12

Watching the Brain Comprehend Discourse

Morton Ann Gernsbacher and David A. Robertson

When the first author (Morton Ann Gernsbacher) was in graduate school, she once sat in on an undergraduate entry-level cognition course to see how one of the department's best teachers conveyed the excitement and exquisiteness of cognitive psychology to an audience of 19-year-olds. This professor began the first day by posing the following challenge: Imagine that you were sent to some faraway planet, and your mission on this planet was to discern the workings of a mysterious—at least to you—structure, what here on earth we might call a building. However, you were prohibited or otherwise unable to enter the structure. How would you discern what work went on inside the structure?

With guidance, the students arrived at recommendations such as to first carefully observe what entered the structure, and then carefully observe what exited the structure, and from those observations infer what work must go on inside the structure. Thus, if one saw sheet metal, rubber, and glass entering the structure and one observed Subaru vehicles leaving the structure, one might infer something different than if one saw denim cloth, thread, and zippers entering the structure and men's blue jeans coming out.

At this point the professor skillfully introduced some of the rudimentary concepts of experimental design: How, if clever enough, one might manipulate certain aspects of the input to the structure while controlling as many extraneous variables as possible, then measure qualitatively or quantitatively the output and thereby allow sharper inferences of what type of work went on inside the structure. At this point students suggested other covert techniques, for example, sneaking in at night or some time when their presence would not be detected—in other words, observing the structure "at rest." Trying to obtain a blueprint or floor plan of the structure was also suggested. Even bombing a section of the structure to see how that affects the output (an idea not so far

---

The research described in this chapter was supported primarily by a grant to Morton Ann Gernsbacher from the National Institute of Health (RO1 NS29926), a University of Wisconsin–Madison Faculty Development Award to Morton Ann Gernsbacher, and a National Institute of Mental Health predoctoral traineeship (T32 MH18931) to David A. Robertson.

removed from our U.S. military reconnaissance) was proposed. These approaches resemble techniques of neuroanatomical inquiry: postmortem analyses, nonfunctional imaging, lesion studies. However, as the professor pointed out, none of those approaches would portray the structure at work. And indeed, over 20 years ago, when the first author sat in on this classroom activity, the opportunity to see the brain at work eluded cognitive psychologists.

Times have changed. Now, the ability to answer questions such as how humans plan, learn, remember, represent experience, and comprehend language is aided by the ability to watch the brain at work (Posner & Raichle, 1994). We can generate visual images of that mysterious structure while these higher level processes are occurring. In this chapter, we provide (a) a claim—that a picture can be worth a thousand milliseconds; (b) a testimonial—that one can find hay in a haystack; and (c) a caveat—from a lesson learned from drinking scotch.

A Claim

The claim we want to make is that "A picture can be worth a thousand milliseconds," which could be subtitled "Why doing brain-imaging experiments can be more than redoing old reaction time experiments with more expensive apparatus." We substantiate this claim with some of our own research. The brain imaging technique we have used most frequently is functional magnetic resonance imaging (fMRI), and the goal of many of our fMRI experiments has been to observe the brain at work while it processed coherent discourse. In one experiment (Robertson et al., 2000), we manipulated a subtle marker of discourse coherence: the English definite article *the*. In languages that use an article system, the definite article signals repeated reference, which typically leads to coherence. For example, the use of the definite article *the* in these two sentences,

> *The conference speaker was talking very fast.*
> *The conference speaker was showing a bunch of overheads.*

suggests that the woman who was showing a lot of overheads was the same as the woman who was talking very fast. In contrast, the use of the indefinite article *a* in these two sentences,

> *A member of the audience was listening.*
> *A member of the audience was yawning.*

makes it unclear whether the member of the audience who was listening was also the member of the audience who was yawning. However, if we substitute the definite article *the* for the indefinite article *a,* this unfortunate situation becomes clear.

> *A member of the audience was listening.*
> *The member of the audience was yawning.*

In fact, the definite article *the* can signal coreference even when the noun it modifies is only a synonym of the previously mentioned noun, as in

> *A member of the audience was getting bored.*
> *The clod began thumbing through his program looking for another talk to attend.*

As these examples illustrate, the definite article *the* promotes coherence.

