Kutas, 1997).

Over temporal sites, both the fast and slow potential activity are sensitive to lexical class and even higher level features of the input. The low-pass filtered data show a phasic positivity beginning at the verb of the clause (see Figure I). Our current working hypothesis is that this phasic positivity is an index, at least in part, of the processing of thematic role information contained in the lexical representation of the verb; loosely this means the positivity is being linked to processes responsible for correctly associating a verb with its noun arguments. If our hypothesis is correct then we would expect the amplitude of the response to the verb to vary systematically with the number of thematic roles assigned by that verb. Specifically, we would expect amplitude differences for intransitive verbs ("sleep") that assign one thematic role, transitive verbs ("kill") that assign two thematic roles and ditransitive verbs ("give") that assign three thematic roles.

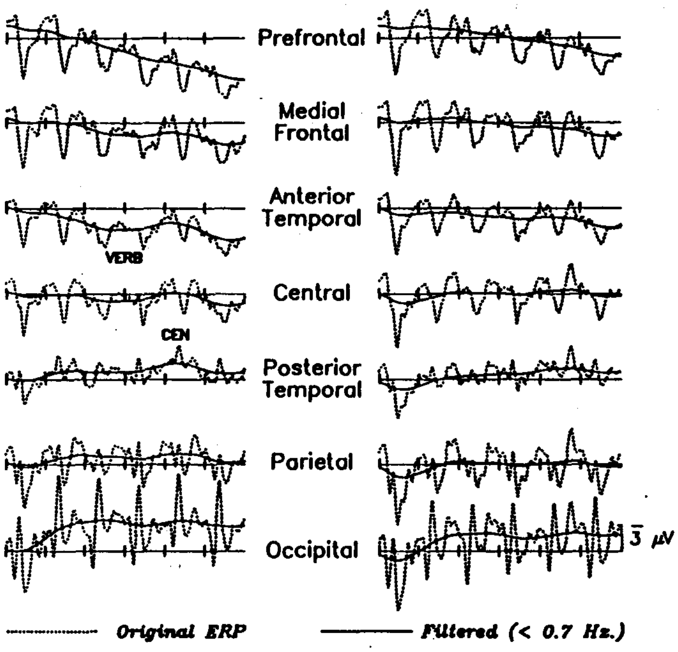

Figure 1. Superimposed are originally recorded cross-clause grand average ERPs and the slow potentials only (low pass filtered at 0.7 Hz) to simple transitive clauses for 7 pairs of sites from anterior to posterior sites across the head over left and right hemispheres separately. Negative is up on this and all subsequent figures. Words were presented one at a time every 500 ms for a duration of 200 ms. The positivity associated with verbs and the clause ending negativity (CEN) are labeled.

Also over temporal sites, following the object of the verb ("child") comes the slow CEN or clause ending negativity (see Figure 1). Clause boundaries have been shown to make very heavy demands on working memory capacity; this is evidenced in generally slowed reading times whether measured by manual responses or eye movements (e.g., Aaronson & Ferres, 1986). If we are correct about our proposed link between the phasic verb positivity and thematic role assignments and the clause ending and the CEN, we would expect to see them both during speech processing as well.

At the same time, over the front of the head there is an extremely slow (< .2 Hz), cumulative positivity. The combination of its frontal maximum and the role of the frontal lobes in executive functions of working memory (Pennington, Bennetto, McAleer, & Roberts, 1996) have led us to suggest that this ultra slow positivity may reflect integration of recently encoded information in working memory with information from long term memory. As with the verb-related positivity and CEN, we would expect this activity to be relatively independent of input modality, insofar as it is related to some aspect of language processing as opposed to reading or listening per se. On the other hand, even if such effects turn out to be modality independent, this does not necessarily mean that they are language-specific.

In brief, we have described four slow potential effects, all of which are laterally asymmetric: a sustained negative shift over occipital sites presumably related to early visual processing, a temporal positivity time-locked to verb processing (possibly thematic role assignment), a negativity that indexes clause ending, and an ultra-slow cumulative positivity over frontal sites that might reflect the processing involved in building a mental model of the sentence. These slow potential effects show that the different temporally overlapping processes that are called into play in multiple brain areas can be monitored via ERP recordings.