Gernsbacher and Robertson (2002) collected empirical data to evaluate these intuitions. We created 10 series of sentences. One version of each series contained only the definite article, and the other version contained only indefinite articles. In one experiment we measured subjects' reading time and their recall, and in another experiment we measured subjects' reading time and their speeded recognition, using a priming-in-item recognition task (McKoon & Ratcliff, 1980). We observed that when the sentences were presented with the definite article, they were read considerably more rapidly, about 25% faster (as Haviland & Clark, 1974, had observed over 25 years earlier); they were perceived as being more narrativelike (as de Villiers, 1974, had observed over 25 years earlier), they were recalled in a more integrative fashion, and they showed a priming-in-item recognition advantage, meaning that recognition of one sentence from a series was speeded if preceded by another sentence from that series. Sentences presented with indefinite articles did not show this priming advantage.

All these data led us to believe that the definite article *the* signals the recurrence of concepts and therefore their interrelations, and that discourse comprehension must involve capturing those interrelations. However, those were just our educated speculations from carefully controlling what we sent into the unknown structure and carefully observing what came out. Therefore, we conducted an fMRI experiment in which we presented several series of sentences that contained only indefinite articles, such as

> *A grandmother sat at a table.*
> *A young child played in a backyard.*
> *A mother talked on a telephone.*
> *A husband drove a tractor.*
> *A grandchild walked up to a door.*
> *A little boy pouted and acted bored.*
> *A grandmother promised to bake cookies.*
> *A wife looked out at a field.*
> *A family was worried about some crops.*

Or we presented the same sentences, with the indefinite articles replaced by the definite article, as in the following examples:

> *The grandmother sat at the table.*
> *The young child played in the backyard.*
> *The mother talked on the telephone.*
> *The husband drove the tractor.*
> *The grandchild walked up to the door.*

> *The little boy pouted and acted bored.*
> *The grandmother promised to bake cookies.*
> *The wife looked out at the field.*
> *The family was worried about the crops.*

Although as psycholinguists we wanted to directly compare these two conditions, the Wisconsin medical physicists with whom we consulted could not believe that such a “subtle” comparison would lead to any observable differences in brain activity. Given their skepticism and to conform to the fMRI literature of that time, we also included a “loose” control condition, in which we replaced the letters of the sentences with nonletter characters while retaining interword spacing and string length (e.g., <`<^#%} |<-|*))#~/ <>*{+-*^~?)*(-.). This combination of tight and loose comparisons turned out to be very fruitful. We collected whole brain functional images in 23 coronal slices from 8 right-handed subjects and observed two main findings.

First, we observed that reading sentences compared with viewing nonletter strings led to a robust region of activation in the left hemisphere extending from the angular gyrus rostrally to the left anterior temporal pole along the middle temporal gyrus; a smaller region of activation was also observed in the right-hemisphere homologue of Wernicke’s area, as shown in Figure 12.1. These data replicate several other sets in the literature, including those of Bavelier et al. (1997), who examined the contrast of reading sentences versus consonant strings. We were encouraged by the very similar results across laboratories despite different tasks and different analysis techniques.

However, in contrast to our comparison between reading sentences and viewing nonletter strings, our comparison between reading sentences that contained definite articles and reading sentences that contained indefinite articles revealed virtually no differences in left-hemisphere activation. Instead,

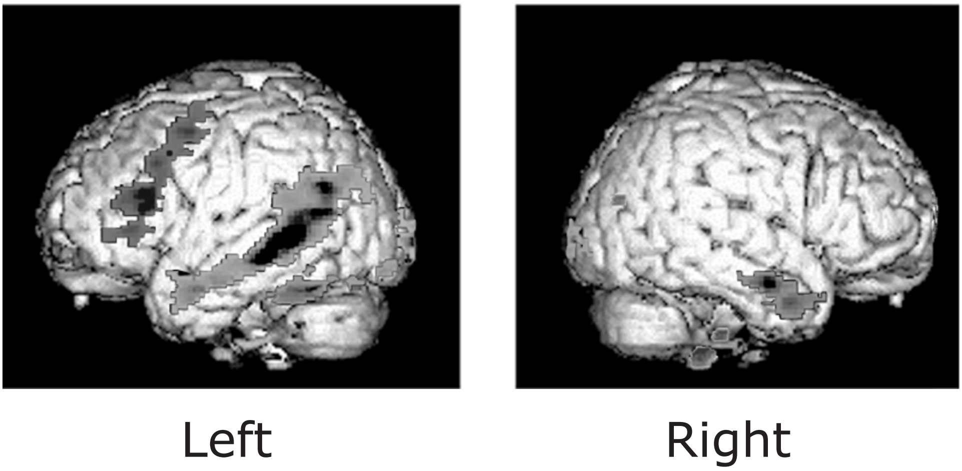

Figure 12.1. Brain activation patterns from Robertson et al. (2000). Illustrated is the contrast between subjects reading sentences (left) and viewing nonletter character strings (right); data are averaged across all subjects.

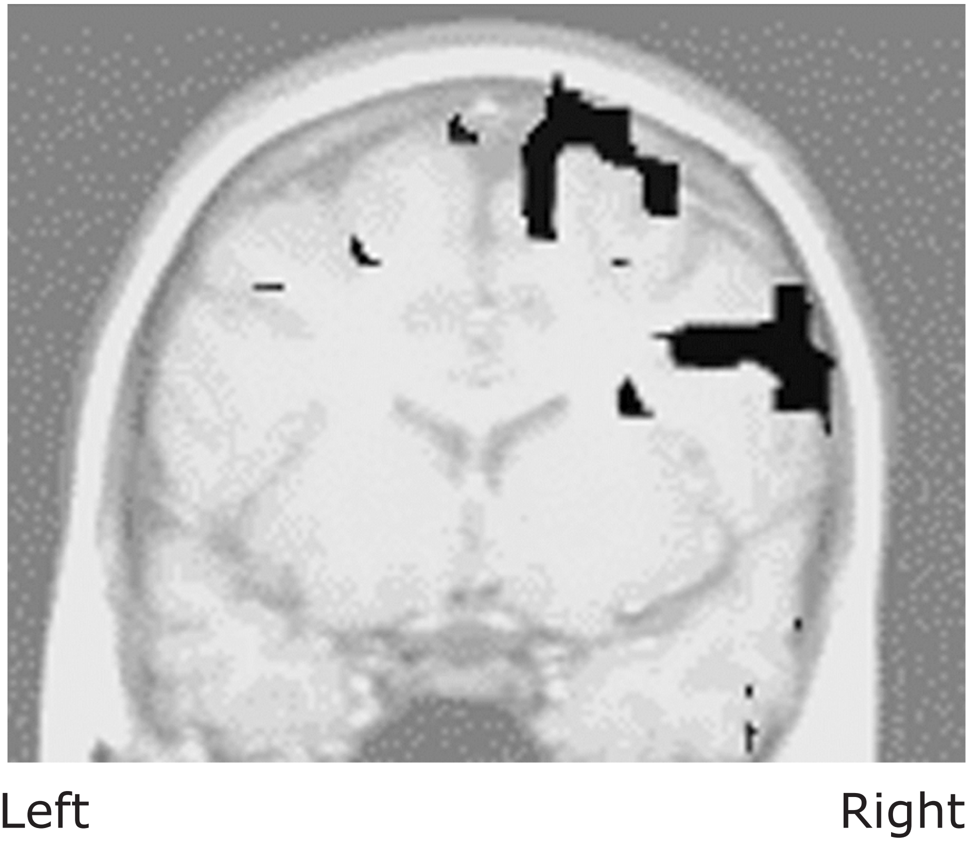

Figure 12.2. Brain activation pattern for one example subject in Robertson et al. (2000). Illustrated is the contrast between reading sentences that contained the definite article (*the* [left]) and sentences that contained indefinite articles (*a, an, some* [right]).

differential activation was observed in frontal regions, particularly in the right superior and medial frontal gyri. Figure 12.2 illustrates the pronouncedly right-lateralized activation in one of our subjects, from a coronal view through the frontal lobes and anterior temporal lobes. This finding corroborates a wealth of lesion studies, which demonstrate that people with right-hemisphere damage are often described as "not getting the point," less able to integrate information, and unaware of inconsistencies; in short, they are described as often oblivious to and challenged by discourse coherence (Brownell, Carroll, Rehak, & Wingfield, 1992; van Lancker & Kempler, 1987; Winner & Gardner, 1977; Zaidel, Zaidel, Oxbury, & Oxbury, 1995).