The pattern of these slow potential effects during the processing of simple transitive sentences is markedly different for "good" and "poorer" comprehenders. Poor and good comprehenders were grouped according to a median split on their scores to comprehension questions following a subset of the sentences. Over occipital sites, poorer comprehenders show larger and more asymmetric early visual components (P1-N1-P2) and larger and more asymmetric slow potentials. By contrast at the prefrontal sites, it is the good rather than poor compehenders who show a larger and more asymmetric slow positivity that builds over the sentences - which we linked to integration. Perhaps there is a tradeoff in resource allocation: more resources devoted to early visual processing (feature extraction, accessing wordforms) by poor comprehenders and more resources devoted to higher order integrative processes by good comprehenders.

We would expect integrative processes, in particular, to be more readily completed for simple transitives than for sentences with more complex syntactic structures that describe more complex discourse relations, such as those with relative clauses. There are however different types of relative clauses; two that we have investigated are subject and object relative sentences. In subject relative sentences, the subject of the main clause is also the subject of the relative clause (e.g., "The professor who advised the dean made a public statement."). In object relative sentences, the subject of the main clause is also the object of the relative clause verb (e.g., "The professor who the dean advised made a public statement."). These sentences are similar in that they both include a relative clause but they differ in the structure of the relative clause and in the role the main clause subject plays in the relative clause. Neither sentence type is ambiguous. However, object relatives are typically considered to be more difficult than the subject relatives (King & Just, 1991).

It has been suggested that object relatives are more difficult primarily because they put a greater load on working memory resources than do subject relatives, although both make greater demands on working memory than do simple transitive sentences. Object relatives require the reader or listener to hold an unattached constituent in working memory for a longer time than in subject relatives. This greater memory load is assumed to lead to processing difficulty (i.e., slower processing or lower accuracy). This has been inferred from reading time measures, taken as subjects press a key to present each word of a sentence one after another, one at a time, at their own pace. In reading time data, the first reliable sign of such a greater memory load occurs at the end of the relative clause ("advised"), which is just about when the load goes away. Moreover, the largest reading time effects are observed right after that at the main clause verb ("made"). Another potentially difficult aspect of object relatives is that they require multiple shifts in focus between the discourse participants. In this example, first the professor is at the center of focus, then the dean, then the professor again. In contrast, the professor remains the focus in subject relative sentences.

Electrophysiological data, at the sentence location where the largest RT effects are generally observed, namely, at the main clause verbs, reveal a much greater left anterior negativity (LAN) for the object relatives (King & Kutas, 1995). This type of effect was first described by Kluender (1991) for wh-Questions versus simple yes-no questions. In wh-Questions, it takes a while before the reader or listener knows who the who is in "Who do you think that ...?" Kluender proposed that this greater LAN for wh-Questions reflects the maintenance of items in working memory (Kluender & Kutas, 1993). An alternative view links the LAN to a disruption of the first stages of syntactic processing (e.g., Friederici & Mecklinger, 1996). Our data do not fit readily with this latter view. We have observed a LAN for both types of relative clauses compared to verbs in sentences that do not have any embedded clauses, so this relative negativity does appear to vary with working memory load. This finding meshes well with the reading time data. As an aside, note that it may well be that there is more than one negativity subsumed by the LAN, one of which may reflect working memory processes. Moreover, it follows that if the LAN reflects working memory load it should be present with non-language materials as well. Although it is difficult to compare ERPs elicited in different experimental paradigms directly, the work of Roesler and his colleagues (Roesler, Heil, & Roder, 1997) shows that increased memory load is associated with greater long-lasting negativities. The negativities are sensitive to the same sorts of manipulations but vary in their scalp maxima as a function of the stimulus materials. We consider this pattern to be non-language specific in its elicitation but language-specific in its scalp distribution.

Back to the issue at hand, one could ask whether the LAN effects at the main clause verb add any new information. Some researchers say that the presence of this LAN effect could be informative, but only if its neural generator were known. However, even without knowledge of neural generators, ERPs can provide information that reading time could not provide. By contrasting sentence length ERPs for the two relative clause sentence types, we can see that processing differences between the two object relative types occur much earlier in the sentence than reading time data suggested (Figure 2). The first sign of a reliable ERP difference occurs as soon as the working memory load is different - at the relative clause, when the reader has to deal with the word "dean" without yet knowing what to do with "professor". The difference is a sustained negativity over central and frontal sites for the object compared to the subject relatives. There is a further ERP differentiation as the reader returns to the main clause, encounters the main clause verb "made", and can in fact finally decide who did what to whom.