Thus, our finding that the processing of coherent discourse is manifested by more right- rather than left-hemisphere activation corroborates the behavioral challenges that people with right-hemisphere brain damage exhibit during discourse comprehension. We can then turn to other functional imaging studies to ask: Why right frontal? Other functional imaging work suggests that right frontal regions support the allocation of attention (e.g., Knight & Grabowesky, 1995); we have suggested that the definite article acts as a cue for coherence. Therefore, the behavioral advantages that we previously observed might be advantages of attention to that cue. Could we have arrived at the same conclusions through reaction time experiments? Granted, this experiment was not as unique as Blakemore, Wolpert, and Frith's (2000) study in which they collected fMRI data to explain why tickling yourself is never as ticklish as someone else tickling you. Nonetheless, we claim that watching the brain at work provides insights that reaction time studies cannot.

A Testimonial

Our testimonial is that one can look for hay in a haystack, or perhaps more accurately, one can find hay in a haystack. We often think of functional imaging as a wonderful vehicle for isolating unique neural circuits, distinguishing independent systems, and otherwise identifying differences, but one can also use the technique to identify commonalities. We illustrate this point by sharing another research study from our lab.

In much of our behavioral work, we have been interested in identifying the common processes involved in comprehending different media (Gernsbacher, 1990, 1997). This is not to deny that there are language-specific processes; rather our goal has been to identify the processes and mechanisms that might be common across the comprehension of various media. We have known for almost 30 years that for adults without any known impairments, written comprehension, measured by performance on a comprehension test after reading, is highly correlated with spoken language comprehension, measured by performance on a comprehension test after listening (Daneman & Carpenter, 1983; Jackson & McClelland, 1979; Palmer, MacLeod, Hunt, & Davidson, 1985; Perfetti & Lesgold, 1977; Sticht, 1972).

For example, Palmer et al. (1985) reported that the correlation between written language comprehension, measured by visually presenting the comprehension sections of the Davis reading test, and spoken language comprehension, measured by auditorily presenting the comprehension sections of the Davis reading test, was .80. The correlations were only slightly lower for written comprehension measured by the Nelson–Denny or the Washington Pre-College Scholastic Aptitude test and spoken comprehension, measured by auditorily presenting the Davis reading test. Over a decade ago, my lab extended these findings by demonstrating that comprehension of written narratives was highly correlated with comprehension of spoken narratives, and more strikingly, that both were highly correlated with comprehension of narratives told without any words, that is, picture-only narratives. We (Gernsbacher, Varner, & Faust, 1990) reported correlations corrected for reliability of .92, .82, and .72 for the relation between written and spoken, written and picture, and picture and spoken comprehension, based on our Multi-Media Comprehension Battery.

These correlations support the hypothesis that many of the cognitive processes and mechanisms underlying discourse comprehension are general enough to be involved in the comprehension of nonverbal media. More discrete laboratory behaviors also support this hypothesis. For example, after viewing a narrative as a movie without dialogue or listening to the narrative as text, comprehenders mark off the same episode structure (Baggett, 1979; Gernsbacher, 1985); comprehenders draw the same inferences after they view nonverbal cartoon sequences as when they read verbal descriptions of those sequences (Baggett, 1975); and when recalling a narrative viewed as a movie without dialog or listened to as a text, comprehenders emphasize, elaborate, and omit the same information (Poulsen, Kintsch, Kintsch, & Premack, 1979).

Therefore, in a recent fMRI experiment (Robertson, Guidotti, & Gernsbacher, 2003) we measured subjects' brain activity while we presented six narratives. Two were presented by means of written sentences, two were presented by means of spoken sentences, and two were presented by means of line-drawn pictures. The subjects' comprehension task was oddball detection. During each narrative segment, which comprised six or seven narrative stimuli—that is, six or seven written sentences, six or seven spoken sentences, or six or seven pictures—one or two oddball stimuli occurred. The subjects' task was to detect these oddballs. For the narratives presented by means of written sentences, these oddball stimuli were sentences that were semantically and syntactically coherent but they did not fit the content of the ongoing narrative. For the narratives presented by means of pictures, oddball stimuli were pictures that were taken from other picture books and therefore did not fit the ongoing narrative. We used as a comparison task the comprehension of unconnected stimuli: By this, we mean written sentences, spoken sentences, or pictures that were unrelated to each other. Oddballs for the comparison task were sentences or pictures that contained a local anomaly, such as the Kutas and Hillyard (1980) sentences (e.g., *He drank his coffee with milk and socks*) or the Biederman, Glass, and Stacy (1973) pictures, in which a common object is misplaced in a familiar scene (e.g., a fire hydrant in a living room scene). We collected whole brain functional images in 23 coronal slices from nine right-handed subjects and observed the activation patterns illustrated in Figure 12.3. The striking similarity among the images representing the brain activity while comprehending picture narratives, written narratives, and spoken narratives suggests that there is a lot of hay in that haystack, but it is not all hay.