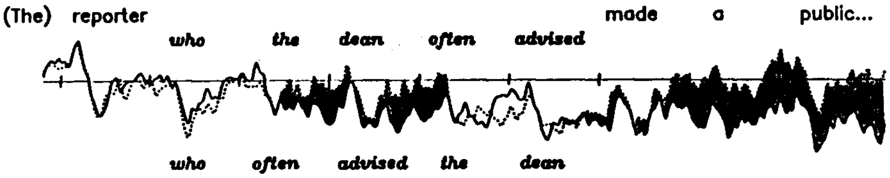

Figure 2. Comparison of the grand average cross-sentence ERPS elicited by subject relative (solid line) and object relative (dotted line) sentences recorded over a left frontal location. the two sentence types are equivalent both before and after the relative clause. The relative clause above the baseline is a sample object relative sentence and that below the baseline is a sample subject relative sentence. Words were presented one at a time every 500 ms for a duration of 200 ms. The shaded area represents the difference between the two conditions on the average; object relatives are reliably more negative than subject relatives here.

A median split of Good (87%) versus Poorer (60%) comprehenders reveals strikingly different patterns, with the good comprehenders showing a large difference between the subject and object relatives and poorer comprehenders showing almost none. In fact, the ERP data suggest that poorer comprehenders find even the subject relative sentences pretty difficult; they generate mostly negativities as if their WM were loaded down even by the embeddings in subject relatives. Moreover, the over-sentence ERPs of poorer comprehenders show a much attenuated version of the slow frontal positivity characteristic of the subject relatives in good comprehenders, suggesting that the poorer comprehenders are not integrating as regularly or as well as the good ones.

One criticism that has been leveled against this work, however, is that reading two words per second is abnormally slow, when normal reading is closer to three to four words per second. Thus, it has been suggested that the ERP differences we have observed between the relative clause sentence types were merely an artifact of our one word at a time, presentation rate. Moreover, even at normal reading rates it has been argued that "real" language is speech, not reading. But segmenting speech raises both theoretical and practical problems; it is difficult to record clear ERP components even from semi-connected speech. However, because the slow potentials we have observed are not triggered by individual word onsets, but rather by the continuous mental processes that extend across larger units such as clauses, we decided to examine the sentence ERP effects as people listened to embedded sentences.

Just as in visually-presented sentences, we observed large ERP differences between the subject and object relative sentence types (Mueller, King, & Kutas, 1996). These were remarkably similar to the visual effects, although somewhat different in their scalp distribution. The auditory effects were more widespread on the scalp, and much larger over the right hemisphere than the visual effects. However, at some electrode sites, such as left frontal regions, the effects during the relative clause are remarkably similar in the two modalities. We view the difference in negativity for subject and object relative sentences for both written and spoken language as a modality-general effect. In other words, we view the presence of a negative difference with a similar timecourse and sensitivity as an index of language processing in general, rather than a process specific to either reading or listening.

As expected, however, there are other aspects of the ERP that are modality specific. The slow potentials reveal that visual and auditory sentences generate dynamically similar but topographically distinct effects. We believe that the negativities across the course of sentences in the visual and auditory modalities are similar; however, since the stimulus input features are different, the particular brain areas involved in their processing are also different to some extent. As a consequence, the relative distributions of potentials across the scalp also differ. For example, in the visual modality, the resting negativity is large at a (left) occipitotemporal site, but absent at a site on top of the head. In the auditory modality, the opposite pattern is seen (Figure 3). We have previously argued that the visual negativity is related to sustained word processing; we make a comparable argument for the auditory data. We also found that the processing differences connected with comprehension skill as reflected in the ERP patterns are similar for reading and listening.