A Caution

The penultimate point we want to make briefly draws back to the old joke about the man who one night drank a lot of scotch and water, and the next morning he had a terrible hangover; the next night he drank a lot of whiskey and water, and the next morning he had a terrible hangover; the third night he drank a lot of bourbon and water, and the next morning he had a terrible hangover. He therefore decided that he just had to stop drinking so much water. Those of us who are corny enough to tell this riddle in experimental methods courses do so in the service of illustrating the importance of one's control condition. When interpreting functional imaging experiments, it is just as crucial to ask what the control or comparison task is as it is when interpreting more traditional cognitive psychology experiments.

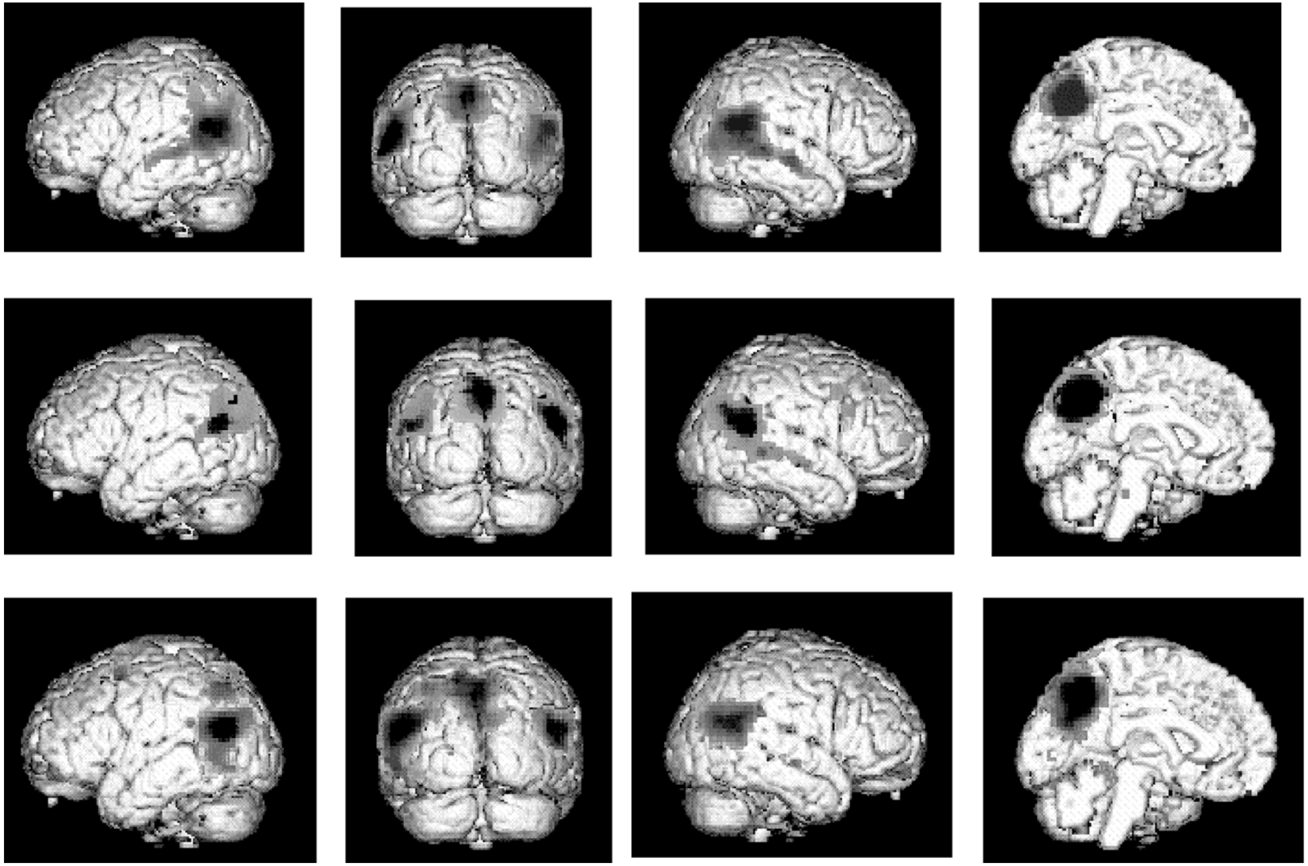

Figure 12.3. Brain activation patterns from Robertson et al. (2003). Top row: data from the picture comprehension task; middle row: data from the written comprehension task; bottom row: data from the spoken comprehension task. The first picture of each row shows a left view, the second picture a posterior view, the third picture a right view, and the fourth picture a medial view. All views are contrasts between comprehending connected stimuli (picture stories, written stories, spoken stories) and comprehending unconnected stimuli. Data are averaged across all subjects.

As we look forward to the future of cognitive psychology, we are inspired to wonder how the field will be affected by the advent of functional imaging. One effect is that we are going to see a lot more color: lots of reds, yellows, and oranges, if those remain our colors for symbolizing the brain at work. To interpret those vivid images of the brain at work, we need to ask the same questions that we have asked since the dawn of experimental psychology, including what was the comparison task? Indeed, this one caution exemplifies our strong belief that good functional brain imaging cannot be done without good cognitive psychology, both methodologically and theoretically, and we predict that soon we shall see that the opposite will be the case as well. Good cognitive psychology will need good functional brain imaging for generating and testing hypotheses.

The Applications

Last, we bring ourselves to discuss briefly the real-world applications of the findings we have reviewed here. Both our empirical data and our methodology suggest direct applications. First, we have argued here and elsewhere for the construct of general comprehension; our neuroimaging data strongly support the hypothesis that there are common brain regions underlying the comprehension of connected discourse—regardless of the medium in which the discourse occurred. Thus, instructional materials and assessments of instructional accomplishment should take advantage of the common mechanisms underlying comprehension of different media. For instance, instruction materials can indeed be multimedia without too great of a concern (particularly in a normative population) that one medium will be comprehended more poorly or differentially than another. Instructional assessment can even be cross-modality as a way to assess a general or "higher level" understanding that is independent of input modality.

Second, our methodology suggests another application, that of using brain imaging for assessing comprehension, development of comprehension, or even impairment of comprehension (Eden & Moats, 2002). Within the more prescribed realm of reading comprehension, we are already seeing brain imaging used this way. For instance, brain imaging has been used to differentiate readers who were described behaviorally as "poor readers as children who retained persistent reading problems in adulthood" as opposed to "poor readers as children who compensated as adults" (Shaywitz et al., 2003). Furthermore, brain imaging has documented changes in neural activity among poor readers who underwent successful behavioral intervention and those who did not (Temple et al., 2003). Thus, cognitive psychologists' current ability to watch the brain at work means that they can also watch the brain work better. These are quite exciting times.

References

Baggett, P. (1975). Memory for explicit and implicit information in picture stories. *Journal of Verbal Learning and Verbal Behavior, 14,* 538–548.

Baggett, P. (1979). Structurally equivalent stories in movie and text and the effect of the medium on recall. *Journal of Verbal Learning and Verbal Behavior, 18,* 333–356.

Bavelier, D., Corina, D., Jezzard, P., Padmanabhan, S., Clark, V. P., Karni, A., et al. (1997). Sentence reading: A functional MRI study at 4 Tesla. *Journal of Cognitive Neuroscience, 9,* 664–686.

Biederman, I., Glass, A. L., & Stacy, E. W., Jr. (1973). Searching for objects in real world scenes. *Journal of Experimental Psychology, 97,* 22–27.

Blakemore, S. J., Wolpert, D., & Frith, C. (2000). Why can't you tickle yourself? *Neuroreport, 3,* R11–R16.

Brownell, H. H., Carroll, J. J., Rehak, A., & Wingfield, A. (1992). The use of pronoun anaphora and speaker mood in the interpretation of conversational utterances by right hemisphere-damaged patients. *Brain and Language, 43,* 121–147.