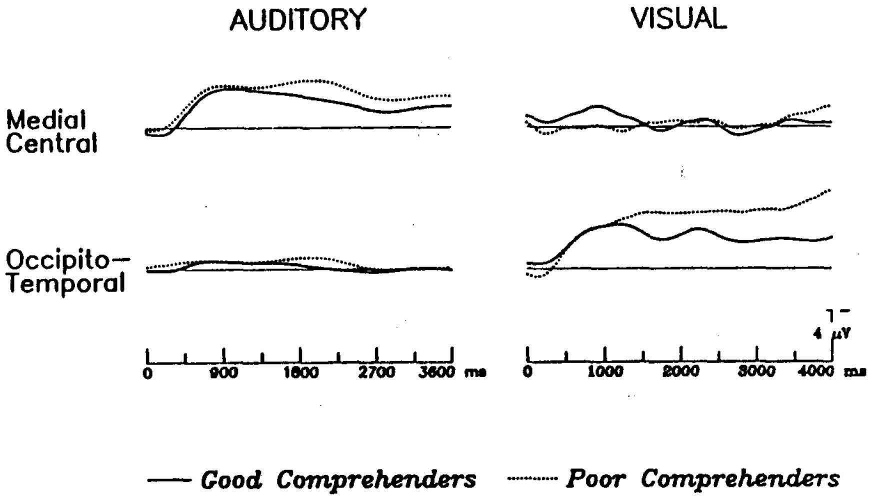

Figure 3. Grand average slow potential ERPS (low pass filtered at 0.7 Hz) to sentences of various structures for good (solid line) and poor (dotted line) comprehenders in the visual and auditory modalities at one medial central and one occipito-temporal recording site. In the auditory modality, the prolonged negative potential over the clause is pronounced over central and occipital sites whereas in the visual modality it is maximal over occipito-temporal sites. Note that the reading (King & Kutas 1995) and listening data (Mueller et al., 1997) are from two different groups of subjects.

With such data in hand, we have begun to examine what types of information (e.g., lexical, semantic, discourse, pragmatic) the brain is sensitive to while processing object relatives. In this way we can begin to understand why some embedded sentences like "The student that the committee that the results surprised summoned listened closely" seem to be much harder to understand than others like "The painting that the artist that everyone loves painted hung in the living room" although they have exactly the same syntactic structure with the same number of embeddings. It is our belief that the difference must rest in the content of the words, in other words, in their meaning. Theories of parsing differ considerably on when meaning is allowed to play any role and exactly what role it is allowed to play. Jill Weckerly, Jeff Elman, and I have pursued this by doing connectionist modeling, reading time studies and ERP studies (see Weckerly & Elman, 1992). All three approaches indicate that a noun's animacy does influence sentence parsing, perhaps by supporting tentative estimates on likelihood of the roles that a noun might play even before any verbs are encountered (Weckerly, 1995).

Thus far I have described how ERP data are used to study language processing in real time. I would like to conclude by relating my comments and data to Professor Mehler's position. Mehler wrote that "language has to be studied as an endowment that is genetically determined and culturally clinched". In other words, that language is innate but influenced by cultural input. If from this we are to understand that language is an inborn capacity that is not learned, then it is not particularly useful since almost all behavior is subject to some kind of learning. If on the other hand we are to understand by this that language is innately specified in the genome, then it is a bit misleading in that every phenotypic feature and hence also every behavioral trait is codetermined by an individual's genome and by environmental variables past and present. PET evidence was offered as "additional data that may make it easier to understand our language processing device (LAD)". I cannot see how PET data, per se, regardless of the task can ever be evidence for a language acquisition device given that the participants in these experiments were molded both by their genetic endowment and their experiences. The existence of a language acquisition device is a hypothesis but in this type of presentation, it seems to be a given. Presumably, this language processing device is a cortical structure that is specific to processing language and no other cognitive function, and thus studying the brain seems a reasonable approach. However, the LAD is not presented as a hypothesis but rather mentioned as if there were already evidence for its existence, although in fact there is no agreement as to what data would constitute evidence for it. Moreover, the brain data are presented as relevant to issues in language acquisition. But, this link hinges on the assumption that language is innately specified, from which it would then follow that this innate specification in the genome would be manifest in cortical organization. But again there is no consensus that language is innately specified. Finally, Professor Mehler presented neural imaging data from bilinguals as a means of determining how their first and second languages compete for this cortical structure (i.e., the LAD); again, this approach is based on the presupposition that such a cortical structure exists. The neuroimaging data do show, not surprisingly, that somewhat different regions of the brain are involved in first and second languages. What they do not show, however, is that there is an LAD in monolinguals or bilinguals; nor could they ever.

By contrast, our research does not start with the assertion that language is innate or that there is a language acquisition device. Rather we take the view that there may be a relatively direct relationship between processes of language comprehension and general cognitive processes, and that linguistic data can suggest constraints on memory access and organization of background knowledge. Work in my laboratory as well as my reading of the literature lead me to conclude that there are multiple areas of the brain involved in language processing. I have shown that the ERP methodology is one technique for temporally tracking sentence comprehension and that so doing reveals that working memory operations and background knowledge are critical to normal comprehension. Our data show that many of the processes of language include large parts of the brain, are neither modality nor language-specific, are quite sensitive to individual differences, and overlap in space and time, such that both dimensions are needed to tease them apart.