Daneman, M., & Carpenter, P. A. (1983). Individual differences in integrating information between and within sentences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 561–585.

de Villiers, P. A. (1974). Imagery and theme in recall of connected discourse. *Journal of Experimental Psychology*, *103*, 263–268.

Eden, G. F., & Moats, L. (2002). The role of neuroscience in the remediation of students with dyslexia. *Nature Neuroscience, 5S,* 1080–1084.

Gernsbacher, M. A. (1985). Surface information loss in comprehension. *Cognitive Psychology, 17,* 324–363.

Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1997). Two decades of structure building. *Discourse Processes, 23,* 265–304.

Gernsbacher, M. A., & Robertson, R. R. W. (2002). The definite article "the" as a cue to map thematic information. In M. M. Louwerse & W. Van Peer (Eds.), *Thematics: Interdisciplinary studies* (pp. 119–136). Philadelphia: John Benjamins.

Gernsbacher, M. A., Varner, K. R., & Faust, M. (1990). Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16,* 430–445.

Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*, *13*, 512–521.

Jackson, M. D., & McClelland, J. L. (1979). Processing determinants of reading speech. *Journal of Experimental Psychology: General, 108,* 151–181.

Knight, R. T., & Grabowesky, M. (1995). Escape from linear time: Prefrontal cortex and conscious experience. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1357–1371). Cambridge, MA: MIT Press.

Kutas, M., & Hillyard, S. A. (1980, January 11). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science, 207,* 203–205.

McKoon, G., & Ratcliff, R. (1980). Priming in item recognition: The organization of propositions in memory for text. *Journal of Verbal Learning and Verbal Behavior*, *19*, 369–386.

Palmer, J., MacLeod, C. M., Hunt, E., & Davidson, J. E. (1985). Information processing correlates of reading. *Journal of Memory and Language, 24,* 59–88.

Perfetti, C. A., & Lesgold, A. L. (1977). Discourse comprehension and sources of individual differences. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension* (pp. 141–183). Hillsdale, NJ: Erlbaum.

Posner, M. I., & Raichle, M. E. (1994). *Images of mind.* New York: Freeman.

Poulsen, D., Kintsch, E., Kintsch, W., & Premack, D. (1979). Children's comprehension and memory for stories. *Journal of Experimental Child Psychology*, *28*, 379–403.

Robertson, D. A., Gernsbacher, M. A., Guidotti, S. J., Robertson, R. R. W., Irwin, W., Mock, B. J., & Campana, M. E. (2000). Functional neuroanatomy of the cognitive process of mapping during discourse comprehension. *Psychological Science, 11,* 255–260.

Robertson, D. A, Guidotti, S. J., & Gernsbacher, M. A. (2003). *A modality-independent network of brain regions supporting narrative comprehension*. Manuscript submitted for publication.

Shaywitz, S. E., Shaywitz, B. A., Fulbright, R. K., Skudlarski, P., Mencle, W. E., Constable, R. T., et al. (2003). Neural systems for compensation and persistence: Young adult outcome of childhood reading disability. *Biological Psychiatry, 54,* 25–33.

Sticht, T. G. (1972). Learning by listening. In R. O. Freedle & J. B. Carroll (Eds.), *Language comprehension and the acquisition of knowledge* (pp. 285–314). Washington, DC: Winston.

Temple, E., Deutsch, G. K., Poldrack, R. A., Miller, S. L., Tallal, P., Merzenich, M. M., & Gabrieli, J. D. E. (2003). Neural deficits in children with dyslexia ameliorated by behavioral remediation: Evidence from functional MRI. *Proceedings of the National Academy of Sciences, 100,* 2860–2865.

van Lancker, D. R., & Kempler, D. (1987). Comprehension of familiar phrases by left- but not right-hemisphere damaged patients. *Brain and Language, 32,* 256–277.

Winner, E., & Gardner, H. (1977). The comprehension of metaphor in brain-damaged patients. *Brain, 100,* 717–729.

Zaidel, D. W., Zaidel, E., Oxbury, S. M., & Oxbury, J. M. (1995). The interpretation of sentence ambiguity in patients with unilateral focal brain surgery. *Brain and Language, 51,* 458–468.