ACKNOWLEDGEMENTS

Thanks to M. Besson and C. van Petten for comments on a previous draft of this manuscript and J. Requin for many discussions throughout the years. Some of the work reported herein was supported by grants from HD22614, MH52893, and AG08313 to M. Kutas.

Les aires de traitement du langage sont largement, mais non aléatoirement, distribuées dans le cerveau. Etablir une correspondence entre ces aires cérébrales et les fonetions langagières, en temps reel, nécessite non seulement de savoir comment ces aires sont représentées dans l'espace mais également de comprendre comme elles sont coordonnées dans le temps. Toute théorie psychologique du traitement du langage doit include une description des representations mentales, des operations qui sons réalisées sur ces representations, de leur décours temporel et de

leurs interactions. Une méthode de choix pour cerner précisément le décours temporel des processus impliqués dans la comprehension des phrases est la technique des potentiels évoqués (ou potentiels lies a 1'événement). Son utilisation dans l'analyse de phrases écrites ou parlees, dont la complexité structurale varie, révèle les roles prédominants joués par la mémoire de travail et les connaissances implicites dans le processus normal de comprehension. Les données montrent que les processus de traitement du langage requièrent de larges zones cerebrates, ne sont spécifiques ni de la modalité sensorielle, ni du langage, sont soumis a des différences interindividuelles et se recouvrent dans Fespace, neural et mental, et dans le temps, de telle manière que des analyses sur ces deux dimensions sont requises pour les différencier et comprendre comment nous comprenons.

REFERENCES

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

Binder, J. R., Frost, J. A., Hammeke, T. A., Cox, R. W., Rao, S. M., & Prieto, T. (1997). Human brain language areas identified by functional magnetic resonance imaging. *Journal of Neuroscience,* 17(1), 353-362.

Chomsky, N. (1966). Review of Skinner's "Verbal behavior". *Language,* 35, 26-58.

Coulson, S. (1997). Semantic leaps: the role of frame-shifting and conceptual blending in meaning construction. Ph.D. dissertation. University of California, San Diego, La Jolla, CA.

Damasio, A. R. (1997). Brain and language: what a difference a decade makes. *Current Opinion in Neurology,* 10, 177-178.

Deacon, T. W. (1996). Prefrontal cortex and symbol learning: Why a brain capable of language evolved only once. In B. Mitrofanovich Velichkovsky & D. M. Rumbaugh (Eds.), *Communicating meaning: The evolution and development of language (pp.* 103-138). Mahwah, NJ: Lawrence Erlbaum Associates.

Friederici, A. D., & Mecklinger, A. (1996). Syntactic parsing as revealed by brain responses: First-pass and second-pass parsing processes. *Journal of Psycholinguistic Research,* 25 (l), 157-176.

Gratton, G., Fabiani, M., & Corballis P. M. (1997). Can we measure correlates of neuronal activity with non-invasive optical methods? *Advances in Experimental Medicine and Biology, 413,* 53-62.

Joanette, Y., & Brownell, H. H. (Eds) (1990). *Discourse ability and brain damage: theoretical and empirical perspectives.* New York: Springer-Verlag.

Joanette, Y., Goulet, P., & Hannequin, D. (1990). *Right hemisphere and verbal communication.* New York: Springer-Verlag.

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

King, J., & Kutas, M. (1995). 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. (1991). Cognitive constraints on variables in syntax. Ph.D. dissertation. University of California, San Diego, La Jolla, CA.

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

Kutas, M., & Federmeier, K. (1998). Minding the body. *Psychophysiology,* 35, 1-16.

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

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.

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

Pennington, B. F., Bennetto, L., McAleer, O., & Roberts, R. J. Jr. (1996). Executive functions and working memory: Theoretical and measurement issues. In G. R. Lyon & N. A. Krasnegor (Eds.), *Attention, memory, and executive function (pp.* 327-348). Baltimore, MD: Paul H. Brookes Publishing Co.

Petersen, S. E., & Fiez, J. A. (1993). The processing of single words studied with positron emission tomography. *Annual Review of Neuroscience, 16,* 509-530.

Richardson, J. T. E., Engle, R. W., Hasher, L., Logie, R. H., Stoltzfus, E. R., & Zachs, R. T. (1996), *Working memory and human cognition.* New York: Oxford University Press.

Roesler, F., Heil, M., & Roder, B. (1997). Slow negative brain potentials as reflections of specific modular resources of cognition. *Biological Psychology, 45 (1-3),* 109-141.

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 *Proceedings of the 14th annual conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Received 27 November, 1997
Accepted 24 April, 1998